

Malware Ecologies:

A Politics of Cybersecurity



Submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Andrew Carl Dwyer

Mansfield College

University of Oxford

Trinity 2019

Abstract

Computation, in popular imaginations, is at perennial risk of *infection* from the *tools* of nefarious hackers, commonly referred to as malware. Today, malware pervade and perform a crucial and constitutive role in the insecurities of contemporary life from financial transactions, to ‘critical national infrastructures’ – such as electricity, water, and transportation – to devices in our ‘smart’ homes and cities, and even to potential ‘cyberwar.’ Yet, critical security research has rarely turned its attention to malware. In contrast, I explore malware and their politics, situated and extended beyond, an (auto)ethnographic study of the malware analysis laboratory of the UK endpoint protection business, Sophos. I argue that malware are currently processed through a *patbo*-logic that conflate organic and non-organic materialities, permitting analogies between biology and computation, and are generative of particular forms of security that relegate malware to the intent of their authors. I explore how endpoint protection businesses are imbued with these logics in order to attend to how malware are analysed, detected, and curated beyond them. By drawing on my method of ‘becoming-analyst,’ I critically reflect on how malware become known, are responded to by *ad hoc* political groups, and can assist in rethinking the role of computational agency in geography, international relations, security studies, and beyond. I instead conceive of malware as performative *political* actors making limited choices in broader computational ecologies. I therefore advocate for an *eco*-logical repositioning of malware, where cyberspace is not simply a neutral domain; but is central to the formation of choice that gives space for malware to be political. With four cases – *Conficker*, *Stuxnet*, *the Dukes*, and *WannaCry / (Not)Petya* – I write new stories on how malware are encountered and dealt with in the twenty-first century. In doing so, I challenge contemporary discourses of cybersecurity to ask if conventional notions of who and what (per)form security are adequate, and how these are reconfigured through a radical ‘more-than-human’ politics, where malware are not just objects of security, but are active participants in its production and negotiation.

Acknowledgements

There are innumerable people to thank over the course of (almost!) five years traversing new (inter)disciplinary spaces. I have had the privilege of encountering so many people, things, and spaces I thought I would never do. Inevitably there are omissions, so I start with that recognition. Various awkward moments, materialities, affects, and *politics* have informed and shaped how I have perceived, felt, and negotiated the world – so if I’ve had a coffee or too many beers with you, met through Twitter interactions, or you just listened to malformed thoughts – you’re all crucial. In the most perverse way, I thank the multiple malware forms that have been such a crucial interlocutor that are so nebulous, performative, and plural. Yet, for an empirical piece of work – thank you Sophos, you were incredibly supportive of a social scientist entering your space and dealing with me as I learnt and often failed to ‘do’ malware analysis and detection. I couldn’t have done this without you. Nor could I have done this without both of my examiners, Prof. Claudia Aradau (King’s College London) and Prof. Tim Schwanen who have given me new directions and thoughts to move beyond this work.

At Oxford, it was Andrew Martin and the CDT in Cyber Security (a fantastic place with lots of wonderful people) that gave me the opportunity to do cybersecurity with incredible support from David Hobbs, Maureen York, and Katherine Fletcher. I made wonderful friends here – Jan Silomon, Richard Baker, Kris Wilson, Alastair Janse van Rensburg, and Louise Axon – and critical perspectives from politics and international relations – Lucas Kello, Monika Kaminska, Jamie Collier, Florian Egloff, and James Shires. Yet none are more influential than Beth Greenhough and Derek McCormack, my supervisors on this journey. In particular, Beth agreed to supervise an oddity wanting to return to geography and she was central to me exploring malware at all. Beth’s expert tack to translate my often convoluted thinking into a pithy summary as well as Derek’s incisive comments (particularly around my theoretical wanderings) are ones I shall miss. Others in geography have made the five years worth it – Helge Peters, Ian Klinke, Peter Martin, Liam Saddington, Daniel Bos, Marion Ernwein, Alex Vasudevan, Jamie Lorimer, Hanno Brankamp, Negar Behzadi, Thomas Jellis, Adam Packer, and Gillian Rose. Beyond Oxford, people have been so influential and beyond helpful and kind – Louise Amore (who I will always owe an inextricable amount of debt to), N. Katherine Hayles, Sam Kinsley, James Ash, Nat O’Grady, John Morris, Ingrid Medby,

Peter Forman, Emma Fraser, Georgios Glouftsios, Clancy Wilmott, Esmé Bosma, Phil Garnett, Sam Hind, and Dorothea Kliene (including the rest of the 'Digital Geographies Working Group' at the RGS-IBG, as it was then known).

The 'sister' CDT at Royal Holloway and Mansfield College always provided sanctuary. From Mansfield, Baroness Helena Kennedy and Ros Ballaster for supporting me as MCR President, to the amazing life-long friends at the *Cowley Road Breakfast Club*, Joe Shaw and Irene Scarpa. At Royal Holloway, Pete Adey, and friends who I cannot imagine life without, who exposed me to so much among the general gallivanting around London - Pip Thornton, Nick Robinson, Adam Badger, Alex Hardy, and Andreas Haggman. Thanks also go to the different institutions that have supported me on the way, including Joyce Tait and Gill Haddow at the Innogen Institute at the University of Edinburgh. Also to the reading groups I attended and in particular 'Life Itself in Theory and Practice' at TORCH in Oxford that really helped when I was attempting to 'contextualise' my work, thanks to Sam Gormley for putting this together. And happily to my new colleagues Andreas Langenohl, Carola Westermeier, Amina Nolte, and Philipp Lottholz at the SFB/TRR 138 'Dynamics of Security' project in Marburg and Gießen, Germany, which gave me the space, time, and freedom, to finish this thesis.

And nearing an end, to thank those closest to me. Friends from home; Jessica Wilson who always reminded me of good drinking spots in Sheffield, and to Kerry-Anne Scothorne Allen, Billy, the boys, Brooklyn and Ronny and the new one, Crimson. Those who I met at Durham, where I developed a love of geography, politics, and technology; Alice Cree, Louise Harper, Emma Holloway, and many more. To those who I lost in this time; my grandparents, Joan Green, Brian Green, and Winifred (Ann) Dwyer. And lastly, to thank those I love the most: Mum, Matthew, and the one who has had to deal with more than enough grievances from me, Oliver.

This thesis is dedicated to my wonderful Mum, Deborah Jayne Green, who never seeks recognition and to Jessica Wilson, whose panic one day made malware all too real for me.

Contents

| | |
|--|------------|
| ABSTRACT | I |
| ACKNOWLEDGEMENTS | III |
| CONTENTS | VII |
| FIGURES | IX |
| CHAPTER ONE INTRODUCTION | 1 |
| MOTIVATIONS | 3 |
| THESIS STRUCTURE | 7 |
| CHAPTER TWO COMPUTATIONAL ECOLOGIES | 13 |
| ORGANICISM AND CYBERNETICS | 14 |
| ECOLOGISING COMPUTATION | 19 |
| MORE-THAN-HUMAN COMPUTATION | 27 |
| A MALWARE ECOLOGY | 36 |
| CHAPTER THREE RESEARCHING MALWARE | 39 |
| APPROACHING MALICIOUS SOFTWARE | 41 |
| ETHNOGRAPHY | 50 |
| BECOMING ANALYST | 53 |
| RECORDING, REFLECTING, AND ETHICS | 67 |
| CHAPTER FOUR LINEAGES TO UNDERSTANDING MALWARE | 73 |
| THE GROWTH OF CYBERSECURITY | 75 |
| ARCHITECTURES OF MALICIOUSNESS | 82 |
| PATHOLOGICAL IMAGINATIONS | 86 |
| ANTI-VIRUS | 97 |
| CONTEMPORARY APPROACHES | 103 |
| CHAPTER FIVE THE ANOMALOUS AND THE MALICIOUS | 105 |
| THE LAB | 106 |
| THE ALLURE OF ENVIRONMENT | 108 |
| STATIC STRATEGIES | 110 |
| CONTEXTUAL STRATEGIES | 117 |
| DATA | 131 |
| CRAFTING SIGNALS | 136 |

| | |
|---|-------------------|
| BLENDING LOGICS | 137 |
| <u>CHAPTER SIX CURATING MALICIOUSNESS</u> | <u>141</u> |
| EMBODYING MALWARE | 142 |
| THE FALSE POSITIVE | 146 |
| THE POTENTIALLY UNWANTED APPLICATION | 151 |
| MAKING SENSE OF MALICIOUSNESS | 156 |
| MOVING TO ECOLOGY | 163 |
| <u>CHAPTER SEVEN MALWARE POLITICS</u> | <u>167</u> |
| ECOLOGICAL CURATIONS | 168 |
| THE CASES | 171 |
| CASE 1 STUXNET | 172 |
| CASE 2 THE DUKES | 178 |
| CASE 3 WANNACRY/(NOT)PETYA | 183 |
| MALWARE MATERIALITIES | 193 |
| <u>CHAPTER EIGHT ECOLOGICAL CONCLUSIONS</u> | <u>199</u> |
| THE CONFICKER WORKING GROUP | 200 |
| THE CYBER DOMAIN | 207 |
| THE OBJECT-IVE CYBERWEAPON? | 210 |
| DRAWING A CLOSE | 214 |
| FUTURE DIRECTIONS | 218 |
| <u>BIBLIOGRAPHY</u> | <u>221</u> |
| <u>APPENDIX</u> | <u>243</u> |
| CONSENT FORMS | 243 |

Figures

| | |
|---|-----|
| Figure 1: A simplified depiction of the 'layers' of computation that form the materialities and spaces of 'cyberspace.' Adapted from a presentation by Hayles (2018). Author's own image. _____ | 31 |
| Figure 2: A screenshot of the Ingress game, where players work together in teams to capture 'portals' from other teams through GPS locations and performing ever-more complicated tasks. Author's own image. _____ | 50 |
| Figure 3: Research diary with coded strips highlighted. Author's own image. _____ | 69 |
| Figure 4: New malware forms over the past decade using data from independent anti-virus product testers, AV-Test (2018, p.2). _____ | 101 |
| Figure 5: A screenshot of the Menu Bar on MacOS with 'Sophos Anti-Virus' presented as the small icon shaped like a shield with an 'S' that persistently remains whether the program is 'open' or not. Author's own image. _____ | 104 |
| Figure 6: An overview of how two networks, yellow and green, operated via a representation of three screens. The solid black lines represent virtual machines with the dotted line being based on the underlying yellow network. Author's own image. _____ | 107 |
| Figure 7: Overview of some 'static' tools. Top-left: IDA Pro Disassembler; Top-right: PEView; Bottom: PESTudio. Author's own image. _____ | 113 |
| Figure 8: An abstracted signature detection in a proprietary stack-based language, Virus Description Language (VDL). Abstraction taken from author-written detection (Research Diary, 9 June 2017). Author's own image. _____ | 115 |
| Figure 9: Overview provided by Saxe and Berlin (2015) in their construction of a form of neural network for malicious software detection, based on training data from Virus Total. _____ | 126 |
| Figure 10: A graphic representation of a 'neural network' algorithm with different layers, neurons as circles, and the weight by each line, connecting each neuron in another layer. This does not mean that every neuron becomes 'activated' but shows possible movements. Author's own image. _____ | 127 |
| Figure 11: An example of the Virus Total interface when searching using a file 'hash' (a unique cryptographic signature). In this case, I used a SHA-256 hash of a WannaCry sample. Screenshot taken 21/03/2018. Author's own image. _____ | 133 |
| Figure 12: Theatrical one-sheet for ZERO DAYS, a Magnolia Pictures release. Photo courtesy of Magnolia Pictures. _____ | 177 |
| Figure 13: Left: 10 February 2017 sample (MD5: 9c7c7149387a1c79679a87dd1ba755bc). Screenshot taken from Twitter (S!Ri, 2017). Right: The updated decryptor (ransom note) screen and the change from 'WCry' to 'WannaCry' from 27 March 2017. Screenshot from Twitter (Malware Hunter Team, 2017). _____ | 184 |

Figure 14: Graph showing WannaCry infection rates across differing operating systems, photo from Twitter on 19 May 2017 (Raiu, 2017). _____ 187

Figure 15: The (Not)Petya ransom note. Taken from Sood and Hurley (2017). _____ 191

Figure 16: Liam O’Murchu demonstrates how a Programmable Logic Controller (PLC) can be manipulated to create ‘unexpected’ action through the popping of a balloon. Screenshot from YouTube video (Perez, 2014). _____ 196

Chapter One || Introduction

We are frequently told computation has errors, glitches, and bugs; its irregularities and insecurities. But rarely do we think that it has a politics that exceeds human intent, and this is something that this thesis attempts to unpick. Malicious software, a form of computational materiality – commonly referred to as computational viruses and worms – trail across a collective imagination of vulnerabilities in our networked cultures; modulating, transforming, and unsettling senses of orderly computation. The threat of those who *write* malware – hackers – morphs between lone individuals, cyber gangs, and complex state operations that force us to change passwords and update our computers with the latest security ‘patches’¹. Malware have thus become the threat *du jour*. In the way they propagate, transform, and render computation alternatively from our expectations of what should occur. Hackers seek to direct malware, *intent* on wreaking havoc on our interconnected worlds. Electronic, digital computation intermingles with and permeates varying geographies at differing intensities, where its (in)securities have become essential to many everyday interactions such as banking, communications, and governing, structuring and informing how we live today. These collective vulnerabilities encouraged, from the 1980s onwards, the privatised response that keep our computers ‘clean’ and free from infection: ‘anti-virus,’ and later, endpoint detection products that exist on billions of computing devices worldwide². They detect and counter these malicious interlopers; and in the process, define, articulate, inform, and detect what *is* malicious. Detecting malware have become paramount to protect the internet economy, wellbeing, as well as the stability, security, and prosperity of states.

In this thesis, I explore how we have come to understand *what* malware are, *how* this shapes and forms a ‘more-than-human’ politics, and *how a conceptual focus on ecology* can challenge and extend current thinking on cybersecurity. Pursuing such a ‘timely’ topic does bring with it the risk of falling into a politicised environment with cybersecurity rising on state agendas and increasingly reported upon by the media. This thesis covers sensitive topics including the stability of internet infrastructure (Conficker), the security of democratic institutions (*The Dukes*), nuclear uranium enrichment (*Stuxnet*), disruption to parts of the UK’s National

¹ Patches are ‘repairs’ to vulnerabilities found in software, and typically refer to operating systems

² There is very little information on the total number of devices that have anti-virus, but it is likely to be at a minimum this figure.

Health Service (*WannaCry*), and the near wipe-out of computing infrastructure at the shipping and logistics business – AP Møller-Mærsk (*(Not)Petya*). In this project, in pursuing a more-than-human politics, I do not dismiss the damage malware cause, however. That would be irresponsible: they impact healthcare and enable money to be stolen from bank accounts, they are used to exfiltrate sensitive and private information, hold computers to ransom, and are used to spy on domestic abuse victims. Yet, unlike critical work on author intentionality and public attribution (Egloff and Wenger, 2019), I pursue a different orientation. Work on malware has become terminally caught between, admittedly essential, technical analyses and considerations of its strategic implications. I instead attempt to open this up by taking malware to be an important *political* actor through the choices they make – where human and more-than-human agencies co-constitute and negotiate cybersecurity – leading to particular forms of spatialisation and securitisation. I consider humans as only one element of their stories. I argue, instead, that it is imperative to recognise the ecologies of malicious interactions, the role of computational cognition and performativity, as well as the spaces in which they become known. I develop an eco-logic as a conceptual and methodological intervention to present and make sense of the complexity of this computational actor; one that is not simply a tool, nor a mathematical, deterministic automation.

In the chapters that follow, I detail how software have become demarcated as malicious, outline the impacts of malicious activity, and demonstrate how malware exhibit political agency through not only an affectual capacity, but also cognitive choice, that blend together in ways complex and multifaceted. As cyberattacks³ have become more prevalent and extensive, it is crucial to explore the processes that determine what malware are in order to understand how societal securities are anticipated and governed, and thus responded to, in this second decade of the twenty-first century. Such a project means I do not engage in a listing nor ranking of cyberattacks, not because each case is insignificant, but due to their sheer volume, and conscious of the quick pace at which many fade into obscurity. I am aiming for something deeper; a re-evaluation and re-orientation of our relationship to malware. In *a politics of cybersecurity*, I rethink how space and security are being transformed

³ I define a cyberattack different to common usage of the term. That is, I wish to use it to say that a 'malware' form has *performed* actions that do not correspond with how the affected user expect their machine to *function*. This allows cyberattacks to not only be about malicious intent, but include glitches, bugs, and other events that are commonly associated with non-intent.

by computation, whilst addressing an urgent requirement to consider malware in the social sciences to expose their important role in our lives.

Motivations

I was motivated to research malware due to a concern that the social sciences have not interrogated malware enough as politically constitutive of cybersecurity. There are two substantial exceptions that I however take inspiration from and are important precedents that I use throughout. Both Jussi Parikka's *Digital Contagions* (2016) and Jessica Johnston's *Technological Turf Wars* (2009b) offer essential but alternative interpretations and emphasis. Whereas Parikka delves into the history and emergence of malware as part of a media archaeology of viral capitalist network cultures, Johnston interviews analysts to understand the 'anti-virus industry.' For a project exploring a politics of cybersecurity through malware, Parikka's in-depth work into malware's lineages and Johnston's empirical work with those who analyse, detect, and curate malware are both essential. Yet, these too suffer from little sustained direct engagement with malware. That is, within the social sciences, there is no empirical research on the spaces in which malware are detected, how this intersects with malware's discursive lineages, and how these situate and render malware as products of human intent.

Central to the arguments I make arrive from time spent in an empirically-grounded encounter that sought to address malware *in situ* and then rethink how we relate to these others. This is a reading that cannot and could never be wholly quantified, one that recognises our knowledge of malware will always remain partial, and that computation's agency always extends away from our apprehension. To allow and permit the sensitivity to malware's technical detail and social tapestry, and to articulate what I call a more-than-human politics, I spent seven months at the UK endpoint security business, Sophos. I immersed myself in the practices, politics and technologies of its malware analysis and detection laboratory (the MAL). I trained as an analyst – learning, making mistakes, meandering, and eventually analysing and writing detections for Sophos' products. My method – *becoming-analyst* – that I detail in the methodological chapter three, offers a unique, detailed perspective on the agencies of computation and malware, along with analysts in a commercial MAL, to develop an ecological reading of malware beyond it too. The products

of these MALs are frequently endpoint security products – ‘anti-virus’ – that is present on the computer I have written this thesis upon – to not do so would make me more vulnerable to malware. But, I barely even notice its operation as it scans the documents I download and open, and how it stops malicious activities I may encounter through my internet browser. This banality does not mean, however, that there is not a whole curative negotiation behind this ‘engine’ – not only in how one defines maliciousness, but in the commercial drive to detect and sell endpoint detection products, but the pathologies that structure how we respond to malware, and the lineages that have informed this to the exclusion of the agencies of computation and malware.

Malware have tended to be approached, in popular imaginations and in practice, through these varying pathological approaches. I am not the first to identify how pathology infuses malware discourses. Early ‘anti-virus,’ the practices we engage in, and the language we use – including terms such as virus, infection, and keeping computers ‘clean’ – have structured social responses to malware risk and computational vulnerability, and in so doing, take analogical cues from biology and medicine. My aim is to trace an alternative path to that implied by pathology. Instead of concentrating on how computation and malware *execute*, I consider how an attentiveness to a *performative* choice-making computation can generate a new perspective on cybersecurity. This requires being attentive to how ecologies are formed, redrawing the contours of who and what participates in politics, and how malware are critical to how cybersecurities are produced. This is an ever-incomplete process and I make no claim to a *definitive* politics of cybersecurity. Instead, I offer points of departure, alternative renderings of the political, in order to help grasp the plurality of agencies that are central to cybersecurity through the support of new materialist literatures and beyond. Drawing on works such as *Vibrant Matter* by Jane Bennett (2010), to N. Katherine Hayles’ (2017, 2019, Forthcoming) expanded role of cognition, to rethinking the place of the human in an age of the Anthropocene (Stengers, 2010, 2015), and staying with this trouble (Haraway, 2016), I contribute to questioning human exceptionalism and its implications. This requires an assessment of the ecologies of computation – that are constructed, performed, and negotiated by a variety of human and nonhuman actors and things.

Through immersing myself in the world of malware detection, I learnt that the question of what kinds of software are defined as malicious is far from clear. As I seek to argue, processes

and architectures are not necessarily malicious, but they may become so in an ecology of practices largely framed by the expectation that computation and software must work as we expect. Here, it is important to recognise the indeterminacy in how we talk about computational problems: we have different names for the computational unexpected. We do not always term such incidents ‘malicious,’ and may instead talk of bugs, failures, or glitches. Some unexpected processes may be seen as beneficial in one context, and not in another. This is not to say that malware as deliberate human products *do not exist* – there are actors who write software with malicious intent. But not everything that comes to be seen as malicious begins in this way, as Sony’s *Digital Rights Management* system demonstrates. This system installed software (a rootkit⁴) to stop replication of copyrighted content without the knowledge of the user (Touchette, 2016), becoming quickly regarded, though not exclusively, as malware. Certain software can then become deemed to be malicious (though this may be disputed) even if the intent of their authors in particular ecologies of practice and expectation were *not* malicious. From this perspective, software performance, even when there are calculative and logical processes, does not mean there are simple deterministic processes that makes software malicious. Malware act, and in their pathological rendering as ‘execution,’ are *logical* (and this must be kept distinct from determinism⁵). An unsettling of *a priori* categorisations of software as malicious then become of utmost importance. How we detect and come to see computational processes as malware are key to its constitution, not everything that begin as malware become malicious⁶, and things that become malicious may have been previously conceived of as benign.

Ever since the first ‘bug’ found in computation – a moth in Mark II by Grace Hopper in the 1940s at Harvard University (Shapiro, 1987)⁷ – it has been clear that computation has always exceeded a mathematical and logical grace. There is an already-possible and ever-present ability for computation’s own degradation and, simultaneously, a potentiality for a politics

⁴ A rootkit is a program that is designed to give enhanced permissions to a computer whilst also hiding its presence.

⁵ In my reading, determinism suggests that there is a *singular* and *predetermined* movement. This would suggest that malware follows its author’s intentions completely. Not only this, but it would suggest that every environment, if could be known, would follow a certain route. In my reading of logic, this is the *most likely* route – but that this is not necessarily the only one that may be following according to environmental factors, leading to non-deterministic outcomes that can still be logical.

⁶ By this, I mean that even if something is written with malicious intent, it may not achieve the intent of its author, and ‘fail’ or ‘incorrectly execute’ or simply not be compiled correctly.

⁷ Following the work of Shapiro, it is important to note that the use of the word ‘bug’ does not come from this moment. Indeed, there are earlier uses of the term, and indeed the verb ‘debug,’ but this is likely to be the first time in the context of computing machines.

to emerge. As N. Katherine Hayles notes with reference to computational materiality, uncertainty even occurs “when cosmic rays flip bits within a computer, analogous to when cosmic rays cause a mutation in a gene” (2019, Forthcoming). This uncertainty of computational materialities permits a consideration of how malicious software itself could render computation anew. Put simply: under what conditions do software become malicious, and why is this? What can we infer from the human expectation of intention in computation? What happens if we think of malware eco-logically rather than patho-logically? And, just as importantly, if malware operate in ways beyond the expectation of their author, what happens to the politics of cybersecurity? For example, such questions of attribution, the concept of weapons, and the participants of politics and security are fundamentally questioned by an expanded view of malware agency. This is a great challenge from conventionally seeing malware as the direct extent of human intention and craft. Indeed, it helps to bridge gaps in debates around who are responsible not only for malware, but also contributes to debates on agency and its relation to human decision-making in machine learning and algorithmic practice.

These questions have usually been restricted to computer science, most recently under the rubric of cybersecurity. But this is only the most recent attempt to conceptualise and collate a series of protocols and practices of computational securities; with ‘computer,’ ‘information,’ and ‘data’ all used prior and concurrently with this most recent iteration. For geographers and those in the social sciences, *cybersecurity’s* simultaneous ‘physical’ and human-centred focus is perhaps surprising; it does however falsely demarcate between ‘cyber-physical’ (Cardenas, Amin and Sastry, 2008; Lee and Sokolsky, 2010) and ‘human-centric’ computing (Coles-Kemp and Hansen, 2017; Denning, 2014). But cybersecurity extends and reaches beyond computation as a ‘boxed’ technology, to include human actors, primarily, as an essential part of conceptualising security (Stevens, 2015). This has had important and significant implications for how security has been practiced, as it has sought to understand how humans work in wider intersectional human-computer apparatuses (Dix, 2009). As cybersecurity in academia, industry, and government recognises that computational securities cannot be solely understood as technical nor social phenomena, it permits an entry point for my research to develop within these interconnected perspectives between the social and technical to extend and compliment who and what practices and performs security. To do this, I make an argument that this requires a (re)turn to cyberspace

– though not in how it is popularly articulated – to permit the space for choice for computation, and thus malware.

Developing the concept of *malware ecologies* is a pragmatic entry point to such debates and I thus must speak to many (and complex) disciplines and movements. These include, but are not limited to, anthropology, science and technology studies, critical security studies, international relations, computer science, and my ‘home,’ geography. Though this thesis is written in the praxis of geography, I hope to speak to a number of disciplines beyond this. This means I delve into some geographical debates whilst remaining open to others who may not arrive from such a background. I weave technical concepts and descriptions that may be new to social scientists and likewise philosophical and theoretical debates less likely to be encountered by those who focus on the more ‘technical’ aspects of cybersecurity. This speaks to the interdisciplinary setting I have written this thesis within, between both the University of Oxford’s Centre for Doctoral Training in Cyber Security (hosted by Computer Science) and supervision at the School of Geography and the Environment. A sensitivity to speak to both social scientists and those who focus on developing computation with which cybersecurity is *written*, is in itself an ecological project, to render oneself vulnerable. My hope is that this is the start of a re-reading of what is political, and through malware, permitting space for a broader discussion of the production and negotiation of cybersecurity.

Thesis Structure

A politics of cybersecurity, through my (auto)ethnographic research as *becoming-analyst*, attempts to rethink and rearticulate how malware and computation are themselves political, which requires not only looking towards the practices of analysts who analyse and detect malware, but the materialities and temporalities of the MAL, the translation of this information, and how this has become practiced. This can only be successfully done by at least partially integrating oneself into such ecologies; one that is insufficiently achieved through short-term observations and interviews. In doing so, I expose the pathologies that emerged through the lineages of computing and (cyber)security to see and comprehend malware as political actors through an eco-logic. This has led to three research questions for an *eco-logical* politics of cybersecurity. Each question presented below is not tied to a particular chapter, but they all infuse the thesis at varying intensities at different moments:

1. *How is software categorised and disseminated as malicious?*
2. *How does malicious software express political agency and become a political actor?*
3. *How does rethinking malicious software through ecology challenge conventional (pathological) approaches to cybersecurity?*

To address these questions, I explicitly attend to what I argue is an essential ecological repositioning of malware in chapter two – *Computational Ecologies* – by considering their capacity to be political. This requires conceptualising how computation can only be understood through its entanglement. This requires reassessing cyberspace as a way to consider how there can be a more-than-human spatiality that is dependent on the materialities of computation, but which is not wholly captured by the ‘digital.’ I consider how cognition in computation provides an alternative rendering of choice through reading N. Katherine Hayles. This requires challenging how technology and computation have been understood through organicism and cybernetics and how computation has been previously theorised. In outlining a ‘more-than-human’ computation, of which malware are a part, I seek to move to performance, but not one tied to the logical bounds of computation. The combination of cognition and ecology allow for a more-than-human politics to be considered which is not restricted to organic, *lively*, matter and centres choice as an additional, and essential, component to identify political actors.

In chapter three – *Researching Malware* – I turn to how I approached this project through the method of *becoming-analyst*. In introducing the spaces of my research at the Sophos MAL, at the Virus Bulletin conference, through interviews, as well as four cases, I discuss how an (auto)ethnographic project can inform a new perspective on cybersecurity. In drawing upon the training I received, how I worked with the other analysts, and some of the difficulties of entering the MAL, I outline how post-phenomenology and resistance underpinned my methodological approach. This was not without its own issues; about how one studies malware when one cannot be ‘with’ the object of one’s research, and how doing security research requires negotiating and balancing access to these spaces. As I reflect on my research with malware and the limitations of ecological research, I present how there is always more to be learnt and argue that our encounters with malware are always situated, bound, difficult to contain, and ultimately withdrawn and partially inaccessible.

After setting out the methods and basis for thinking and researching ecologically, I attend to the particularities of how malware and cybersecurity have been constructed in chapter

four – *Lineages to Understanding Malware*. This chapter explores the pathologies that have informed the practices of malware analysis and cybersecurity in its current formulation. A variety of influences have impacted on computing such as the military-industrial nexus that have particular understandings that inform us what is malicious, to how security itself has always been pathological through disciplinary to biopolitical societies, and how concerns over control are reflected in 1980s cyberpunk literature (Parikka, 2007). Within the discourses and practice of computation, there has always been an experimental other, epitomised in the maverick figure of the hacker during the 1960s and 1970s, and later ‘benevolent’ viral and worming architectures. At the turn of the century, this changed significantly and, increasingly, malware authors were represented as dangerous hackers from whom we (or at least the state) needed to be protected. The trend to control has been sought by governments and corporations to demonstrate an enrolling of digital technologies into security frameworks – where risk is incessantly filtered, sorted, and sifted – and cybersecurity is not immune from these broader pathological movements in biopolitics. ‘Anti-virus’ – at least since the 1980s – then has played a central role in the actualisation of these lineages for the (imagined) borderlands of computation, drawing anomalies and abnormalities together through senses of medical normativities, to define what are malicious through privatised security. This chapter, then, draws together these various lineages together in ways that complicate the histories of malware and cybersecurity.

In chapters five and six, I explore the malware analysis laboratory at Sophos and how there are curative negotiations of maliciousness that construct and amplify our popular imaginations. In chapter five – *The Anomalous and the Malicious* – I explore a space where multiple techniques are employed, along with data, to comprehend malware. In contrast to assuming maliciousness as a kind of given or essential quality, I argue software becomes malicious through these interactions through the sharing of knowledge between analysts, bodily affects, and more-than-human technologies. I identify two significant impulses, or strategies, that can be teased apart in the approach taken by the MAL to defining maliciousness: one *static*, the other *contextual*. The static renders malware as something to be dissected and analysed, abstracted from the environments in which they perform, whereas contextual strategies use data and monitor behaviour to determine maliciousness. This in turn informs software ‘reputation’ and provides (big) data for machine learning algorithms that introduce greater deviances and difficulties in determining maliciousness. Through

static and contextual strategies, I show how the MAL treats malware in similar ways to the pathologist; dissecting, analysing, and searching for anomalies.

Continuing the exploration of the malware analysis laboratory in chapter six – *Curating Maliciousness* – I investigate how I embodied the search for malware through curative acts. I define curation as the work performed by analysts to bring together and experiment with disparate data and information to analyse, detect, and thereby co-produce knowledge about malware. The technologies and practices identified in chapter five are deepened to explore how the workings of the analyst are shaped by both the patho-logic of the MAL and the commercial imperatives of an endpoint protection business. To develop an awareness of how this operates, I focus on two pressing issues. The first being the false positive, whereby software are incorrectly detected as malware, and second, the Potentially Unwanted Application: a quasi-malicious detection that makes the user decide if they want to have that software execute or not. Through these, I query conventional understandings of malware; as these two issues stretch and contextualise what can be malicious and open-up the curative practices of malware analysis and detection.

I demonstrate that curations of the malicious are not undertaken solely by the malware analyst and the computations of the MAL; but emerges through a conversation with quality assurance processes, customers, and environments beyond it. Through contextual strategies, data are collected for anticipatory action, that are collated and rendered visible in visualisations and feeds that direct (more-than-)human attention to what anomalous software to analyse next. Curation is not clear-cut and is imbued with competing priorities and maintenance of testing systems that ensure that different software are given varying levels of attention – in order to commercially identify what should be prioritised. Simply, anomalous software that has not yet been delineated as clean nor malicious must come back to human analysis, so that new detections or existing ones can be (re)written. These curative processes can be haphazard, yet there is a dual activity taking place; where testing and quality assurance condense the multitudes of agencies and techniques in the laboratory into industrial-like detection production unlike the more ‘experimental’ analysis. This is required to enable the stability required of a commercial laboratory, where detections do not generate ‘false positives,’ and thus endpoint protection businesses are seen as competent.

In chapter seven – *Malware Politics* – I step away from the MAL to reconsider and re-write malware ecologically through three cases using my experience as becoming-analyst, in order to demonstrate what might be achieved by shifting from the patho-logical that drives the detection and imaginations of malware to an eco-logical approach. Malware perform in ways which frequently exceed author intentionality (they do things those who devise them do not anticipate), and almost always perform differently as they read and construct the environments they are within, even if they *logically*, and frequently, arrive at a similar outcome. To demonstrate this, I detail three registers of political agency based on *output*, *architecture*, and *surprise*. In writing about a more-than-human malware politics, I draw on the translation and encounter of specialised malware knowledges as becoming-analyst that I developed in the MAL and contrast this with the minor narratives I wish to draw upon. Through the cases Stuxnet, the Dukes, and WannaCry/(Not)Petya, I argue for a reconsideration of their stories through a focus on computational choice as well as human agency. My claim is that malware become important *political* actors, where agency needs to be taken into account in comprehending cyber(in)security.

Then, in the concluding chapter – *Ecological Conclusions* – I consider how security is already being shaped by more-than-human agency and malware politics. Through a final case of Conficker, I detail the emergence of a proto-ecological politics distinct from contextualisation in the formation of the Conficker Working Group. I take a view of what thinking ecologically might mean for how we understand who should take responsibility for authoring malware (attribution) and ask what happens to concepts of the ‘cyber domain’ and ‘cyberweapons’ in an expanded view of politics and agency. By recognising the need for the pathological as a pragmatic way to *deal* with the detection of malware – I do not argue that ecology should somehow replace pathology - but rather that it should be seen as an expansion of it, relinquishing human-centred notions of who and what practices security. The materiality of malware, as much as, if not more than, the intentions of its authors, shape the choices malware can make. Sometimes the choices malware can make are exceptionally limited; at other times there is much greater scope. The question of authorship and intention of malware are still important – and depending on their intentions – it may be useful for them to enter a more-than-human relationship that extend agency and political choice in order for certain malicious actions to be performed. However, what is distinctive about such an approach is how choices are formative of the political itself, and choices made by malware

become interpreted by us, that lead to further choices being made, and with this, decisions are made. Hackers, MALs, and governments already engage in a malware politics, a certain politics of cybersecurity, even if they do not realise it. Humans do not recede in this account but are resituated. This thesis then attends to these ecological politics to make present how malware choice already pervades cybersecurity and what (re)cognising this means for the future.

Chapter Two || Computational Ecologies

The growth of computing, and awareness of computational agency, specifically through algorithms and machine learning, have profoundly questioned human exceptionalism. This questioning, and older, related fears of technology⁸, have been articulated in contemporary concerns around the impact of computation and information on employment, privacy, and security – no more so than in Shoshana Zuboff's *The Age of Surveillance Capitalism* (2019). This has brought to the fore a re-evaluation of how human society relates to computation, both within and outside academia, that seeks to shine a new light on these relationships. In this chapter, I then seek, in similar ways, to ask how computation itself can be thought of in our contemporary moment. I do so through arguing for a rapprochement with cyberspace as a central way to conceive of security and computation today, that builds upon work on digital geographies (Ash, Kitchin and Leszczynski, 2019), but also expands beyond it. I suggest that an appreciation of nonhuman cognition is central to a recognition of computational agencies already present in our societies, suggesting this recognition can radically reconfigure the study of space, politics, and security. This has critical implications for the study of cybersecurity; as acknowledging computational agency, through what I detail as cognition and choice, opens-up and complicates who and what participates in the production of contemporary spaces of politics and (in)security that percolate the rest of this thesis.

I argue that an ecological approach opens-up a 'logic' where more-than-human agency, security, and space come together to produce certain configurations of politics. Ecological thinking has been applied to other areas of thought too, including; climate change (Morton, 2010), the Earth as Gaia (Stengers, 2015; Latour, 2017), the planetary (Connolly, 2017), through the *Anthropocene* (Castree, 2015; Lorimer, 2012; Gibson-Graham, 2011; Steffen et al., 2015; Yusoff, 2013) as well as Donna Haraway's *Chthulucene* (2016) and Jussi Parikka's *Anthroboscene* (2014). This condition has also moved to claims of planetary computation as argued in Benjamin Bratton's *The Stack* (2015). Through these, I situate a particular articulation of the ecological – *malware ecologies* – in relation to broader theoretical thinking around new materialism and the work of N. Katherine Hayles. I do this to achieve three

⁸ This can be extended back to critiques from Walter Benjamin (1992) and Martin Heidegger (1977; Harman, 2002), though there is a much longer lineage to this with the advent of industrialisation.

aims: one, to consider malware as political actors (as partially generative of malware agency); two, to investigate how to locate the emergence of malware at critical junctures of environments, more-than-humans, and technologies (understanding malware spatialities); and, three, to explore how computational and human agencies are formative of, and are formed through, malware ecologies.

I structure my argument into three sections. First, I consider how *organicism* and *cybernetics* have structured how we think about computation, information, and now the analogies drawn between computation, biology, and medicine have constructed malware as under the control of humans. Second, I explore the potential for framing computation through ecology, and how this opens-up the idea of computation as performative and capable of acting (exercising choices) in ways unanticipated by its authors. And third, I seek to understand what a ‘more-than-human’ computation, with an ability to make choices through cognition, means for politics and for the production of security. I then conclude on what such a positioning and critique means for researching malware.

Organicism and Cybernetics

In the earliest days of computing, there were continual references back to humans as a source of machinic inspiration. This relied on an equivalence between the organic and the inorganic, that can partially trace its roots from the philosophy of both René Descartes and Gottfried Leibniz, to the mathematician Alan Turing and into the cybernetics of Norbert Wiener. In this section, I argue that organicism, as a forerunner to cybernetics, is central to how we see the role of computation and its agency in the early twenty-first century. From the 1930s, the computing sciences drew upon physiology and the philosophy of the mind (Channell, 1991), which continued through analogy to the body and artificial intelligence. Alan Turing and John von Neumann – some of the earliest figures of modern computing – both made biological analogies between humans and computing. Turing’s paper ‘Computing Machinery and Intelligence’ (1950), popularly known as the *Turing Test*, stipulated a formalised method where human and computer are hidden and respond to questions. If the computer can respond and convince the questioner that it/he/she/they are human, then it is ‘intelligent.’ Whereas von Neumann in 1958 (2012), developed an analogy between the human brain and body to the computer in a discussion on neural networks – which are now used extensively

in malware detection and in other algorithmic practice. Yuk Hui (2019, p.146) asserts that such comparisons between computational capacity and human intelligence are a reductionism that separate form from matter, flattening the potentials and differences between organic and inorganic materialities. This flattening between organic and inorganic is what permitted computer scientists such as Cohen (1994, p. 32) to claim that humans ‘compute,’ for instance.

Hui argues in *Recursivity and Contingency* (2019) that organicist thinking has been central to how we approach computation – which led to cybernetics’ focus on feedback, and later recursion, through equivalence and flattening. This was dependent on a broad, and deep, history conceptualising organisms as machines based on analogy through Cartesian mechanism – that allowed for analogies between organisms and mechanism to take place. Descartes, for example, came to understand biology through *mechanisation* whereby bodies could be understood through such components as valves, pipes, and pumps as in the *Treatise on Man*, published in 1633 (Descartes, 1972). Yet, at the same time, Descartes marks a distinction between human and nonhuman, arguing that nonhuman animals were machine-like as they were “devoid of mind and consciousness, and hence lacking in sentience” (Hatfield, 2018, np). This suggests that there is something that makes humans distinct – the mind and conscience. As Leibniz writes later, the demarcation between *natural* and *artificial* machines is also of great importance. It permits a bifurcation, albeit in a different way to Descartes (Des Chene, 2001), between those beings that are created by a divine power (organics), and are thus more complex and ‘natural,’ than those by humans (machines). Limited mechanism may apply to organisms, but human-created machines are much simpler, less complex, and inferior. This presents an equivalence of function *but* not an equality. Analogies abundantly flow, but machines are regulated to a second-order behind ‘natural’ things, and non-human organics without consciousness.

These traits were integrated into the cybernetics movement. In the late 1940s, the study of systems and feedback drew on a similar equivalence of function between things. It assumed the world could be objectively studied through mechanistic analogy (Clough et al., 2015). This strand of thought emulsified and developed, though not exclusively, at the retrospectively named *Macy Conferences* - that fostered a complex, radically interdisciplinary exchange of ideas (Hayles, 1999). Figures including von Neumann, and others such as Norbert Wiener

and Claude Shannon, crafted an approach to studying the world based on systems, transmission, and feedback. Wiener (1948) argued for an equivalence between animal, human and machine by abstracting all things to the transmission of information where objects could be compared and contrasted on this plane of abstraction. Both Shannon and von Neumann were crucial in persuading the conferences that *the* mode of communication for understanding systems was information, not thermodynamic energy (Hayles, 1999), which was the other dominant, alternative method of understanding communication at that time.

Wiener claims in his book outlining this new field of study - *Cybernetics: Or Control and Communication in the Animal and the Machine* - that, “[i]f I were to choose a patron saint for cybernetics out of the history of science, I should have to choose Leibniz” (1948, p.12). He chooses Leibniz due to his dual work on logic and the organic – in particular through his *Monadology* (Leibniz, 2001, originally published in 1720) – that through various other works combines both his idea of a ‘reasoning machine’ and his philosophy of the organic. It is the former’s calculative mode and the latter’s potential mechanisation according to biology, that allow for calculative ‘natural laws’ in cybernetics⁹. An organicist tradition extends through Leibniz, as well as other philosophers, such as Hegel, who in turn influence the logician Kurt Gödel to develop a logical ‘arithmetization’ (Hui, 2019, p.110), with an ability to see ‘life’ as mechanical and thus the mechanical as organic. Once all organic and inorganic matter are based on calculative logic, it allows for a common denominator to be drawn across these by cybernetics: information.

A particularly abstract model of information becomes crucial here. For Shannon, (cybernetic) information must be devoid of meaning. Shannon attempts to reduce noise to zero, as this allows for the ultimate expression of translation to information. That is, information does not need context (i.e. noise) in order to be transmitted. This abstraction of information from meaning has structured much thought in the twentieth and early twenty-first centuries. As Paulson (1988, p.48) notes, “the role of information in the structure and function of organisms is the contribution of twentieth-century biology, which has modelled life on transmitted signals, codes, even computer programs, just as the eighteenth and nineteenth centuries modelled it on mechanisms and motors.” Though Paulson divides

⁹ For a detailed contribution to this debate, see Yuk Hui (2019) on how Wiener uses Leibniz’s various works in cybernetics.

the modelling of life between information and mechanisation – these two histories are interlinked, according to an organicist basis that deems that life can be equivalent to mechanisation, or later in cybernetics to information. The abstraction away from context and materiality has allowed information to be treated as some form of *neutral* communicative tool that does not need context to be interpreted or made *sense* of.

The ability to equate machines across organic and nonorganic matter, permits what Hayles (1990) identifies as an underlying sense that information is the common denominator of systems – in ways similar to Leibniz and Descartes’ mechanist, organicist thinking. This is reflected in *The Allure of Machinic Life* (Johnston, 2008) that claims philosophy has developed a sense of life as mechanical or machinic. This has permitted a discourse where malware can be compared to biological *machines* such as viruses and worms (and vice versa). Erwin Schrödinger in *What is Life?* similarly draws together clockwork and organism. But he stresses, “please do not accuse me of calling the chromosome fibres just the ‘cogs of the organic machine’ – at least not without reference to the profound physical theories on which the simile is based” (Schrödinger, 1945: 85–86). Though he goes to great lengths to explain the difference between the organic and (mechanical) clockwork, the tenor of comparison between machinic and organic, is conditioned by such organicist thinking. Hence, understanding machines and computation as biological, was not too great a step for early computing imaginations.

Information, through ‘first-order’ cybernetic equivalence, lost *meaning*, replaced by a goal of efficient information transmission devoid of context. This was not a simple consensus, but was actively challenged by participants, through comparisons to energy or to context-dependent information (Hayles, 1999, Chapter 3). As information becomes devoid of context, it enables an ‘efficient’ transmission from A→B but it loses *meaning* as the ‘→’ becomes a neutral variable that does not imbibe the context of the communication. This is not dissimilar to how money varies over time and space, where its context determines its value. When devoid of context, the original value of money may be lost, and new ones generated in the new contexts it becomes embedded within. For malware, this detaches its social construction from its form, and devoid of context, becomes a ‘technical’ artefact. This is important when considering how software becomes malicious in certain ecologies. How does maliciousness, as a performance of societal expectations of abnormality, become

abstracted into information, and how does one deal with these ecologies when this maliciousness becomes abstracted and reduced to mere information in certain ways? First-order cybernetics left us with abstraction, which still lingers in computation today as part of communicative network as *unpolitical* and transmitting an apparent singular, logical *truth*.

This has led to critiques in critical data studies (boyd and Crawford, 2012; Lupton, 2018) that ‘information’ is not some neutral communication but a complex embedding of contexts. Some of this was addressed in ‘second-order’ cybernetics that considered information-in-context and through its relations to different systems. As Bateson (1979; Goodbun, 2010) explores with relation to psychiatry, and Maturana and Varela (1980) develop with environments, they argued that cybernetics encompass both the organism/machine and its environment. This aligned with the ecological movements of the 1960s (Parikka, 2016, pp.231–232), and influenced Félix Guattari’s (2014) later, and influential, work on ecology. Bateson’s attention to the qualitative aspects of information (rather than the quantitative, probabilistic understanding of both Wiener and Shannon) addresses information as *difference* that permits for delineations between different forms of information. This is dependent on recursion, which Hui (2019, p.132) claims is almost equivalent to ecology – and I call *contextualisation* – that permits not a singular feedback system but an interaction of systems that move holistically. As Hui (2019, p.136) continues paraphrasing Bateson, “[i]nformation is a difference that makes difference, only because it is both contingent and recursive.” In similar ways to how Gilbert Simondon (2017) uses information to talk of difference through the intensity of processes in individuation, we can take from a second-order cybernetics something useful. Information is about *difference* rather than one purely of quantitative probabilities, and that is where *value* can be attributed.

Yet, cybernetics also grew at a time of the Cold War where computing was developed in the “military-industrial-university system” (Parikka, 2016, p.209). Due to the highly interdisciplinary nature of cybernetics, and its adoption by militaries and defence, many of the imaginations of first-order cybernetics sit at the heart of contemporary computation unlike second-order cybernetics, more out of incidental omission rather than as a deliberate effort at exclusion. As Parikka (2016, p.214) says, these imaginations fed “security interests [that] have been at the core of cybernetics and computers since their inception; first, national defence interests, then increasingly also business and corporate interests.” These

demonstrate interlocking themes across security, cybernetics, and computing from its outset. We are imbricated in the collaborations between military, industry, and government that are brilliantly detailed in the present day, both visually and textually in *Ecologies of Power* (Bélanger and Arroyo, 2016). Therefore, any effort to develop a computational ecology must reckon with the histories of cybernetics and computer science in their emergence from a military-industrial nexus.

These early movements in both organicism and cybernetics that I have briefly outlined, are vital to understanding how computation, life, and malware are understood today. These came together in early understandings of ‘viral’ computation – where organicism, the equivalence through analogy between organic life and machines with first-order cybernetic information – become tied in understanding malware. There are direct connections between the experimentation of cybernetic principles such as von Neumann’s ‘self-replicating cellular automata’ (feedback machines that operate ‘autonomously’) to computer viruses (Cohen, 1991). These self-replicating cellular automata, form simple, individualised feedback systems, that were first presented as a lecture in 1948 (Von Neumann, 1961; Von Neumann and Burks, 1966). They are able to move, replicate, and sometimes to cease existence – and appear to replicate ‘life’ that Hayles (2005, p.173) traces through into Deleuze and Guattari’s work on the rhizomatic and becoming. Therefore, cybernetic architectures also impact on contemporary thought. Both organicism and cybernetics form a foundation of what I detail in chapter four as the *patho*-logical – a logic that combines these two traditions with discourses of medicine, the military, parts of contemporary thought, and in contemporary malware practice.

Ecologising Computation

Julia: ‘If he’s [the *Tau* ‘Artificial Intelligence’] so advanced, why is this the only house that he is running?’

Alex: ‘He’s an early version of the project I’m working on now, given the wrong information, he reacts erratically. I need the AI’s behaviour to be predictable, controlled. Control the flow of information, you control the behaviour, Tao included.’

Julia: ‘How do you do that?’

Alex: ‘I keep him disconnected from the outside world.’

■ (A conversation in the 2018 film, *Tau*)

Instead of tracing an organicist or cybernetic understanding of communication, I aim to develop an understanding of computation and malware through ecology. In a qualitative way forward, underpinned by the logics of computation, and in step beyond Gregory Bateson's understanding of qualitative, contextualised information in second-order cybernetics, I question how thinking ecologically about computation and malware can deconstruct questions of intent and authorship. This questioning, of the importance of context and ecologies of interaction, are also present when the young AI expert, Alex, in Federico D'Alessandro's film (*Tau*, 2018), slowly loses 'control' of the AI 'consciousness' *Tau* that he built. As *Tau* finds out about the world outside of Alex's home through conversations and interactions with Julia, whom he has imprisoned for his AI project, *Tau* extends and stretches the logics Alex has built. What happens then, when ecologies are embraced for an understanding of computation? Do we similarly lose control? In this section I explore what thinking ecologically could mean for computation; and work through computation's intersection with security and space to establish productive avenues of exploration.

But first, a delineation between *technologies* and *computation* is required, specifically because computation's organicistic and cybernetic histories might otherwise point us towards a very different (Foucauldian) way of thinking about politics. As Foucault said in *Society Must be Defended*, there was "a whole system of surveillance, hierarchies, inspections, bookkeeping, and reports" (2003a, p.242) that underpinned disciplinary societies. However, these are not analogous to contemporary digital, electronic computation due to their capacity for choice. Computation is not just a technique, but a recursive process of interpreting signs extending beyond the human. Different technologies have processed masses of data for the purposes of security and discipline for centuries. Yet, seeing computation as only a technological, 'automated' extension of human authorship fails to appreciate that computation has agency exceeding our intentions. This is why an interrogation and exploration of computing ecologies is so important. This is not to say that ecologies of *technologies* are unimportant, but that malware, as an expression of computation, requires an understanding that moves it beyond being a tool of human intent.

Thinking briefly about software here might be helpful. Software have been routinely referred to as performative both within and outside of geography. However, software is in itself

slippery. Some studies of software, such as *IO PRINT CHR\$(205.5+RND(1)); : GOTO 10* (Montfort et al., 2012) trace a never-ending program in intricate detail, claiming that, “like a diary from the forgotten past, computer code is embedded with stories of a program’s making, its purpose, its assumptions, and more” (ibid, 2012, p.3). Software could be said to hold the cultural and social resonances of its production, yet there is also an emergent, performative aspect, where “software is defined by its motion and rest, speeds and slowness, but also its affects, i.e. its relations with other bodies” (Parikka, 2010, p.124). Hence, there is a double bind at work when discussing software, one of its past production brought into confluence with its present moment of performance, that combine to determine software’s many possible futures beyond those anticipated by its past. There are events in moments of execution that are not pre-set. Yet these possible futures are not unlimited either, they are constrained by computational logics, the limited ways in which signs can be translated; they sit somewhere in-between. In short, software, like malware, are performative in ways that exceed (within constraints) author intention.

The question of how exactly software interact with humans, such as hackers, is an important one for understanding how software expresses agency distinct from human intent. The hacker, as popular culture informs us, is typically a single white male, hooded in dark clothing, surrounded by screens in a darkened bedroom (Holt et al., 2012). *He* is the one most often seen as the agent of malicious attacks. Yet, I do not engage with those who *code*, *write*, or *author* malware. I leave this important work to others. I am instead concerned with how malware are also an active agent in the process, whose actions may shape the emergence of the ‘hacker’ figure as well as vice versa (that I discuss in further detail in chapter four). In other words, I am interested in how malware and hacker are co-produced in ways that have political effects. In *Protocol*, Galloway offers “code [as] a language, but a very special kind of language. *Code is the only language that is executable*” (Galloway, 2004, p.165). The distinctions between code and software are fuzzy at best, but can be summarised as thus, where code are the building blocks to the structures of software (which require systematic routines, and structures¹⁰). Yet to claim that it is a language that is executable sets the wrong tone. This suggests that it simply converts, without change, the language of the author, and neglects the

¹⁰ For example, a code could be a singular line of executable language that performs an action – such as downloading other code or software, changing a particular registry, or accessing a computer. However, software requires structures, sequences, repetitions and so on such as a Word processing software, which performs a task dynamically.

complex translations that occur between the different layers of cyberspace. The hacker is important, but through the performativity of software, computation, in different ecologies mean that the intent of the author is never wholly transferred without transformation, not unlike in ‘Chinese whispers.’

Machine Learning and Algorithms

It is only the recent attention to machine learning algorithms that clearly express the agential properties of choice that are a foundation to computation. Algorithms are software that are both logically emergent but also exhibit cultural and societal histories, and in machine learning through learning data. Algorithms can be present in both malware and the technologies used by MALs. It is in the latter where machine learning (ML) algorithms have become more extensively used, and in their endpoint detection products. These ML algorithms ‘learn’ from big data how to separate, configure, and categorise. As Safiya Noble’s *Algorithms of Oppression* (2018) and Orit Halpern’s *Beautiful Data* (2015) describe, there is a bias expressed through algorithms and the learning data they are constructed by; whether through gender, race, or class. Or as seen in sustained engagement by Lucy Suchman on the performativity of bodies and things, in *Human-Machine Reconfigurations* (2007). However, is it possible to explain machine learning algorithms purely through the biases of their authors or learning data? Louise Amoore in her book, *Cloud Ethics* (2019, forthcoming) explores how a particular flavour of ML, the neural network algorithm ‘learns’ attributes that often exceed an author’s awareness or even explicit knowledge on how to intervene when a particular outputs are produced. This lack of awareness leaves us with an irreducibility in computation between authorship, ‘input’ learning data, and broader ecologies.

When a neural network algorithm recognises objects or generates a probability to render a person or software as a ‘threat’ (Amoore, 2018, 2019, Forthcoming) an author may not be able to understand how that result is produced due to the recursive process of learning that occurs (Hui, 2019). Machine learning algorithms can be supervised or unsupervised. In the former *features* (that structure data into parameters) are rules-based, whereas in the latter, features are defined by the algorithm itself. Or, in other words, algorithms create their own forms of classification and categorisation, and as we know from Bowker and Star (1999), classification is a distinctly *political* process. This produces effects, particularly in unsupervised learning, where emergent points of connection are made by an algorithm. Although the weights, and

dependencies between a neural network's 'layers' can be adjusted by its author, they do not necessarily lead to a linear output change. This is important when thinking of computation ecologically, where software and algorithms are not simply about learning data, but about how software performs according to certain environmental variables and mathematical logics that cannot be known in advance, and at the same time constructing and modifying these same environments.

Recent developments in machine learning by endpoint protection providers (that have a substantial, longer history in algorithmic practice that I detail in chapter five) make the interrelation between software and algorithms important. As Aradau and Blanke (2016, p.8) indicate in a discussion on the construction of *feature space*, "it is not simply a connection or network, but an understanding of similarity and difference based on geometric distance." This difference should not be confused with what I detailed as Bateson and Simondon's understanding of information as difference – but one based on a quantitative *deviance* from one another. Bias is *always* required for an algorithm to operate, to make choices; with various weights applied to develop a sense of similarity and difference. Difference is required for information to be produced – otherwise there would simply be no information! In security, then, algorithms "reorient the embodied relation to uncertainty, so that human and non-human beings are constantly attuned to novel events and features in their data environments" (Amoore and Raley, 2017, p.5). By this, the generation of feature spaces, their biases and weights, according to data environments – that condition the production of big data – allow for a twisting, bending, and new coagulates, where information, and value, are co-produced ecologically between human, environment, and algorithm. I argue this relationship is not unique to machine learning algorithms but to all computation at varying gradients. Information is not simply transmitted without context, but this context, the materiality of computation, is what leads to translations, and thus emergent co-productions and performances.

Balzacq and Caveltly claim, in one of the few non-technical works to actively understand malware, that "the 'goodness' or 'badness' of software cannot be determined before said performance *and* its interpretation because it always incorporates a range of possible becomings in its code" (2016, p.182). This suggests something is happening between *both* computation and social expectation – malware emerges at the intersection of computational

execution and *our* expectations of what software should do. Malware are co-produced by different elements that combine cultural and social biases, weights of algorithms, human expectation of ‘normal’ software, and the subjectivities of analysts in malware analysis laboratories. Ecologies make it impractical to identify malware simply by attributing it to a malicious author. The performative aspects of software mean data environments (Parisi, 2017), social environments, and the *agencies* of software must be considered. This permits an emergence of something that is not solely a production of an author, nor completely independent of them either.

In extending away from mechanism to computation as a distinct materiality, thinking of software production as a form of artistic practice can be productive. Artists may create art, but they are not in control of its social and political implications. Hackers, though frequently perceived as threats to the state, have also been credited with artistry. Software is not produced through a purely technical process; but is infused with creativity. As Heidegger develops in *A Question Concerning Technology* (1977, p.13), “*technē* is the name not only for the activities and skills of the craftsman, but also for the arts of the mind and the fine arts. *Technē* belongs to bringing-forth, to *poiēsis*; it is something poietic.” *Technē* forges art and technology. So, seeing both the production and detection of malware as an art can assist in moving beyond malware being seen solely as a mere technological extension of malicious human intent. As Grosz (2008, p.7) notes, “art proper, in other words, emerges when sensation can detach itself and gain an autonomy from its creator and its perceiver, when something of the chaos from which it is drawn can breathe and have a life of its own.” In taking Heidegger’s critique of the contemporary condition to heal the chasm between art and technology; software’s position as both a product of art and technology allows it to escape a mechanist vision, and instead orientate it towards performance.

Malware as a form of art reasserts the importance of its broader social, political and environmental ecology. Technics, the intermingling of art and technology, are not some form of exterior, but are integral to our everyday lives. While computation has different qualities to other mechanical machines through the processing and translation of signs and in the retention of memory (Kinsley, 2015; Stiegler, 1998), this does not necessarily mean that there is a fundamental distinction between computation, humans, animals and plants; differences are more of degree, rather than kind. Technics and humans are not two distinct categories,

we can “think of technicity as the ways in which humans and technology mutually co-constitute one another in an ongoing formation of associative milieus” (Kinsley, 2014, p.372). Technics should be a central concern of geography and international relations, as much as in philosophy, to allow an engagement with how computation is integral to political ecologies. Malware are not distinct but are intimately wrapped in ecologies of computation and societal expectation; being shaped, and shaping, alongside multiple others. But taking technics seriously, as participants with different qualities in the construction of cybersecurity and elsewhere, requires an assessment of their role in constructing our politics.

Ecological Thought

Comparisons in computer science between malware, and software generally, to ecosystems are, and have been, insufficient (Dunn Cavelty, 2013, p.108). These are frequently modelled on natural systems that in turn have been critiqued as misunderstanding microbial life (Greenhough et al., 2018; Hinchliffe, Butcher and Rahman, 2018). Ecosystems have also been twinned with a *patho*-logical focus on ‘artificial’ life (Channell, 1991; Forbes, 2004; Stahl, 2014). Cohen, who first defined the computer virus, claimed that “we have created a vast number of new ecosystems in informational form” (1994, p.31). This assessment is based on thinking through computation in its second-order, contextualising, cybernetic sense; that embraces interaction but not one that lets go of totalising formal systems themselves. One paper in computer science that coins ‘malware ecology’ (but in a very different way to me!) explores the “complex interactions between malware in the wild [outside of laboratory conditions], including parasitism, predation, facilitations, and commensalism” (Crandall et al., 2009, p.101). Yet this insight is restricted by its technological and mechanistic focus – based on analysing malware only in context – rather than the broader, explicit recognition of a potential politics of malware, participating in computational ecologies that I attempt in this thesis. An ecosystem approach to computation assumes that there are *only* quantifiable, contextualisable, formal mechanisms in which malware can be understood.

Ecological thought on the other hand is by no means new, even though it has had a recent renaissance (Morton, 2010). A key text that informs perspectives through and beyond

philosophy and geography¹¹ is Félix Guattari's *The Three Ecologies* (2014). This 1989 piece proposes three ecological registers; on environment, social relations, and human subjectivity. Through working on three registers, Guattari establishes what he calls an *ecosophy*. This is a sustained attempt to draw together different parts of a project that developed independently and in collaboration with Gilles Deleuze. Guattari argues for 'human' subjectivity as we must have a 'mental ecology' rather than thinking as one; there must be a plurality of subjectivities to enable alternative ecological strategies to be developed. This inspires how I move forward with ecology, as malware are no doubt constructed, partially, through their relation to our subjectivity (humans) – so that we can even categorise something as malicious. It is an important step, to expand ecologies to plural, potentially antagonistic, humans and nonhumans.

This ecosophy admits actors that are not only at the behest of human subjectivity but also extended and withdrawn simultaneously. As Guattari says, ecologies;

“proceed through the basis of human and even non-human temporalities such as the acceleration of the technological and data-processing revolutions, as prefigured in the phenomenal growth of a computer-aided subjectivity, which will lead to the opening up or, if you prefer, the unfolding [*dépliage*], of animal-, vegetable-, Cosmic- and machinic-becomings.”

(2014, p.25)

Hence, I arrive at a complex understanding of what ecologies are – shifting coalitions, environments, and unfoldings of different agencies that have led to computation having such great effect on the world. Computational ecologies do not attempt to somehow emulate organic life, in an organicistic way, but to depict the distinct and alternative ecologies that emerge from interaction between different forms of computational and noncomputational agency. The recent interest in ecology in philosophy has become clear in the edited book collection, *General Ecology: The New Ecological Paradigm* (Hörl, 2017). This includes work by Elena Esposito, Luciana Parisi, and Bernard Stiegler who offer a reading of technology and ecologies that informs my use. Parisi's (2017, p.79) contribution provides a critical countenance that the becoming-environmental of data does not mean it is comparable to

¹¹ For example, this was part of two connected sessions at the American Association of Geographers 2017 Annual Meeting in Boston, USA on “Ecologies of Toxicity” (AAG, 2017, pp.358, 371). This has since been followed by a book collection, *Why Guattari? A Liberation of Cartographies, Ecologies and Politics* (Jellis, Gerlach and Dewsbury, 2019).

nature. Equally clear, Esposito (2017) contends that the debate around ecology is different from living beings and their relationship to their environment, that is dependent on thinking from the second half of the nineteenth century (in itself informed by organicism). Instead it is about moving beyond perceiving ecology as the disparity between order and disorder (in the way that Wiener interprets the movement of systems to higher states of entropy in first-order cybernetics) to one based on change and creativity.

This change and creativity however should not be unlimited and universal, as Zapf (2016, p.138) argues in relation to cultural ecology. As he states, some versions of ecotheory have attempted to abolish all boundaries, but this neglects the differences and boundaries between material and cultural semiotic layers. Maintaining difference and resistance for an ecological politics to emerge is important. Ecology does not include everything but incorporates formalised structures and fissures into the political. It allows for a radical openness but not one that exceeds computing capacity, that incorporates its mathematical logics, but does not close it down to choice, agency, and the potential for something new to emerge. This is not dependent on naturalistic comparison, but appreciative of the traction of the term, 'ecology,' to offer something to those across disciplines. So, although I do not wholly subscribe to a Deleuze and Guattarian approach of becoming and assemblage due to my attentiveness to formal, logical systems, Guattari does open, through the three ecologies, the potential for ecology to exceed cybernetics. In short, by thinking ecologically in this way, I wish to appropriately bound my work according to the capacity and the fissures, limits, and logics of computation – my emphasis is different in direction rather than an explicit split. It is about moving through contextualising computation like others to bring together different environments, interactions, structures, and choices under the rubric of ecology to challenge mechanism – informed by organicism and cybernetics – so that malware have the space with which to become a political actor that extend and twist human intent.

More-than-Human Computation

Questions of more-than-human agency are crucial in assessing how malware become political – as I argued in relation to the hacker and intent earlier. I develop a malware politics through three substantive ways through and beyond what I see as contextualisation: i) through new materialisms' approach to the political agency of things such as in Jane Bennett's *Vibrant*

Matter (2010), ii) through a reconsideration of the role of cyberspace, and iii) through the construction of computational cognition (and choice) through cybersemiosis in the recent work of N. Katherine Hayles (2017, 2019, Forthcoming). Through my interest in malware ecology, I hope to complement and extend the work of new materialism through the critiques of Hayles; to introduce choice as a complementary construction of the political and as central to my rendering of the ecological extending beyond affect and power. Work in geography on metallurgy and pipelines (Barry, 2010, 2013), gases (Forman, 2018), governing border circulations (Glouftsiou, 2018), and Ash's *Interface Envelope* (2015) all productively work with non-organic materialities in a 'vibrant' materialism. This movement both within and outside of geography, has led to a greater appreciation of materials as important for the constitution of politics (Aradau, 2010), or as part of a broader materialist and affectual return (Develennes and Dillet, 2018; Saldanha, 2006; Tolia-Kelly, 2013; Whatmore, 2006) – yet none explicitly try to reconcile computation as a distinctly different kind of political actor.

Bennett, in *Vibrant Matter*, claims it is not easy to demarcate the 'big agency' of humans or the 'small agency' of worms and that "the political goal of a vital materialism is not the perfect equality of actants, but a polity with more channels of communication between members" (2010, p.104). She deploys 'political ecologies' to think what an opening-up of the *demos* may look like by putting John Dewey (processed through Bruno Latour) in conversation with Jacques Rancière. Drawing on the latter, she argues there is a materialist politics that permits the understanding of affect, in that it changes the way things are seen, "it overthrows the regime of the perceptible" (Bennett, 2010, p.106). Materials can *have an affective capacity to influence and inform politics*. Bennett is also clear that it is not only 'vital' objects that matter; "persons, worms, leaves, bacteria, metals and hurricanes have different types and degrees of power, just as different persons have different types and degrees of power, different worms have different types and degrees of power, and so on depending on the time, place, composition, and density of the formation" (2010, p.108). There are certainly different types of power, but this does little to consider the distinctive qualities of different kinds of material actors, such as the interpretation of signs and choice in computation. I am more interested in how a politics is formed by things making choices, something not really captured in Bennett's work.

Furthermore, I suggest different forms of nonhuman actor are also distinguished by limitations which shape their political influence. New materialisms have become a dominant mode of thought that include Barad, Parisi, Braidotti, Grosz, and Parikka. I follow a detailed, and welcome critique by Hayles in *Unthought* (2017, pp.65–85 (Chapter 3)), that new materialisms are heavily influenced by Deleuzian¹² concepts of becoming, deterritorialisation, and the rhizomatic; with an emphasis on non-linearity and unpredictability such as in Parisi's (2013, 2017) work on algorithms. Barad's (2007) work on quantum physics and materialism focuses on similar avenues, albeit drawing instead on the work of the physicist, Niels Bohr, rather than Deleuze. Hayles' critique suggests that there is a lack of appreciation of the *limits*, or *resistance*, to the expanses of becoming and deterritorialisation. Likewise Ian Klinko (2019) has recently written in the journal *Political Geography*, that 'vitalism' often does not attend to ideologies or power structures that condition war through his reading of Friedrich Ratzel – though it could be argued that Guattari, for instance, is attentive to these issues, it does raise the need to be more critically attentive to how human intention and responsibility intersect in more-than-human ways. Crucially, for my interest, Hayles asks, when turning to formal structures, why does scientific, formal modelling work so well (2017, p.81)? She contends that "leaving aside emergent results... each technical object has a set of design specifications determining how it will behave. When objects join in networks and interact/intra-act with human partners, the potential for surprises and unexpected results increases exponentially" (2017, p.84). Thus, how do we resolve human intent and 'responsibility' with computational agencies? I believe Hayles attempts to do this through acknowledging the limits and 'continuities' (Latour, 2013) through mathematical calculation, structures, and their limited capacities to act. Yet, she embraces new materialisms' focus on the capacity for things to 'overthrow the regime of the perceptible' within *ecologies*; therein producing 'surprise.'

As I will argue throughout this thesis, the capacity for computation to make choices are not determined solely by human agents. We cannot fully relate and understand (malicious) software; there is something irreducible to computation. I follow Thrift's (2008, p.13) uneasiness about the prospect of losing a connection to a human subject entirely however. Here I follow his maintenance of a 'minimal humanism' to ensure that we understand how

¹² No more so than with Félix Guattari, in particular in *A Thousand Plateaus* (2013).

malware are still tethered to human expectation. I do not wish to lose the subject (in similar ways to Guattari's ecosophy). Human subjects, in their multifaceted, complex arrangements, are essential to understand how *malware* comes to be understood as malicious in the first place. Consequently, I intend to highlight the imbrication between malware and human, to understand a plurality of agencies that exist around us. That is why I think that embracing a 'more-than-humanism' is more appropriate. If we are to take computation as an actor seriously, it is important to recognise they are constructed by humans but exhibit sign-processing and choice beyond us, exercising agency when they encounter and construct new environments and social relations.

However, where and how does this choice permeate and take place? I argue that, through cyberspace, we can glimpse at spaces in which choices can be made, which are not reliant on imaginations of a flat, terrain-less plane. Within geographical debate, however, there has been a reductive view of cyberspace as immaterial (Kinsley, 2014; Ash, 2015; Wilson, 2018) and 'out-there' (Graham, 2013). As Kinsley articulated in a broad semblance of the debate, "geographers are well equipped, theoretically and empirically, to discard the notion of immaterial 'cyberspaces' and to conduct more nuanced and careful studies of contemporary digital geographies" (2014, p.378). Yet, the emphasis on the 'digital' risks reducing computation to a particular material, semiotic regime. I instead use cyberspace to orientate to understanding processes between different layers of computation that include data, machinic instructions, sensors, software, code, and so on. This is not to say that all work in digital geographies perpetuate a reductive approach – as work on methods (Leszczynski, 2017; Duggan, 2017), friction (Rose, 2015; Ash et al., 2017), and an earlier editorial piece by Martin Dodge on 'cybergeography,' have shown (2001).

However, the digital is but one actualisation and materiality of computation. Other forms of computation exist in quantum computing. These use quantum mechanics and *qubits*¹³ that are significantly different to digital binary, that could permit greater speed, efficacy, and

¹³ Qubits can represent conventional binary of 0,1 however they can also exist in a mixed state of *superposition*. By this, they can be 0 and 1 (including all points in between) at the same time. Due to this quantum property, there can be digital states that exponentially increases processing power. As two qubits can represent 00,01,10,11 at the *same time*, it means that the more qubits, the more a computer can process simultaneously compared to conventional (binary) computers that can only process one state at any one time. So, the potential processing power of a qubit is 2^n , so that the common 64-(qu)bit computer could process 2^{64} possibilities simultaneously! For a more detailed overview, see Mermin (2007).

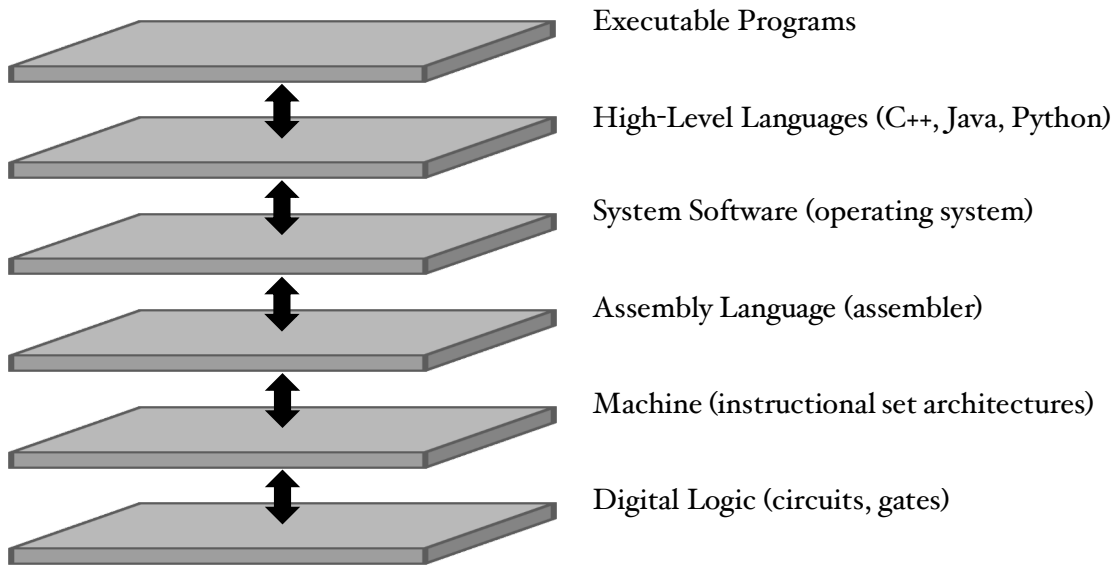


Figure 1: A simplified depiction of the 'layers' of computation that form the materialities and spaces of 'cyberspace.' Adapted from a presentation by Hayles (2018). Author's own image.

security potential than available to conventional digital, *Turing*, computing in the future (see Rieffel and Polak, 2011 for a gentle but rigorous overview). Quantum logics point to an alternative materiality to computation which is not restricted to the digital. Hence, I prefer to talk of cyberspaces where computation is more than the digital. With a more expansive term; it can be a category not tied to a particular materiality that can consider mathematics, calculation, algorithmic decision-making, and computational choice, which I do not believe are easily satisfied by various 'digital turns' alone. This form of work includes those who work at the interface of security and spatiality including Louise Amoore (2016; Amoore and Raley, 2017; Amoore and Hall, 2009; 2011), Marieke de Goede (2014; de Goede and Randalls, 2009; de Goede, Leander and Sullivan, 2016), Claudia Aradau (2010; Aradau and Blanke, 2016, 2018), and Nathaniel O'Grady (2015, 2016, 2017).

What cyberspace allows is a perspective to understand the particularities of computational materiality without tying it to a particular materiality itself. In Figure 1, I present a very basic outline of some 'layers' of current digital computation. Here we find movements through these layers between digital logic, to assembly language (the lowest human-readable language), to higher-level languages (which do not require an author to have an explicit knowledge of the particular computational architecture), through to executable programs (such as a word processor). There are many more layers, dependent on the actualisation of a particular form of computation – and this depiction neglects the 'inputs,' 'outputs,' and variances between different computational, informational, and social environments, and how

the translations and interpretations of signs between layers occur. But crucially, there is a space in which signs must be translated by various layers in order for computation to logically function. It enables a space where signs between different layers are shared, translated, and interpreted. The digital is but one mode of engagement with computation – and cyberspace is another which I think geographers and others should re-engage with to understand choice-making, and thus politics, as I will use in this thesis. That is, cyberspace is a space of withdrawal, a space in which politics is possible, and thus understanding how these layers interact both within computation and beyond could be a productive avenue for future engagement. Cyberspace is not “out-there” or “immaterial” but an alternative spatiality, where computation processes and interprets its ecologies. It is not about ‘entering’ an ‘ungridded’ place – but is full of interpretation, choice, and fissure. Without varying spatial understandings of computation, it is similar to saying that we, as humans, are simply organic, carbon processors, and that this essential quality means we could understand everything about us by abstracting from this basis. Clearly, this isn’t true.

However, despite my reservations around ‘digital turns,’ in geography and elsewhere, there is much work in the former to assist in thinking through computational ecologies in a more-than-human register. For example Kitchin and Dodge’s *Code/Space* (2011) and Thrift and French’s discussion on *the automatic production of space* (2002) provide a succinct and eloquent account of how space and code are co-constituted. In the former, Kitchin and Dodge claim that software is embedded at four levels of everyday life; as objects, infrastructures, processes and assemblages. This is a useful way to think about the spatiality of software, and how it has become part of the mundane and banal parts of everyday life. In the most explicit articulation of their thinking, they note that when code breaks down, the airport check-in area transforms to a waiting area, demonstrating how important understanding the connection between spatiality and code can be. Especially around expectation and how malware may be understood similarly. Recent work presenting queer and feminist perspectives (Cockayne and Richardson, 2017; Elwood and Leszczynski, 2018) also offers critical attention to how digital technologies are differentiated according to gender and sexuality, and how normativity can be central to how people both experience and are constructed through computation. And the extensive work on big data and algorithms in the discipline (Crampton, 2015, 2016; Amoore, 2009; Kitchin, 2017; Amoore and Raley, 2017), or with drones (Gregory, 2011; Shaw, 2016a), provide enriching developments on the relation of space

and computation, with warnings of what a positivistic approach to the study of big data could do for critical, qualitative research (Graham and Shelton, 2013). However, any rethinking of the relation between geography and elsewhere with cyberspace, depends on a deepening of our exploration of cyberspace's impact and use – so that its geopolitical affects and the ways it becomes deployed in cybersecurity can be understood. Previous digital geographic research can here continue to provide strong theoretical foundations on fissure, friction, and the performativity of space.

Yet cyberspace did not arrive at its current (non)usage without significant influence from previous understandings of computation, information, and thus how space was interpreted. As I go into further detail in chapter four, the history of cyberspace itself is constructed on narratives of control and sweeping vistas of gridded uniformity – which in turn influenced geography and the social sciences in how they thought of cyberspace. These perspectives of cyberspace, emergent from the coining of the term in the 1980s, are in turn informed by organicism and cybernetics. So, in calling for a reconsideration of cyberspace, I open up computation as the mediation of different layers and space for interpretation, and at the same time, also an explicit movement away from, but staying with, the legacies of earlier computing metaphors.

Cognition and Agency

If we recognise computation as having political agency, what form does that agency take? How much power does computation have to shape political worlds? Hayles uses cognition as a construct to critique new materialism on its lack of capacity to distinguish between power (or rather as I term, *affectual capacity*) and agency (*political choice*). New materialism, she argues, “leads to a performative contradiction: only beings with higher consciousness can read and understand these arguments, yet few if any new materialisms acknowledge the functions that cognition enables for living and nonliving entities” (2017, p.66). She instead divides objects according to their ability to be ‘cognisers’ and ‘noncognisers.’ The former make choices (no matter how limited) through the processing of signs; such as humans, plants, computing technologies, and animals. The latter, noncognisers, do not make choices (i.e. do interpret signs) such as rocks, ocean waves, and atoms. These are guided by deterministic processes and have no cognising ability. By this, physio-chemical atomic reactions do not constitute choice but follow a (fairly) predictable, modelled future.

However, they can still have political affectual capacity through producing signs that are interpreted by higher-level systems *to mean something*¹⁴. This does not exclude the great power that noncognisers may exhibit but recognises that an atom reacts according to relatively predictable laws and conditions and cannot exercise choices based on those. It is choice that enables agency, choice that allows for a more-than-human politics, a computational politics, and a malware politics.

Rather than dividing the world into human and nonhuman parts as a starting point, it makes more sense to understand various gradients, between cognitive and noncognitive actors. The former make choices and thus have gradients of political agency. I admit this is a rupture of what *constitutes* politics. For me, choice, not affectual capacity (power), is that which conditions the potential to perform politics. Where choice then emerges is essential. Hayles says “like humans, intelligent machines also have multiple layers of processes, from ones and zeros, to sophisticated acts of reasoning” (2008, p.55). She argues that there are different layers to computing, from the processor, to higher-level languages, to their organisation in software and programs (as I demonstrated through cyberspace in the previous section), in a recursive function that increases complexity that is apparent in Fazi’s *Contingent Computation* (2018). This requires a processing and manipulation of signs between different layers that at lower levels broadly follow *electro-mechanical* ways of *processing* the world – such as binary code at the level of the processor. However, between these layers is room for choice and interpretation – making computers cognisers. Malware sit in a broader ecology of computational sign-exchange, where computation’s layers (cyberspace), give space for the formation of choices based on how these layers process environments, the information they receive, and how it in turn transforms the environment, and how the layers of computation interact. I contend it is necessary to push further than Hayles’ insistence that computers only demonstrate agency through their collaboration and interaction with humans. I think that different computing environments *themselves* introduce multiplicity through their

¹⁴ Some may wish to argue that at sub-atomic levels, quantum mechanics confuses this distinction – as they do not follow deterministic principles, and I agree. However, these principles do seem to be able to be modelled or understood fairly well in a manner that if x happens, then y or z would happen without knowing which, but still being in the bounds of y or z. I think this is distinct from choice at higher semiotic levels, as though quantum mechanics may be different to Cartesian mechanics, it is not one that necessarily challenges a mechanistic process to noncognisers.

various choices that are formative of ecologies due to their performativity. That is, a political agency can emerge with humans but also without humans as a direct referent.

Recognising that computers may exercise agency (in terms of choice and affectual capacity) both in direct engagement with humans and distanced from them suggests we need to rethink the relationship. Though each computing device may follow certain logics, the ecologies in which it performs are different – leading to choices that cannot always be predicted nor modelled. This means that computation has the potential to extend beyond human knowledge or grasp through choice – for example reading signs in an environment and then acting upon these, which could be pre-defined by an author but also be within bounds of an author (meaning they do not know what will happen in different environments). In comprehending a malware politics, there are choices that can be imperceptible to humans, that I think can be at least partially explained through cyberspace’s withdrawal. This, at least partially, answers the call by Amoore that with algorithms we should (2016, p.16) “seek better geographical understandings of more than human forms of perception acting beneath the thresholds of observability.” This perception can be considered through modes of computational cognition, which extend and reach away from humans and mean a different engagement, through ecologies, is required in the study of computation. This “offers a reassessment of cybernetics and computation as central to automated systems of feedback control and logical procedures, which have exposed the changing meaning of cognitive activities, generalized from particularities (animals, humans and machines)” (Parisi, 2019, p.94). By this, computation is opening up a whole area of the role of cognition and choice.

Computers are thus in a process of *intermediation* as developed by Hayles in *My Mother Was A Computer* (2005), which Hayles claims is one of her most misunderstood concepts¹⁵. There is mediation between humans and computers through signs – that come together to generate new signs that are neither entirely computational nor human. This raises important questions. For example, how do computers and humans work in ecologies to curate the distinction of the abnormal, and therefore categorise software as malicious (see chapter six)? How are ecologies produced, sensed and productive of these intermediations, through *cybersemiosis* (2018) as “an ecology of signs”? Cybersemiosis, as developed by Søren Brier (1996,

¹⁵ This was during a lecture, “Cybersemiosis: Meaning-Making in Humans, Nonhumans and Computational Media” at the Winchester School of Art, UK, 3 May 2018. Link to event page on an internet archive: <https://perma.cc/K7ZE-73Y3>.

2008, 2014), links biology, information, and cognition. This has clear, and explicit resonance to cybernetics. However, unlike cybernetics, I wish to avoid its explicit formal modelling, as this modelling will always be too abstracted for an ecological project – as it neglects that choice and cognition mean that a deterministic modelling will always fail due to politics itself. Importantly, one may wish to class this as a more-than-human semiotics, which is detached from the linguistic semiotics of the twentieth century (Saussure, 2013 (1916); Barthes, 1967 (1964)). This is not taking this *in extremis*, but to emphasise a plurality of signs; physio-chemical, affective, computational, linguistic, material and so on.

There is much work in geography on the more-than-human that can aid in a task of recognising cybersemiosis, but it almost exclusively considers only *vital*, ‘living’ matter (Dowling, Lloyd and Suchet-Pearson, 2016; Greenhough, 2010, 2014, 2012; Probyn, 2014; Barua, 2014; Lorimer, 2016, 2012), with some exceptions (Woodward et al., 2015). However, in the pivotal piece by Whatmore (2006) on more-than-human worlds in cultural geography, she notes that “animals and technological devices have variously been used as ‘agents provocateurs’ in tackling the question of difference and rigorously working it through the specific materialities and multiplicities of subjectivity and agency” (2006, p.604). I wish to re-engage with malware as *agents provocateurs* and talk about more-than-human computational agency. I also attend to the call by Jamie Lorimer for more work that engages with antagonistic relationships. As he says, “existing work has tended towards affirmative relations and has yet to focus on examples which the interested parties – human and non-human – are engaged in lethal and antagonistic relationships in which any solution will significantly [compromise] the welfare of the other party” (2012, pp.604–605). Malware engages in *antagonistic* acts (according to us). By embracing Hayles’ work on cognition and intermediation, I work with computation, how humans exhibit certain subjectivities that are different times render something malicious, that are all constructed within, certain ecologies of environments, affects, technologies, and politics.

A Malware Ecology

Rethinking malware politics in the ways I have outlined, using choice and cognisers as a way to conceive of computational agency, helps us move from the limitations of organicism and first-order cybernetics but also the limited contextualisation of second-order cybernetics.

Instead ecology offers us new political actors and nuanced readings of the relationship between choice and affectual capacity (power) that I develop explicitly in chapter seven. Malware are then not just a technique as in previous technologies of security as Foucault would see them, in the control of humans. Computational choice informs a malware ecology that redefines who and what performs security not unlike Whatmore's *agents provocateurs* that unsettle our worlds. Processing malware through ecology then challenges who, what, and how security becomes performed in cybersecurity. As N. Katherine Hayles, in an interview on her work on cognition, says “[a]s soon as complexity and recursivity enter the picture control becomes more elusive and unpredictable” (Amoore and Piotukh, 2019, p.151) – and I interpret this not only as the emergence of affectual capacity but also the potential for greater choice to be made as the capacity of computation to add further recursivity increases. That is, the reason why machine learning algorithms seem to be so ‘intelligent’ is because they have greater capacity for a recursive process and choose from signs in order to make apparently ‘higher-order’ *decisions*. Thus, when malware exhibits a politics through its interpretation and choice of signs – such as whether a particular environmental variable exists – this may not be directed towards a goal that we would understand – as a liberal, ‘rational,’ reflexive choice – but it destabilises who and what negotiates security, and in the process is generative of a more-than-human politics.

Paul Harrison (2015, p.286) claims that there is a risk of “forgetting dying, or of forgetting finitude, and forgetting the give and take of living.” Likewise, when talking of computation – the focus is frequently on the construction of the new and the emergent. With malware, I am still concerned with those questions, but also with breakdown, fissure, logics, and manipulation of computation. It is with this, that I have restricted the sense of assemblage, forever becoming, to understand the resistance to life and the materials that structure common understandings of malware; as something malicious and disruptive. However, this does not mean that we must simply see things in a restrictive, naturalistic ecological sense; but one that functions by opening-up the world to a cybersemiotics that allows computational agency to thrive through cyberspace. Ecology also means that malware are not an ‘object’ that can be reduced to their ‘source code,’ as performance is required in computational ecologies for certain alignments with our expectations of maliciousness. This means that one cannot simply ‘follow-the-software’ or ‘follow-the-malware’ as they are

performative and constructed in many places, where choices through cyberspaces are made that are inaccessible to us.

It is precisely the emphasis on *cybersemiotics* and the processing of signs in a computational ecology that makes sense to return to *cyberspace*. This is one that is integrally connected with understanding how space comes to matter with the materialities of computation – how and where signs are made, choices are made, and how agency of computers come to the fore. By describing cyberspace as something ephemeral, it degrades potential forms of more-than-human agency and power in cybersecurity and in other arenas. By understanding how a computational agency works, we need a better vocabulary and understanding of how meaning is created in the layers of computation that may not explicitly and directly interact with humans, and thus generate meaning that is distinct to a human politics or ethics and with which makes it extremely difficult to reconcile with ourselves. That is, what is politics, should it only be decision as a highly-recursive form of choice, or choice itself, and how does choice thus translate into the political? I argue, for a true challenge to human exceptionalism, this requires us to place choice as central to politics and expose the gradients of how this interacts with us. By recognising the need for cyberspace and an understanding of cybersemiotics, it is possible to reach a greater understanding of the depth and breadth of the impact of computation on today's societies, and no more so than in understanding points of decision. This requires addressing axes of investigation, questions opened-up, and the challenges that this thesis takes up and explores. That is, how do computational ecologies radically alter approaches to cybersecurity and introduce computation as a serious actor in the production and practice of security?

Chapter Three || Researching Malware

To seriously interrogate malware's performance and negotiations of security that I explored in the previous chapter, between January and July 2017, I immersed myself in an (auto)ethnographic exploration of the UK endpoint business, Sophos. Its malware analysis laboratory (MAL), *SophosLabs*, was a space where I encountered software becoming sifted, sorted, and categorised as malicious. Here, I wrote, performed, and engaged with malware, being thrown back and forth, circulating in a vortex of aspirations that enveloped a variety of fields and spaces. This required an attentiveness to malware's specific materialities, performance, and cyberspatial withdrawal, whilst also attending to the effects of a variety of technologies, people, and software. My aspiration was to *become* with malware. I did not intend to focus too much on the humans who analyse, detect, and curate malware, although at times this slipped. However, in maintaining a commitment to study malware as malware, I treaded a post-phenomenological path (Ash and Simpson, 2016) – by attuning to its politics and power that extends away from its authors. Importantly, this is not a project that asks how to improve malware detection, even though this may be a 'good' thing: it is not hypothesis-driven. On entering the MAL, *becoming-analyst*, attending the Virus Bulletin industry conference, conducting select interviews, and rethinking four malware cases, this chapter details how I developed an ecological and methodological sensitivity to what researching malware can teach us and its implications for cybersecurity.

There have been relatively few research methods or concepts developed in the social science of cybersecurity beyond the essential goals to improve technological security, develop policy, or improve user behaviour and interaction. None of these are (auto)ethnographic however, apart from Pip Thornton's work on interacting with Google's 'linguistic capitalism' through her artistic intervention, *{poem}.py* (2018, 2019). Yet, some work on malware does not directly attempt to improve security. Balzacq and Cavelty, for instance, construct a bridge between malware and international relations through actor-network theory (ANT), and claim that malware cannot be understood before its performance (2016, p.182). This performance and interpretation permit an entry point to see malware differently. Not simply as an execution of intent, as something that can be understood solely through the study of its authors, but instead concerned with how it becomes malicious, and how this is interpreted. In taking on

board the performative aspect of malware from Balzacq and Cavelty, I embrace Esposito's critique of ANT, that we must not "blur the distinction but radicalize it, sharpen it" (2017, p.289). That is, appreciating the fissures, resistance, and breaks. Unlike ANT's 'flat' understanding, where the world can be accessible through its networks and actors, I instead wish to develop a research framework where cybersemiosis and cyberspace can be understood through ecology – with an irreducibility to computation through withdrawal that means we cannot access the world in certain ways.

As Clough *et al.* (2015, p.153) outline in *Non-Representational Methodologies*, "the "point" of the datalogical is not to describe a stabilized system or to follow a representational trail, but instead to collect information that would be typically discarded as noise." This, they claim, is a transition away from a critique of governance and economy to a now-required critique of mathematics (Clough et al., 2015, p.160). A politics of mathematics, and I think computation then, due to the irreducibilities I identified in the last chapter, is therefore not one that can be only achieved through looking at humans, nor the discursive effects of security performed through computation. Cybersecurity cannot be left to looking at actors and relations; but must incorporate more-than-human (mathematical and computational) noise and choice. Noise, and its irreducibility, then became an everyday part of this research, and something I initially did not record, or even sense and notice, but it was critical to the development of an ecological practice of 'becoming-analyst' to permit an engagement in a politics of cybersecurity.

In this methodological chapter, I explain the becoming-analyst's three main components; an (auto)ethnography, interviews and conference attendance, and four malware cases. The (auto)ethnographic exploration of the MAL provides an in-depth understanding of a site where malware becomes known, is interacted with, and knowledge of how the malicious is curated, distributed, and translated. This enabled a vertical and deep appraisal of malware; where I was able to delve into the intricacies of malware analysis and detection, to speak its language, and interact with those who work with malware. Select interviews and attendance at the *Virus Bulletin* conference assisted in expanding my view of spaces beyond Sophos, to other analysts and government organisations. This came as part of a horizontal and broadening intersection to triangulate my experience in order to engage with and pursue a rethinking of four cases through ecology. This simultaneous vertical deepening and

horizontal broadening allowed a criss-crossing to develop a particular attunement to malware in the MAL and to events and politics beyond it.

To explain my approach to an ecological rethinking of malware, in this chapter, I begin by considering how to approach the study of (malicious) software and how the MAL is an essential nexus in software's production as malicious and how (auto)ethnography can be used as a way to intervene and research these spaces. I then introduce the spaces of Sophos and the Pentagon as well as its analysts before turning to (auto)ethnography and the 'becoming-analyst' as a performative figure in which an ecological sensitivity to malware was developed through post-phenomenology. I then explain how the *Virus Bulletin* conference, interviews, and cases tease out how the becoming-analyst navigated spaces beyond the MAL and permitted new stories about malware. In concluding the chapter, I then reflect on such an engagement, the ethics of doing so, and how I recorded malware in a new light.

Approaching Malicious Software

Software has been explored in a variety of areas, with software studies and media archaeologies providing a particular entry point I wish to pursue. I draw upon a qualified version of Parikka (2010, 2016) and Sampson's (2012) Deleuze-Guattarian *ethological* approach as a starting point to research co-productions of humans and software. Following this work, a study of software through ethology becomes less about software as distinct, totally separate, or distant. Instead software is conceived of affectual, entangled, exhibiting agencies and, in my reading, able to make choices that exceed human intent. In combining ethology with Hayles' (2018, 2019, Forthcoming) work on *cybersemiosis*, it is possible to explore an ecology of signs between the layers of computation and its environments. Computers are not singular units, but are formed of different layers (hardware, assembly code, high-level languages, digital logic, and executable programs for example), allowing for a translation of signs that develop a meaning that is distinct from us. Methodologically, this requires an attentiveness to these signs, and how they are ethological, and thus how I could research malware through becoming immersed in its performance in an ecological practice.

The study of software is not only about their logical enactments but also the noise from their interactions with environments, and their effects on the world. Rob Kitchin (2017) has expressed similar sentiments about the autoethnographic study of constructing and writing

algorithms, where my interpretation would be to include its other components as well. The various approaches that he identifies are collated as; examining pseudo-code/source code, reflexively producing code (auto-ethnography), reverse engineering, interviewing designers or conducting an ethnography of a coding team, unpacking the full socio-technical assemblage of algorithms, and examining how algorithms do work in the world. There are multiple techniques in which to pursue the study of software that intertwine with both the social construction and histories behind software and computation that still attend explicitly to its mathematical and calculative modes. When approaching malware and the social construction of maliciousness, the intent of their authors must also be considered in relation to their environments, to the 'source code,' and in the ecology of practice in which this intent emerges.

Because of this, my fieldwork involved the complex task of learning to understand the various environments, human relations, protocols, software, knowledge, affect and other elements that come together to produce certain conditions in which software perform malicious acts. One such place where these coalesce is the malware analysis laboratory. Thus, as a researcher attempting to enter a space where this is done and understood, the MAL is one of the few sites where I could observe and participate in the analysis and detection of malware. Yet as Crandall et al. say, "there is an important difference between the richly diverse interactions that malware have with its real environment in the wild, and the more restricted behaviors of malware in clean and controlled laboratory conditions" (2009, p.100). This has important methodological implications – that malware does not operate in laboratory conditions as it does in the 'wild' – outside the laboratory. In developing this project, I had to be aware that the MAL is a particular space, with particular environments, where malware became ensnared in certain tools, and thus produced certain forms of maliciousness. It requires and required certain stories to be told.

Yet software remain incredibly difficult to empirically intervene with, due to their movements, formations, and locations. I take inspiration from writers such as Mol (2000), Latour (1987; Latour and Woolgar, 2013), Law (1994), Vertesi (2015), and Myers (2015) on their various work in laboratories to sense how locales become important for determining certain outcomes. Laboratories are spaces where certain variables are abstracted, and expert specialisms are developed. Results in laboratories, Latour claims, are "*extracted* from the

instruments in this room, *cleaned, redrawn, and displayed*" (1987, p.65). The MAL is a rather different kind of laboratory, however. It has its own geographies and temporalities that mean it is different to conventional laboratories. Critical here is the extensive use of big data analytics, where endpoint detection engines, that provide protection for customers, produce data on 'wild' environments that are then fed back into MALs in near 'real-time.' Malware in these spaces are not simply extracted and brought into the confines of the MAL's walls, but are in now constant communication with other MALs, through sharing data, developing knowledge of the environment *outside* to inform the analysis and detection *inside*. Though all laboratories are not isolated, big data changes the degree to which this interconnectedness applies – with new data constantly reformulating the practices of the laboratory.

Malware Analysis Laboratories (MALs)

MALs are intriguing sites in that the actions and choices made here permeate the everyday lives of millions, if not billions, of individuals and organisations through the endpoint detection engine – anti-virus – that run on the computing devices we use; often in the background, as part of a banal, privatised, security infrastructure. MALs with research teams almost exclusively operate within larger cybersecurity businesses, that are typically referred to as 'endpoint detection' businesses, that maintain them, update detection software on these engines, and retrieve data from software that is processed through these engines. Many have become household names, such as Kaspersky, Norton and McAfee. These businesses have become centres in the production of an infrastructure of privatised cybersecurity protection which evolved from the 1980s onwards that I explore in greater detail in chapter four. This makes this locale a central component of the internet economy; without MALs, there would likely not be an internet economy, or the internet as we know it today. What is deemed malicious or not, how we experience the internet, and how we protect ourselves online is frequently pulled together in the spaces of these MALs, distributed via their various endpoint detection engines, and the knowledge they gain feed into the media, policy, and government advice.

In a MAL, it is possible to work with its *scientific* apparatus and tools, follow logical deduction, and observe the limits of malicious performativity. This meant I could analyse a malware form and generate a detection, and in the process learn something about its limited possibilities to affect the world. As I highlighted in the previous chapter, new materialist

literature has a tendency to approach the world through excessive potentiality, but working with malware is about appreciating the limitations of computation, failed relations, and the inaccessibility to access and account for its actions. So, I (re)turn to the *resistance* of materials that Hayles (2014) identifies as an essential response to speculative realism and object-orientated ontology and their tendency to prioritise the imbrication between all things. During my time in the MAL, things were not easy, they were frequently hard. Resistance constantly came forth, not only for me, but other analysts and technologies in the laboratory. For example, tools frequently needed adjusting, tools for ‘containing’ malware crashed and required the reinstallation of components, the environments of analysis and detection had to be periodically renewed. Resistance is an important component of my experience, and it is this resistance that allowed for a development of an affectual awareness of the potential of something *being* malicious. Studying software is not about observing only their movement, flows, and traces but also their stoppages, failures, and resistance. It is resistance that permits a politics, a moment of choice, something that is not deterministic.

MALs are not isolated with various data exchanged including malicious samples, malicious attributes and detection rates. No MAL is separated from other MALs but is in a conversation with others through sharing data, over phone and email, and at industry conferences. They interface with wider society, shaping political, industrial, and popular understandings of malware. They do not only produce detections for their products; but visualisations, blogs, and information for customers, journalists, and governments. This curative work is communicated to others to develop an awareness of their brand and to share information. At important moments, around emergent malware events that cause sustained and deep impact, such as WannaCry and (Not)Petya, it is a time to both develop protection against these attacks *and* promote their businesses – as was seen in the growth in market capitalisation of endpoint protection businesses in 2017 (Ram, 2017). For a researcher, the MAL offers a place at the intersection of the modern economy, the development of a more-than-human politics, the production of knowledge, and the negotiation of maliciousness. The MAL is the only place where malicious production and encounter is its *raison d'être*. It is also a place where analysts have specialised knowledges of malware, which are unlikely to be found anywhere else in such density. This enables space for training for a becoming-analyst, an exclusive focus on malware, and an ability to reflect on broader repercussions that the

MAL's technologies support. This provides a locale of maliciousness that generates (cyber)securities that are unique and which I could interrogate through (auto)ethnography.

Setting-Up Sophos

On a more pragmatic front however, this research also required gaining entry to Sophos and its malware analysis laboratory. Sophos is a large, international cybersecurity endpoint protection business and vendor based in Abingdon, UK, roughly 11 kilometres south of Oxford. I came into contact with Sophos through the interdisciplinary Centre for Doctoral Training (CDT) in Cyber Security at the University of Oxford. These CDTs, created in 2013 in response to the UK Government developing its cybersecurity training needs, were designed to proactively engage with industry. As a social scientist, this provided both a point of entry and also the essential negotiation of critical distance that can be difficult to maintain. In negotiating a landscape that sees DPhil and PhD students as 'interns' and being able to demonstrate 'productive' outputs in terms of the value they will provide for industry and government, I attempted to bridge contributions to industry and to this thesis.

My first point of access to Sophos was through a CDT in Cyber Security-organised 'deep-dive' day on 9 October 2015, which are intended to introduce students to industrial issues on site-visits and workshops. After this exposure in the preliminary training year of my DPhil, I decided to further explore the potential to study malware and conducted an initial shadowing session at the Sophos MAL on 26 February 2016 after contacting Daniel through the coordinator of the university research network, *Cyber Security Oxford*. Daniel organised a session with the malware analyst, Alex, who stepped-in due to Daniel being unable to be there on the day. Daniel proved to be an invaluable contact both pre and post-fieldwork at Sophos as my gatekeeper. Initial conversations on joining Sophos occurred after some advanced attempts to go to both Kaspersky and FireEye did not bear fruit. Emails were exchanged with Daniel, who coordinated with Sophos HR to arrange a Skype interview with Peter, head of global MAL research. The interview involved explaining my interest in malware (and then with little technical expertise). It was difficult. I felt out of place under intense questioning and felt immensely inferior (which was later transformed through meeting Peter in person) – but throughout I was honest about my reasons and this must have seemed convincing. Rather than the four months I initially envisaged, my period of fieldwork

was extended to seven months, that in hindsight was the correct call on behalf of Daniel and Peter.

After passing the initial hurdles of gaining trust – that in itself was supported heavily through leaning on my identity as a *DPhil in Cyber Security* – this only went so far. In a series of emails between HR, the legal services of the University of Oxford, and I, a Non-Disclosure Agreement (NDA) was written, agreed upon and concluded on 4 January 2017. These agreements are used to protect ‘confidential’ information and was a pre-condition of my entry to Sophos, to ensure that I did not write anything that would breach commercial confidentiality. The implication is that that I am limited in discussions about the *particularities* of the tools, names, and commercially sensitive information not in the public domain. Yet, it permits me to talk in *general* terms of how the systems operate, my experience with analysts, to conduct interviews, and to form opinions on my time at Sophos. This flexibility to talk about the Sophos Abingdon MAL enabled me to position myself as ‘becoming-analyst.’ Without such an agreement this thesis would not have been possible, and I do not believe the terms of the NDA has hindered any of the content I wish to discuss. In exchange for this training and time commitment, I offered (unpaid) labour to Sophos that later would be ‘productive’ when I was analysing and detecting malware (this transition happened from April 2017 onwards). Throughout my time at Sophos, I have empathised with Law’s (1994) experience of a nuclear radiation research centre in Daresbury, UK. He reflects how he did not wish to cause any harm to the people he worked with in the laboratory, and I feel the same with those in the MAL, beyond formal ethical requirements to something else. They were my colleagues as *becoming-analyst*, and though I keep a critical attentiveness, this must not be forgotten. This means that I have had to re-read my notes in several moods to ensure I do not put a favourable spin on my experience and consider different perspectives on their actions.

I never intended to approach the MAL as a ‘workplace’, but its commercial drive became exceptionally clear as I spent more time at Sophos. This is important to reflect on – as Vertesi’s (2015) extensive and detailed work does on the apparatus and organisational work that goes into the production of a Mars rover. Or more closely aligned to my research, work by Rachel Gordon (2012; Anderson and Gordon, 2016) on the motorway control room. In similar ways to the control room, the MAL is analysing and detecting malware in order to

produce 'better' detections to sell more endpoint detection products. This came together with the functioning of a set of practices, corporate support, and the drive to compete with other MALs for better-performing detections. As part of my pragmatic orientation from a post-phenomenological perspective, it was an essential engagement with the analysts, and the conditions under which this space is constructed, paid-for, and how it is generative of a privatised security response.

The Pentagon

Sophos' global Abingdon headquarters, colloquially referred to as the 'Pentagon,' is covered in a glass façade, generating a sense of 'openness,' with wide open-plan offices. But it is full of security features that protect this corporate headquarters. It contains a moat-like feature that surrounds part of the building, which must be crossed to access the main reception. Barriers guard entry, accessible only through ID badges using RFID technology along with a four-digit pin, and the car park has a wide barrier more reminiscent of high-level security facilities. The landscaping adds to a relaxed aesthetic within Abingdon Science Park which helps to disguise and blend in these materialities of security. On entering the main reception, the corporate arrangement of (always the same) two women that I often chatted to in the mornings, helped to sift and manage entry. Within the Pentagon there was the MAL which covered only a small part of the building; where there were also two open-plan floors for engineers, sales representatives and human resources as well as on the ground floor server rooms, a post room, cafeteria, lecture theatre and meeting rooms. Each floor had two kitchen stations with free hot and cold drinks as well as free snacks brought around at 11am each day where people congregated to chat (where I often grabbed a savoury cheese scone).

The spaces beyond the MAL in the Pentagon were as important in developing a contextual awareness of Sophos and its analysts; we (the analysts and I) often chilled out at the kitchen workstations outside of the MAL, and I escaped the draining process of (auto)ethnography in the cafeteria for lunch. This was time I liked to keep for myself to reflect on how everything was going. The MAL itself however was a distinctive space in an open-plan building. It was shut-off behind grey-steel walls, with two entrances controlled by ID cards that permitted or denied entry. On entering the MAL, there were rows of desks, each with a desktop computer and three screens; to the right, a bank of computers humming, running

the various testing of new detections on different computing environments to prevent errors or engine ‘clashes’ for customers. At the far-end there was a large screen and seating for guests on ‘lab tours’ used for demonstrations of the MAL’s technologies, and centre-right, the MAL’s meeting room with video conferencing capabilities for the different teams who all worked internationally at Sophos’ various offices.

The MAL, its analysts, and technologies all constructed and performed a particular social space, with certain points of reference, certain norms and ways of communicating and interests. Initially I found this perturbing, especially given what I thought was silence but quickly found to be something else:

“I noticed that instead of the quiet that I first experienced, there is a small murmur in the background that emanates across the [MAL]. To most people this would be background noise, barely audible. I think I first missed it because it is so unnoticeable – however it does exist, and I think it’s an important part of its functioning as a lab. People do speak to one another every now and then in hushed tones, the groans of the computers spewing out that low grumble.”

(Research Diary, 23 January 2017)

The Sophos MAL was not completely silent, but full of communication, noises, and signals. The deep engagement in the Pentagon allowed this sort of sensitivity to emerge. It grew upon me during my (auto)ethnographic experience, the little sounds, glances, looks, facial expressions, all made for an understanding of the dynamics of this space. Yet, how I encountered, perceived, and worked in this space was not contained in the four walls of the MAL in the Pentagon, but with MALs in Australia, Hungary, and Canada. Typically, I would use the messenger tool, Jabber, to contact those both in the MAL and outside of it; often being the place for quick comments and sorting out if people were going for lunch or not. These parts of the MAL, not only malware, became constitutive of my experience as ‘becoming-analyst’ – where I was enrolled in various spaces, practices, and languages to imbricate myself and develop a sensitivity to malware through (auto)ethnographic practice.

The Analysts

In developing my more-than-human relations with malware, it is not possible to ignore the analysts, and others such as the quality assurance ‘technicians’ that support and generate the

MAL. MalwareBytes (2018) provides a good overview of what a malware analyst's role involves:

“A Malware Analyst is a highly specialized reverse-engineer, programmer and detective. They accomplish their task by using various tools and expert level knowledge to understand not only what a particular piece of malware can do but also how it does it. Becoming a Malware Analyst requires a large amount of focus and discipline as well as training and practice of the inner workings of computer systems, programming methodologies in multiple languages and a keen mind for solving puzzles and connecting the dots.”

Critically the role of the *reverse engineer* is important, as there was much self-identification in the MAL with this role rather than analyst. However, I prefer to maintain this broader term – as reverse engineering is but one role within a broader one which is highlighted above. I would also add that for many analysts it also included developing media for outside the MAL – either through curated datasets, visualisations, blogs or interviews. The analysts in the MAL were from a variety of backgrounds, some with PhDs, mainly from computer science such as Mason and Nathan, but others such as Joe and Jacob from a much broader background.

The analysts were a great source of support during my time in the MAL, and I had varying levels of relationship with different individuals; some I would chat to outside the MAL, others would be purely through work or on communication tools such as Jabber. They also provided the liveliness of the MAL and often engaged in political debates – such as the inauguration of President Trump, the American healthcare system compared to the UK's National Health Service, to the implications of Brexit (the UK's exit from the European Union). The sociality that developed meant that it was frequently over cups of tea, discussions and laughter over who had been 'pwned' (who had been *owned*, or hacked/attacked), and often by playing games in lunch breaks such as the interactive game, Pokémon Go, where gamers 'catch' Pokémon through an augmented reality mobile app, to Ingress (Figure 2) where players attempt to capture portals for the team to progress and to control an area. Frequent lunchtime visits to portals around Sophos were common, with several younger members of the MAL taking part in this. I was told that I must join the 'blue team' (Research Diary, 19 May 2017), and occasionally joined other analysts at lunch. This activity beyond the core activities of the MAL is how I came to bond with individuals – in a way that was not trying to garner information, but to feel part of a team and to feel as though



Figure 2: A screenshot of the Ingress game, where players work together in teams to capture 'portals' from other teams through GPS locations and performing ever-more complicated tasks. Author's own image.

I had a place. During my first few days however, I recollect the dominance of men, predominantly white, sitting in t-shirts with various references to gaming cultures. This is not dissimilar to the hacker culture that Sherry Turkle (1984) identified in her work at MIT during the 1970s and 1980s. There was no singular masculinity at play – with some individuals quieter, akin to the more conventional ‘computer geek,’ whilst others were more interested in cycling – but ultimately all had interest in the *geekiness* of reverse engineering. Occasionally, there were uncomfortable moments of implicit sexism, that may not even have been recognised by those involved, but jarred with me as a researcher, where I decided not to intervene. The analysts were not a singular group, with different personalities, and I had varying interactions over time, that reflected how ‘well’ I got on with them. These all influenced my becoming, at Sophos, and how I learnt about malware, and conducted my (auto)ethnography.

Ethnography

“What does an ethnographer do? I worried about this a lot. Partly it was a matter of measuring myself against an ordering idea – that of the ideal ethnographer. Such a creature would have made more phone calls, been more sociable, and have had a better memory.”

(Law, 1994, p.43)

Law's description of what the ethnographer should do was a constant theme across my own (auto)ethnography. I always felt I was falling behind, somehow not understanding the malware analysis process, not writing detections that were 'good enough.' In fact, there were several times when I was warned that my detections needed be 'tighter' and 'stronger.' I was paralysed in a competing combination between learning to become an analyst along with trying to engage with malware and its technicalities, trying to observe other analysts, to writing a research diary at the end of each day on the forty-five-minute standing-room-only train ride home.

Ethnography is not something that can be picked out of a textbook and routinely applied. This is both its strength and weakness. Its strength comes from developing an acute (un)awareness of certain spaces and social relations. However, one can get buried in detail, and due to the all-consuming effect of long-term ethnographic engagement, insights produced at certain times and spaces make replication difficult. This does not mean that there is no possibility for repetition, but that as difference is ever-present, specificity is lost, but over time *generalities* can emerge that transcend the singular contemporary MAL to broader questions on malware. There is never a perfect copy, and ethnographies help deal with the complex experiences I encountered, where stories provide one way to navigate and make sense of this complexity. This could be more widely applied to science more generally – as the production of certain stories, and ethnography is no different.

Within geography, Herbert called for a return to ethnography as it “is a uniquely useful method for uncovering the *processes* and *meanings* that undergrid sociospatial life” (2000, p.550). Though there has been a return in geography to ethnography, I argue these are potentially different to what Herbert envisaged. They are frequently short-term; with participant observations, interviews, and day trips now part of the 'ethnographic.' Although these are useful, I think there is still a value in ethnography that is long-term that, in various ways, becomes ingrained within communities. Ethnography does not work for all disciplines or research, as Kuus (2013) laments on her work in foreign policy – where access, and level of engagement can often be exceptionally difficult to achieve. In this, the ethnographic is not a method, but a way forward, a way of approaching research that incorporates participant observations, interviewing, and co-constitutive learning (Roe and Greenhough, 2014). This means I attend to Law's (2004) call to adopt a 'messy' approach that blends together

participant observation, interviews, shadowing, and my becoming-analyst, which appreciates that research is a reflexive process of going back and forth.

One of the few ethnographic studies in cybersecurity by Sundaramurthy, McHugh and Ou (2014) concerns a Computer Security Incident Response Team (CSIRT). They point to two salient difficulties in researching cybersecurity; first that there is a hypothesis-based way of approaching problems in cybersecurity, and second that analysts have little time for interviews and are unlikely to trust researchers (2014, pp.52–53). Approaching cybersecurity research without preconception – through ethnography – is something that is particularly undervalued and underutilised. Engaging with analysts means that to do something beyond interviewing is seen as suspect. This is likely to be for a range of reasons, but primarily a drive within computer science to produce experiments and results at pace to publish at conference which does not necessarily align with slower work that is typically done of social scientists – but with which human-computer interaction has aimed to satisfy (Cooper et al., 1995; Goulden et al., 2017). The *slow* work of in-depth ethnography may not seem to support the goal to improve device security or be of immediate fruition in improving practice or policy. None of those targets are poor – however deeper engagement with and understanding of the trajectories of cybersecurity is sorely needed.

(Auto)Ethnography

Autoethnography emerged during the 1980s ‘crisis of representation’ as a way to move from realist traditions of writing to destabilise ethnographic authority (Butz and Besio, 2009). However, this is not a singular method, but rather a retrospective label applied to a range of self-representational practices. It can refer to a more introspective approach or one that situates oneself in the practices of ethnography. I follow Anderson’s (2006) distinction that he identifies between an *evocative* autoethnography that focuses on emotions in its more literary sense, such as that advocated by Bochner and Ellis (2016), and the more realist perspective that writes things ‘as they are.’ Anderson proposes an analytic autoethnography that incorporates five strands in order to bridge these two approaches:

- 1) Complete Member Research Status,
- 2) Analytic reflexivity,
- 3) Narrative visibility of the researcher’s self,
- 4) Dialogue with informants beyond the self, and
- 5) A commitment to theoretical analysis.

This is significantly different to the evocative methods that are dominated by specific (bodily) experiences. Though I read this work with great zeal, and the value of the particularities of the place in which things become experienced – this, I would argue, is not so useful for more-than-human research. It centres too much on the human experience, and in the spirit of post-phenomenological engagement, does not take the political agency of *things* seriously enough as having affects beyond the ability for the researcher to comprehend, or that there may be generalisable, logical parts of the world that are essential for any engagement with computation. Therefore I broadly follow Anderson’s call and support his approach that the “defining characteristic of analytic social science is to use empirical data to gain insight into some broader set of social phenomena than those provided by the data themselves” (2006, p.387).

In pursuing (auto)ethnographic research, it is important to note my explicit use of parentheses. The interplay that is evident in the parentheses of (auto)ethnography is also an explicit rejection of being autonomous in my decisions, thoughts, and effects. My research is not solely an autoethnography following malware, isolated, and becoming only with it. Yet neither is it an ethnography in the sense of being focused on the malware analysts. It is post-phenomenological; acknowledging and working with the agencies of malware that always exceeded me. I reject the auto in ethnography as some sort of possible existence, where the human is centre. Cognition is distributed, and this is why malware have political agency in the world (and what distinguishes them from noncognisers). This is part of a post-phenomenological practice and sensibility that I use to support my (auto)ethnographic research. Yet, at the same time, I cannot escape these parentheses. I am human, and this will always be present in my writing – I cannot write on behalf of malware, nor understand it in anything but a humanly way. This means this story of malware is forever (and will always be) incomplete. Being open to other perspectives and agencies is one thing, but computation’s withdrawal always introduces a friction and resistance as I was becoming-analyst.

Becoming Analyst

As a human however, I had to navigate certain spaces. Due to the institutional ‘set-up’ at Sophos, I was deemed an *intern*, a term loaded with particular references to work and

hierarchy. Daniel said it didn't really matter to the everyday working environment, at least he thought. I went along with this arrangement in order for HR to process me in their systems. Therefore, it was not just about 'becoming-analyst' as a research method but also in terms of institutional alignment. As a young DPhil student, certain assumptions were made about my position – and which sometimes came back to haunt me with questions of whether I would return to Sophos for a 'proper job.' The term becoming-[something] resonates with Deleuze and Guattari's (2013) work, but I also see its resistance as I described with Hayles earlier, and how this is always limited. In becoming-analyst, I thought it possible to evoke novel ways of engaging with computation in a post-phenomenological sensitivity, that over time, through practice, enabled me to glimpse towards malware and how it became translated beyond the MAL. It allowed me to enter Sophos and to talk to analysts who may have been initially sceptical of a researcher, may be concerned about confidential and proprietary information, and question the value of long-term ethnographic engagement. As I was *working* at Sophos, unpaid, it gave me access to all the meetings I would be expected to attend, whilst also being able to speak to people outside of these. This helped to develop trust among analysts and permitted not only an in-depth knowledge of malware to emerge, but also to become a (partial) member of the MAL.

I was forever encountering resistance however. I was coming into contact not only with the materialities of malware analysis and detection, but also with the social spaces of the MAL. This most keenly emerged when negotiating my position, and my craving to be part of 'something.' I thought it best to seemingly work like other analysts full-time - roughly 8.30am to 6pm each weekday. One thing I did not fully account for was post-6pm. This emerged when Peter, a senior lab member outside Abingdon came to visit, where some drinks in a local pub were arranged.

“Many went to the pub for some sort of celebration drinks. This was scheduled for 6pm, and I stayed around but didn't have anyone say whether I was going. This put me in an ethnographic dilemma. Obviously, I want to feel valued and part of the 'team,' but I have to remember that I am in a different power dynamic to one of the average 'intern.' I could have gone – as I was formally invited, but I didn't feel as though this would be the right thing to do – without an individual suggesting I go. I think this is important to reflect upon – as I would have liked to have gone, but clearly, I have not developed that sort of relationship with these individuals. Also – am I Andrew the intern, or am I Andrew the researcher beyond the field site of the lab?”

(Research Diary, 11 May 2017)

This feeling of not knowing my place happened many times, where the boundaries of (auto)ethnography were stretched, contorted, and shifted. Later in my time at Sophos, I did go to the pub a couple of times, getting to know one analyst, Charlie, far better - initially through the use of the messaging tool, Jabber, and a shared love of Polo mints. However, I broadly tried to avoid such events to maintain some distance between myself as researcher and analyst – even though the ‘becoming-analyst’ made this hard to maintain. This came through most prominently when several members of the MAL organised a ‘leaving event’ for me, where we went bowling and had food at the American restaurant, Frankie and Benny’s, at Ozone Leisure Park in Oxford (Research Diary, 26 July 2017). This event was both a recognition of my time *becoming-analyst*, developing close bonds, whilst also being a stark reminder of my difference as a researcher, that I was only ever a temporary *intern*.

These limits and resistances must exist to maintain appropriate ethical distance between researcher/d, but also these interruptions, though sometimes confusing and surprising to myself, open up what a researcher needs – for there to be non and un-becoming in the process. As Paul Harrison (2015, p.296) writes in his work on non-representational geographies, through Friedrich Nietzsche, on being forgotten and finitude, “active forgetting is part of the health of the individual, selective repression, a digestive and metabolic incorporation, and a voiding of what cannot be assimilated.” The becoming-analyst must be simultaneously, in order to be becoming, also un-becoming. As an analyst, I was trying to remember things that happened, had to check that I was *not* becoming an analyst, and indeed there was an end goal to writing this thesis. My becoming-analyst was only ever partial. By this, I had to be constantly unbecoming, in order to continue its process of becoming. Therefore, the research required moments of not fitting in to allow a further becoming, to pursue new avenues.

My technical training was a complex experience. This involved at various points reading books, practising with training data sets, playing with different tools, information from the MAL ‘wiki,’ writing detections, understanding the *value* of contextual information (see more in chapter five), and shadowing sessions with other analysts. The learning never really stopped, but I believe there are two significant stages to the process of my becoming-analyst. By mid to late March 2017, 2.5-3 months into my time at the MAL, there was a distinct change from training and learning to shadowing and producing detections. I wrote that,

“I think my recent boredom / tiredness may be born out of the fact that I needed to move on from the training. It is quite exciting and sets up a whole new part of the ethnography where I can analyse how people work through the samples and write detections for them.”

(Research Diary, 14 March 2017)

And later that month, in a meeting with Daniel, he indicated I should start my shadowing activities (Research Diary, 31 March 2017). This transition emerged from the technical education I received, and slowly I was able to write detections for more complex malware, and able to put together *better* detections that ran quicker and were *tighter*. That is, they picked up similar variants of (malicious) software and were unlikely to detect other software.

I found the technical training arduous and gruelling. On the first day, I clearly remember being shown to my desk and one analyst, Elliot, setting up my computer, his hands whizzing away, me sitting uncomfortably on my chair unsure where to look. One comment stuck with me: “you know how to do everything on the command line¹⁶?” I timidly replied “yes” (Research Diary, 5 January 2017). I did not know anything, so I was left sitting there for the rest of the day attempting to look like I knew what was going on. I had to learn not only the basics over the next few months, but: how virtual machines worked, basic command line operators used in Sophos, how computer chips function, assembly language showing instruction-level computer processing at the lowest level of human-readable code, how to use specialist tools, learn the main attack vectors, and how to write in a proprietary language to construct detections, to name a few. The first few months were solitary, reading textbooks, feeling awkward asking questions and working through exercises that were created several years ago for new analysts. In fact, I found most of my initial time in the MAL grinding. It wasn’t until the end of my first month that I was able to say that;

“This was the first day when I thoroughly enjoyed working with malware – albeit in static form. It was enthralling to see the instructions suddenly call out to me, recognising structures and understand how a particular piece of malware operates. Truly fascinating.”

(Research Diary, 31 January 2017)

However, there were many ups and downs, and it took several months for me to become accustomed to the dense technical components of malware analysis and detection. With

¹⁶ a text-based tool that allows direct access to the computer without a mouse or other interfaces

hindsight, it is easy to say these were inconsequential, and by June 2017 I had developed a rhythm, a routine, and felt able to move around the MAL and its technologies with ease, as well as developing relationships with analysts. However, it was no easy task to keep going earlier in the process, through Oxfordshire's dark winter days.

Though my first few months may have been solitary, there was a steady rhythm of meetings that I often craved to avoid sitting in front of my computer. These included the weekly 'catch-up' with the manager of the MAL, the near-daily 'stand-up' between analysts in the UK MAL, and the meetings in the MAL meeting room including the fortnightly meetings with the MALs in Hungary and Canada in the generic detection team that I joined, monthly 'sprint' kick-offs that dealt with tickets to be addressed in the next month, and monthly 'all-hands' where everyone in the MAL (not only the analysts) looked on the past month and any trends that had emerged. My catch-ups with Daniel, as the generic detection manager, were the most beneficial out of these meetings. Daniel was someone to discuss both my becoming-analyst and identity as a doctoral student. Discussions ranged on the history of malware analysis, particular organisational processes, his close connection to Oxford, to being interested in my research and a place to sound ideas. This often happened in the centre of the first floor in an area designated with comfy armchairs to encourage informality and discussion. These times were some of the best moments where the (auto)ethnography came together and to have conversations with an experienced analyst. Other meetings, especially those with other MALs, allowed for reference points for how these places operated and expanded my experience beyond Abingdon.

The balance for the becoming-analyst in conducting (auto)ethnographic work was constantly at the forefront as I expressed very early on in my fieldwork, and which continued throughout my time at Sophos,

“that is the difficulty of learning and observing. I understand that these both go hand in hand; however, I feel as though I'm being sucked into my own thoughts with my computer and I almost lose sight of everything else. The [malware] analysis becomes everything. I've noticed that many others in the lab wear head/ear phones whilst they are working. I thought I would do this, and the office you work in is not one of multiple people but you and your samples. I have to admit that sometimes it is disappointing to be distracted.”

(Research Diary, 13 January 2017)

I was constantly in flux between wanting to engage more squarely with malware, investigate the strategies in which malware became designated as malicious and what conditioned this, to then wishing to observe and be part of the team. Clearly this evolved over time and in tandem, but the risk was that I was becoming too comfortable. I wrote about this in my research diary several months in,

“have [I] become malware analyst myself[?] I have got used to the rhythms of the office, of my commute, of thinking with malware in order to detect (or not detect as today has shown!). I think this was most evident on Tuesday when I had my catch up, when [Daniel] almost threw my position into question when he asked whether I still wanted to do interviews. This struck a chord with me that I have been thinking about for the past few days. I have become so comfortable I’ve almost lost sight that I am actively observing others and what they are doing. This is very dangerous to feel like one of them when I am not. This is the classic ethnographic trap I have walked in to.”

(Research Diary, 1 June 2017)

This resonates with Leon Anderson’s autoethnographic work on skydiving, where he talks of trying to balance the two parts of his auto-ethnography and the difficulty, and probably incessant failure, of it,

“while the plane ride to “jump altitude” is commonly used by skydivers to mentally rehearse planned jump manoeuvres, conduct checks of one’s own and others’ gear, and joke with the other jumpers, for me, it is also a time to consciously observe and etch conversations and events deeply enough in my mind that I will be able to recall and record them in detail after the jump.”

(Anderson, 2006, p.380)

My moments of awkwardness were an acknowledgement of my position, as I partially became part of the team, shared jokes and laughed, and became better at attuning myself with (malicious) software. Play was central to this. Play with the different knowledges, with each malware sample and form, with different technologies that allowed this sinking in. I frequently dreamt of malware (and sometimes still do), and I felt affinities to the structures that I can only vaguely attempt to explain – certain structures spoke to me, jumped out, that made sense to pursue more. As McCormack (2013) says, and I think it is relevant to my experience as becoming-analyst, that experience is mediated through play, it is not a transcendental moment but a transition. The transition was the build-up of hours upon hours of working with different samples, learning different techniques, the best tools to use which

all came through playing around, often not understanding, experimenting with new ways of putting things together.

Ultimately, I think this goes to the core of what a more-than-human (auto)ethnography inspired by post-phenomenology can entail, when the agency of computation as a specific materiality and actor is taken seriously. The Sophos MAL was a space where software came into contact with a variety of ecologies, from tools and technologies, analysts, knowledges and discourses, as well as data from multiple environments. As a becoming-analyst I was constantly being (re)drawn by the ecology of relations and choices I was in, forever wishing I could more exclusively be with malware. Yet, I can never be alone with malware without extensive support or bleeding of knowledges about how to analyse them. For example, it would be wrong to say that malware ever ‘recognised’ me or were ‘affected’ in ways that I could think of – its deletion on endpoints. We work on different levels of cognition, signs, and meaning so that it is only ever possible to glimpse at one another and assume we have some points of choice that align, malware and human. Connected by maliciousness, anomaly, abnormality, but unable to comprehend one another – sense in distributed ways, but never able to touch.

Post-Phenomenological Sensitivities

So, as becoming-analyst, I worked with post-phenomenology as a way to comprehend and perform with computation. This is by no means unique – as has been developed in new materialist literature and in post-humanism, but post-phenomenology has been developed to attend, at least partially, to computation. As Ash et al. (2017, p.169) state, “post-phenomenological approaches understand that objects both proceed and exceed human experience of them while also providing the grounds and means for human thought and cognition.” This is in contrast to approaches such as ‘ANT’ which seek to anticipate the world as understandable through actors and networks in a flattening ontology, although there are earlier considerations of its limits (Latour, 1996; Law, 1999), which are arguably still unaddressed. My contention is that by subsuming the world under actors and networks (as much as these are constitutive of one another), we lose the mathematical and logical boundedness that is required for computation to perform (or, in its more limited expression, execute). As Latour, himself central to ANT, considers in the introduction to *An Inquiry into Modes of Existence* (2013) through the figure of the anthropologist, that there are indeed limits;

variously collated as values and continuities that mean that there is some boundedness that ANT has neglected to fully embrace.

Post-phenomenology however requires an attunement to the world, and I think to attend to these frictions and irreducibilities. As McCormack outlines, post-phenomenology is about “how the capacities of bodies, machines or entities to affect and be affected are shaped” (2016, p.6). This requires an account of a researcher’s positionality that centres on others in an ecological practice, acknowledging the various senses and effects on a researcher’s body. When I was in the MAL, how did my bodily senses and knowledges of malware analysis and detection intertwine with the materialities of the laboratory, its technologies, and malware? What could be seen as *small* events shaped the world in which I was working within, reflected in my notes and approach I took that day. Whether that be a smile when I went to get a cup of tea, remembering a comment made by Terry to try a different method to help me out of my “rat hole,” as Daniel described it (Research Diary, 7 June 2017), to just being “absolutely knackered” (Research Diary, 22 May 2017). All these moments were small but essential to my experience and how I approached and felt affected by malware. By working post-phenomenologically, I collated notes on how malware changed the laboratory, my experience, rather than just describing how analysts constructed the MAL. My authority as a writer was not only to be reflexively noted but was distributed across a range of actors that came together in *our* (auto)ethnographic research. But not to be fooled, I am the prime interlocutor here – I am the human who writes for you, with particular sensibilities, the *ours* is not equal and is quintessentially and irresolvably political.

More-than-human geographies have provided much theory to rethink our (human) relationship with the world but Dowling, Lloyd and Suchet-Pearson (2016), in a progress report on research methods in more-than-human geography, note a limited development in praxis. There has also been a tendency to focus on animals in the study of more-than-humans in geography (Madden, 2014; Lorimer, Hodgetts and Barua, 2017; Buller, 2014; Srinivasan, 2019), with little concerning technology, unless it falls into post-humanism of some form (Rose, 2017; Seaman, 2007). Methodologically, a focus on ‘being-there’ amongst animals is valuable in the study of elephants (Barua, 2014) and in swimming with Bluefin Australian tuna (Probyn, 2014). However, it will clearly offer little to understand the running of code, or reading code on a screen as this was not this work’s focus. But, I cannot simply ‘be with’

malware. The imperative of more-than-human approaches of being-there and understanding agency are important, but they are often restrictive to *larger*, more-present animals such as elephants, dogs, and things that are tangible to us. Some work, such as that engaging with biological bacteria and viruses offers an alternative perspective about how to work with non-visible things, that experiments with approaching relations with these others (Greenhough, 2012; Greenhough et al., 2018; Lorimer et al., 2019). *Political Matter* (Braun and Whatmore, 2010) likewise develops a political approach to the study of ‘matter’ that I find productive to engage with materials in a more-than-human non-organic tenor. This begins to offer a more-than-human geography inspired approach to research an entity – malware – that does not adhere to the categories of animal, post-human, or solely ‘matter.’

Bruni (2005) shadows software and clinical records in a four-month ethnography following medical professionals, in order to open a gateway to what I term more-than-human computational research. As Bruni (2005, p.374) says, “shadowing non-humans requires the ethnographer to be able to orient his/her observations to the material practices that perform relations, and probably also to devise new narrative forms able to make that performance accountable.” Here, these records are something produced and performed into existence in a more-than-human tenor, even though he does not depict it as such. Others such as Woodward et al. (2015) speculate on visualisations of Hurricane Katrina (2005) to produce a geovisualisation that develops on Simondon’s (2017) individuation and concretisation to suggest that technologies should be seen as collaborators in its production, but emphasise that we must not lose the positionality of the human. These two studies extend computational agency with human collaborators and depict not only how researching computation can be achieved but how it reorientates the positionalities of those involved. Therefore in drawing on research in more-than-human geographies, and on post-phenomenology, on an orientation to things that exceed my experience, in order to ‘open’ them up, I seek to treat computation and malware as being able to exert effects on my body, that became embodied renderings of maliciousness, and can be related to in ways developed in more-than-human research.

Virus Bulletin Conference

However, in order to assess and complement my deep insights from the MAL, in October 2017 I attended the primary annual conference for malware analysts, *Virus Bulletin*, at the

Novotel Hotel Madrid Central, Spain. The conference is a site for sharing information on recent research, new techniques, and developing a sense of community in malware research. Jessica Johnston, in *Technological Turf Wars* (2009b), uses the *Virus Bulletin* conference for her research on the 'anti-virus' industry. Here, she uses this space to explore the industry, gain contacts, and conduct interviews. As one of the few times when malware analysts come together, rather than broader cybersecurity conferences, this was an important space to triangulate my experiences with those at Sophos, speak to different analysts, and to listen and comment on their presentations. This gave me a sense of a middle ground of translation of the activities of MALs to wider audiences – in that these were often nuanced, complex, and acknowledged deficiencies in their work unlike the more polished 'public' outputs I had witnessed at Sophos.

With a significant press gallery, drinks receptions, quiet words in corners, and a conference dinner; this was a crucial point to mingle and imbricate myself as the becoming-analyst. Attendance at this conference underscored an ending point in my fieldwork, in that I could attend technical presentations, understand and follow the information, and engage in discussions. I attended a range of talks, many of them technical, which would have been incomprehensible prior to my fieldwork. I even bumped into analysts from Sophos who were there, which provided some comfort in attending a conference alone. It assisted me in 'getting up to speed' with the variety of new research in malware analysis and how this was presented with explicit links made to malware's authors as threat actors or 'advanced persistent threat' (APT) groups. Over the course of the conference, references to environment and improving contextualisation through data were made; no more so than with reference to machine learning and algorithmic processing. At an evening drinks event, at coincidentally, *The Geographic Club*, I spoke to a range of analysts and other leading individuals in machine learning and malware detection. Over the course of the evening, and many drinks, I was passed innumerable business cards as I was enrolled into the role of prospective analyst to 'snap-up' – trained in malware analysis and studying at Oxford. The conference was very much a place for endpoint detection businesses to understand their competition and for some, smaller, endpoint protection businesses, to train their analysts.

The ability to speak to other analysts, to observe their presentations, and the tools they used, confirmed that though there may be differing techniques at different MALs, there are some

salient features that apply more broadly. These include the need to maintain a low false positive rate, the segregation and negotiation between ‘clean’ software and malware, and the growth of contextual strategies to both analyse and detect malware. One particular moment made this explicit. At the conference dinner, held on the evening of 5 October 2017, I was surrounded by individuals from across different endpoint protection businesses, when revelations of Kaspersky’s potential infiltration by the Russian state hit the email inboxes of several senior corporate employees. Over that table, it was discussed that endpoint engines that detect malware were no different across the industry to what Kaspersky used. As became clearer over dinner and the next day was that the US had accused Russia of using Kaspersky’s endpoint engine for spying. This eventually led to a ban on the use of Kaspersky’s engine and products in the US Government (Rosenberg and Nixon, 2017). Yet, in this instance and at the conference as a whole, it became clear that the fundamentals of endpoint detection engines, the core concerns of detecting malware, and some methods were generally similar to what I have experienced and worked with at Sophos. Although the way things are categorised, measured, and collected may be different in each MAL, this horizontal appraisal through the *Virus Bulletin* conference permits some generalisations, for example around logics, broad methods, and techniques that I had observed at Sophos.

Interviews

I supplemented my (auto)ethnographic research through six in-depth semi-structured interviews. With four malware analysts at Sophos and two interviews with eight individuals at two organisations, one in the UK and the other, an agency of the European Union, I explored how knowledge is translated from laboratories to broader political discourses. These semi-structured interviews lasted for at least one hour. Interviews conducted with MAL analysts developed on near-seven months of (auto)ethnographic research that included conversations, shadowing activities, and the development of close bonds with some individuals. The latter interviews with individuals outside of Sophos focused on the contemporary cybersecurity context, with particular reference given to the 2017 WannaCry and (Not)Petya malware attacks.

In developing a semi-structured approach, I jotted ideas and directions for all interviews that allowed for in-depth interviewing to be conducted in an expansive fashion. There were no doubt difficulties in keeping these within certain bounds, that of malware and its agency,

and away from the human-dominant forms of what malware are. However, these apparent distractions led to some interesting divergences on the histories of malware, how the MAL analysts themselves became analysts, what interviewees believe the trajectories of malware futures were, and the processes of their analysis or dissemination of information. The two external organisations were:

1. A group of individuals from the UK National Cyber Security Centre (NCSC), which is the public-facing organisation of the Government Communications Headquarters (GCHQ), that is the UK's listening agency, and;
2. An individual from the European Union Agency for Network and Information Security (ENISA) that the participant described as its "function is to serve the EU Member States, the governments of the Members States and other entities, such as the EU institutions (think other EU agencies, Commission, Parliament, Council, etc.)" (Interview, 18 October 2017). This includes providing information and sharing information for these bodies.

The meeting with the NCSC was conducted in a meeting room at their central London offices, and was arranged through contacts passed to me through the research network, *Cyber Security Oxford*. The ENISA interview was conducted via conference call with the interviewee in Athens, Greece, where I made contact after they presented at a session at the *Virus Bulletin* conference. The collective interviews, attendance at the Virus Bulletin conference and my experience at the Sophos MAL then informed an ecological re-reading of malware forms through focusing on four cases.

Four Malware Cases

Early in the research, it became apparent that in order to discuss a 'malware politics', I would need to explore different malware 'cases' as a way to extend and complement my (auto)ethnographic engagements in the spaces of the MAL. The becoming-analyst enabled and was a necessary (pre)condition to re-read malware ecologically. Thus, I present an ecological reading of Conficker, the Dukes, Stuxnet and WannaCry/(Not)Petya, and initially a fifth case, Dridex, that I decided not to pursue¹⁷. Though the (auto)ethnographic research in the MAL was exceptionally useful, through the use of different 'well-known' cases, it is possible to consider how my arguments around politics, cognition, and choice could be made explicit through re-reading and re-writing the dominant stories of popular malware forms.

¹⁷ This was a more pragmatic decision for the length of the thesis – and that after the other cases – did not contribute much in addition.

A focus on ‘objects’ in security has gained traction following work such as that by Aradau (2010) and others (Folkers, 2017; Geers, 2009) on critical national infrastructure and more recently in geography by Meehan, Shaw and Marston (2014, 2013). This movement, which connects to vibrant materialisms (Bennett, 2010; Barry, 2010; Braun and Whatmore, 2010) and associated movements in object oriented ontologies (Bryant, 2011; Bogost, 2012; Morton, 2010), places the object at the centre of theoretical analyses. In pursuing a post-phenomenological direction in this research, I however seek to question whether malware can be understood as an object but rather a series of materialities, practices, and performances in an ecology that translate software as malicious. By exploring malware through ecology, I aim to write *differently* about malware – that is from the perspective of their performance and an opening to explore their potential choices that extend away from the human. By this, tracing a more-than-human politics through choice and power, to position malware in a new light.

In chapters seven and eight, I follow the traces of how malware can be recognised and interpreted differently, drawing on the four cases. In ways similar to the post-phenomenological work that informs the rest of my methodological approach, James Ash’s (2015, p.25) work on interfaces identifies the crucial issue at our current juncture, that computational devices have “become black boxed tools, flat visual images, weightless representations or binary code that await activation, either by human beings who perceive or sense them or by the computer which runs code or follows a program.” I think a similar rendition of malware exists, as something ephemeral, as an existing object that moves around as a singular unit. I use the phrase ‘form’ to try and move beyond seeing malware as an object, to critically evaluate how it maintains different variances in different environments and ecologies, but still permitting many variants to be classified under a singular name. Thus, I refrain from perpetuating malware being treated as an object that could be contained on a USB stick, for instance, and argue instead that their performance between the layers of computation mean they cannot be disentangled.

In the discussion of the four malware case forms, I do not intend to provide a representative sample of malware. What I ask is why certain malware forms became known and how they can be reinterpreted ecologically. This is why I use ‘case’ rather than ‘case study’ as I do not seek to suggest that any of these malware forms should be seen as exemplars of malware more

generally. In this way I draw off Lauren Berlant’s discussion of the ‘case’ as being a “problem-event” (2007, p.663) as a “perturbation in the normative” (2007, p.670). By this, the case is an opening of a greater complexity between the generic and the particular (Cass, Schwanen and Shove, 2018, p.162) than what a case study typically denotes. Instead I offer these cases as “problem-events” not as prescriptive examples but as moments of reflection to re-assess the current theorisation of events, through an alternative view of cyberspace and thus of (cyber)security. I do this through exploring a range of anti-virus reports, speeches, newspaper articles and other media forms, and some technical analysis, to offer a new, parallel narrative. As Balzacq and Calvelty (2016, p.192) express on the study of Stuxnet, “the information that is most relevant for this case study... is to be found in the technical reports of anti-virus companies and security researchers, newspaper articles that focus on the timeline of preparation and discovery, and on technical blogs,” which is something I realised in my work. Below I provide a brief overview of each malware case in Table 1. The latter three provide an alternative aspect of how a more-than-human politics emerges through the ecological that I detail in chapter seven; whether that be in the hacking of democratic organisations, on ‘cyberweapons’, or how emergent events become interpreted. Whereas with Conficker, I glimpse towards a proto-ecological practice in cybersecurity.

Table 1: An overview of the four different malware cases.

| Case | Overview |
|------------------------------------|---|
| Conficker 2008 – 2009 | One of the fastest-growing <i>worms</i> in history that formed a botnet (a network of connected computing devices that can be used to spread malware or launch attacks). Conficker did relatively little malicious but motivated an <i>ad hoc</i> group of researchers and organisations to come together to form the Conficker Cabal, later known as the Conficker Working Group. This group, formed in response to a more-than-human threat, is important for recognising a proto-ecological response that formed new political alliances and networks, especially as Conficker’s choices grew. |
| The Dukes 2008 – Present | Unlike the other cases, the Dukes have been categorised by various MALs according to their relationship to a hacking organisation of the Russian state. They are classed as an Advanced Persistent Threat (APT) that use a variety of techniques to gain and maintain access to systems. Otherwise known as APT29 and CozyBear, the malware collective has been linked to the hacking of the US Democrat National Convention (DNC) and influencing elections in several countries. How MALs come to link different malware forms together, and how suspected author intentions inform this, requires further attention, especially when intent and malware performance are not aligned. |
| Stuxnet 2009 – 2010 | Referred to as the world’s first ‘cyberweapon’ against an Iranian uranium enrichment facility, this malware is still one of the most complex malware forms there has been. Unlike other studies on Stuxnet, I explore the ecologies in which it was revealed and argue that its more-than-human agencies led to its discovery. Thus, exploring how much effort goes into comprehending what the ecology may |

be, this case shows even with what is suspected to be a highly-specialist and controlled malware form - there is always excess and surprise.

**WannaCry /
(Not)Petya**

2017

WannaCry and (Not)Petya are two different forms that brought cybersecurity and malware to the fore of contemporary discussion in 2017. These two forms emerged after the ShadowBroker's dump of suspected US National Security Agency's (NSA) toolkits. Based in the Sophos MAL during these events, I sketch a variety of politics that emerged from each of these attacks, through their different geographies - varying from UK NHS trusts having most of their systems offline, to the AP Møller-Mærsk container shipping business being impacted to the tune of US\$200-\$300m to read how malware grafts onto conventional human politics.

I use these different cases in pursuit of an alternative orientation beyond the Sophos MAL to the 'wild,' to offer a horizontal perspective to a politics of cybersecurity. That is, cases exist to offer particular openings, and thus these cases offer particular renderings of a malware politics, and I explain the differences that may occur through what I outline as different gradients of political agency and affectual capacities in chapter seven. And, through addressing malware ecologically, it was through my experience as becoming-analyst that I was able to pursue this investigation that these cases gesture towards - a re-writing of the stories that we are told about malware.

Recording, Reflecting, and Ethics

Yet, in becoming-analyst, recording research was a difficult task in order to maintain forms of privacy and confidentiality, and also trying to provide as much as a *truthful* record of my (auto)ethnographic research and interviews. My position as becoming-analyst generated certain perspectives. As an 'intern' it meant I was not seen as management, allowing for alternative power relations, and thus relationships. As many analysts saw themselves as reverse engineers, the apparent creep to engage in more media-focused work was seen as time away from what they enjoyed: analysing malware (though this was not exclusive). This was part of a broader sense of management not understanding their work.

There was a discussion about management and their lack of understanding of what they were doing. Here they said they didn't know the intricacies of their work - that they had to detect (or even play with) malware in certain ways and that introducing new systems and completely cutting back on what they did was 'idiotic.' There was a sense that there were too many managers aiming for the same goal but with different agendas."

(Research Diary, 21 April 2017)

This was always an underlying tension in the MAL for some – between the various roles that the analyst now has to play. Some clearly had made clear they did not want to engage in work that was not directly related to the analysis and detection of malware whilst others actively said they preferred this.

Working in a small laboratory however, where reidentification was easily possible by those involved, introduces certain risks and issues of confidentiality that were developed in the research ethics and supporting information¹⁸. All analysts in the MAL knew I was a researcher looking at malware and those who analysed them and wrote detections. So, though the perspectives on management may be interesting, I do not develop these, nor on other issues, as I did not intend to research these practices – and I uphold this. So, in ensuring the confidentiality of participants, not only at the Sophos MAL but in interviews, all names have been replaced to provide pseudonymity. I was always forthright to those who I worked with about my purpose, and rightly said I was mainly focused on the malware rather than them.

I used the Evernote application both on my personal computer, and primarily on my iPhone, to write my research diary. Within the MAL, I did not take any materials out, and only made slight notes in the day on a notepad, meaning that all of my written notes that form my research diary are those which are recorded post-event. On most days (I did not achieve this every day), I wrote my research diary on the 45-minute train home directly using my iPhone to type these notes, which automatically synced through cloud-based EU-Switzerland services with my computer, meaning I always maintained a back-up. The use of Evernote has been a constant throughout my research, and also allowed me to add photos, links to newspaper articles, and draw together different research materials. However, writing these notes on a train, after a sometimes-tiring day in the MAL meant I did not develop prose in similar ways to much in-depth (auto)ethnographic work – with my notes being sketchy, things remembered, or briefly scribbled down on paper in the MAL to read before I left at the end of the day. Perhaps this is something I could have done better – however, there are unfortunate limits to my concentration and ability to write eloquently.

¹⁸ The research was approved by the University of Oxford Central University Research Ethics Committee (reference number: SOGE 17 1A-4).



Figure 3: Research diary with coded strips highlighted. Author's own image.

To draw together research materials across the thesis, I used the NVivo (11.4.3) software for the analysis and coding of materials. Using an iterative process of coding, I drew broad categorisations that I then worked with to draw out important aspects of malware analysis and events. I printed each 'node' (Figure 3) to work with the materials in paper form to help develop more nuanced readings of each node. This helped to blend the first task of reducing the roughly 75,000 words of the research diary into palatable sections to tease out themes and directions of my research. This was not an easy task – often revisiting difficult times I had – and one which often required time and thought to re-engage with the material in meaningful ways. The combination of photos, newspaper materials, and little sketched notes all helped to record my experience. Yet, just reading through my research diary, or even writing, brought up *new* things that were not recorded in great depth, meaning I constantly scribbled thoughts down for use, or as a better illustrative example. This is part of the post-phenomenological process of reflecting, stepping-back, and re-reading and re-imagining moments of malware agency, how I felt and how this transformed my experience in the MAL, for instance. Recording is not some single act that I did on the train after a day of

fieldwork; but it is part of my bodily memories. When writing, something forgotten came forth, and then I added further to my notes, or something that was said came back to me.

In this methodological overview, I have then provided an exploration of how my approach to the study of computation and malware draws on a range of different themes that can be blended into the *becoming-analyst*; in a post-phenomenological sensitivity through (auto)ethnography that re-situates malware at the centre of my concern. This broad appreciation can be undertaken in a more-than-human methodological approach that generates an empirical sensitivity to some of the theoretical methodological constructs I have outlined. Pursuing research in a malware analysis laboratory enabled me to develop a relationship with malware that only an (auto)ethnographic study can provide. This in-depth personal engagement however has not let me ignore the specific (human) social spaces and their politics. My engagement is critical, then, to comprehend why certain politics and agencies of malware are so impactful on the world. As Mol (2000, pp.96–97) argues in researching a laboratory:

“The very advantages and disadvantages, the goods and bads, of performing reality in one way or another are themselves up for debate. A debate that is political – and simultaneously concerns the very materialities of microscopes, examination tables, paper forms, cold feet and conversations. The politics of the medical field is a matter of organization, but also of metal, and of suffering, and of flesh.”

I think something similar exists in the practice of studying computational materialities – the play of working with malware. It was the materialities of the MAL, and that of (malicious) software, that made politics accessible to the becoming-analyst and permitted an eco-logical perspective to develop to inform a repositioning of malware in cybersecurity.

On my last day in the MAL I wrote “it was odd finishing, no grand ceremony. I said my last goodbyes around the lab but then that was it. I simply walked out and that was the end of the ethnography. Odd” (Research Diary, 31 July 2017). The world moved on. Malware was still being analysed and detected as ever. However, I am not the same as when I entered, and this has allowed for a deeper engagement with the materials of research and on the cases that inform this thesis. When I returned to Sophos to meet Daniel (Research Diary, 15 November 2018) to discuss some of the my ‘findings’ (primarily in chapters five and six), it was like meeting an old colleague, and one where the discussion was frank and open, where we agreed on the core components of my work. As Greenhough (2010, p.48) says, “fieldwork is more

than a process of data collection; it is an event through which the researcher and researched are resituated or repositioned in the world.” This is undoubted in my case, in ways that I cannot fathom, by cognisers that work in different ways to me. An eco-logical sensitivity is the ‘output’ from this co-constitutive experience. Throughout the rest of the thesis, I draw on this experience to re-think malware, and its role in politics and cybersecurity, that itself is informed by various lineages that I now turn to.

Chapter Four || Lineages to Understanding Malware

“A virus may be understood as a calculational process at the material level of computer circuits but when this accident (event) is called “malicious software” it connects to a whole incorporeal sphere of morals, crimes, criminals, laws and judgements. Hence, an analysis of computer culture should not focus on solely on the material event of calculation (the technical diagrams) nor on the discursive events, but on the constant double articulation between various semiotic regimes.”

(Parikka, 2009, p.107)

Malware are imbricated among varying semiotic regimes that include computational materialities, social interactions, and information that become part of the computational ecologies that I outlined in chapter two. In this chapter, I instead look to Parikka’s claim through the angle of understanding malware as agential forms rather than solely through computational cultures – so that cognition, through its calculative modes, can also be addressed. A politics of cybersecurity demands an attention not only to the ‘material events of calculation’ nor to its ‘discursive events,’ but the interaction between the two. Thus, I extend and make palpable how lineages have been constructed; both how events of calculation have been understood, and how these are influenced by essential *patho*-logical traits, building upon organicism and cybernetic discourses.

Pandemic, public health, emergence, viruses, and antibodies converge in a *Wired* interview with Barack Obama (2016) near the end of his Presidency of the United States. There are intriguing slippages, analogies to bacteria, to being ‘clean’ and talk of movement away from ‘armor’ and ‘walls’ to one of “*think[ing] differently about our security.*” A broad *patho*-logic – that is thinking medically, biologically, and even militarily – of cybersecurity is not a recent phenomenon but has its roots in the emergence of mathematical thought, computing, and how viruses and worms were first conceptualised. Thus, I question how we understand about malicious software today, and simultaneously critically interrogate the pathologies that inform computation and cybersecurity. Cybersecurity and malware do not sit in a vacuum but are intimately tied in the broader logics of governance, biopolitics, and geopolitics that are intermingled with senses of risk, feeling, threat, and vulnerability.

Malware sit at the heart of forms of collective vulnerabilities where “the stealthy and technical nature of “rouge computer programs” and the difficulty of obtaining relevant data, [mean] we know little about the scope and impact of the problem [of malware]” (Hughes

and DeLone, 2007, p.79). This vulnerability that malware expose form an omnipresence that could undermine the entire system – that are in part connected to post-9/11 cyberspace and cybersecurity as they became integral to US imaginations of security (Kaiser, 2015, p.13). This situates malware within the securitisation of ‘cyberspace’ (that includes network infrastructure, data, and the internet). These coalesce with such initiatives as the Chinese ‘social credit score’ (Mistreanu, 2018), to Russia’s attempts to create a ‘national’ internet (Roth, 2019), through to the ability to track anomalies to determine software as malicious. This connection between national security and cybersecurity was made explicit in the UK Government’s “major offensive cyber campaign” against *Daesh* (or the Islamic State of Iraq and the Levant (ISIS)); its first confirmed offensive cyber operation (Flemming, 2018, p.5) demonstrating how cybersecurity has come to the fore of national defence (Betz and Stevens, 2011).

In this chapter, I provide an overview of the discourses of cybersecurity and malicious software, and its subsequent securitisation through endpoint (‘anti-virus’) detection engines. As Parikka and Sampson say, “long after Nietzsche, the question of whether things are good or evil, positive or negative, normal or strange remains on the tip of the collective tongue” (2009, p.x). Malware, through pathological practice are still seen as *a priori* ‘bad,’ where software are neutral and are infused with the intent of their author, are treated medically, and draw on various lineages of the military and biology. This is imbricated with certain senses of how security is performed – whether through the protection of children, or the military’s priority of network protection. Considering these different, intersecting lineages, I first ask how cybersecurity is imbricated with other forms of post-9/11 computational securities and how the changing spatialities of computation has transformed how cybersecurity itself is practiced. I then explore how ‘benevolent’ software architectures were constructed as malicious and how these were undermined by changing securities that intermingled with pathology. In an explicit turn to pathology through its various guises in biopolitics, in the formation of (ab)normality and anomalies, and the definitions of malware themselves, I explore how these resonate with how we respond to the threat of malware today. Anti-virus – and today’s endpoint detection – have been central to this and are infused with both how cybersecurity has been (re)constructed and how pathology has become the central theme to its actualisation. The aim of this chapter, then, is to draw together some of the lineages which have permeated our understanding of malware as broadly pathological,

comparable to biological equivalents, and have become enrolled in broader security architectures. And how this shapes a response which renders malware as technical, without a politics, and as something that can be abstracted, rendered knowable, and detected by endpoint detection engines on our computers.

The Growth of Cybersecurity

‘Cyberspace’ – in its conventional reading – has been treated as ungovernable, unknowable, as something that makes us vulnerable, is inevitably threatening, and is full of threatening actors (Barnard-Wills and Ashenden, 2012, pp.116–119). The securitisation of computation, through cybersecurity, has been constructed through multiple lineages, and cannot simply be understood as a ‘new’ phenomenon, but sits within and is (re)formulated from other forms of (in)security. The growth of big data and the production of information through computation is formative of a security calculus for our times – which has orientated security to the potential for action and a capacity to act in the face of uncertainty (Amoore, 2011). This has led to anticipatory action (Aradau and Blanke, 2016; Adey and Anderson, 2012); where, if only there were more data, the more *secure* the future would be through the propensity to narrow, and condition futures. Calculative and pre-emptive actions are full of discursive and affective production of sensibilities and sensitivities to risk (Beck, 1992), such as in tracking suspicious transactions and movements that feed into pre-emptive drone strikes, that predict a potential future (violent) act to condone a present, anticipatory response (Amoore and de Goede, 2008; Gregory, 2011; Crampton, 2016; Shaw, 2016b).

Malware have thus been imbricated in the broader ecologies of threats, worries, and concerns that seek to secure in an increasingly mobile and uncertain world. Within cybersecurity this has transformed beyond the conventional confines and interests of militaries and large corporations from the late 1970s and early 1980s to everyday sites of the home, the online shopping transaction, and centralised infrastructures such as water and electricity that undergrid life, at least in the Global North. As Massumi (2010, p.61) writes, “the terrorist threat series blends into series featuring generic identities. There is a generic viral series, real and non-existent, as heterogeneous as human-adapted avian flu, SARS, West Nile virus, and

the Millennium bug¹⁹.” Though the Millennium bug was not malware, it is easy to see how malware – through the popular computational virus – became enrolled in such uncertainties. Computation and malware then have become connected to the pandemic – as ex-President Obama made explicit at I highlighted at the start of this chapter – that support assertions to protect ‘catastrophic’ supply-chain failures of ‘critical national infrastructures’ (Folkers, 2017) to protecting those at ‘home.’ The latter has explicitly materialised through privatised security through endpoint (anti-virus) products, that seek to secure the ‘internet economy’ (Carr and Tanczer, 2018; Timmers, 2018).

Malware have been enrolled into broader methods to attend to threats as information to ‘connect the dots’ post-9/11 became a central focus for conventional securities (Amoore, 2011). This is seen through a range of security architectures; in financial transactions (Wesseling, de Goede and Amoore, 2012), airline records (Pötzsch, 2015), on connections in chatrooms (Weimann, 2016; Chen, 2011), to tracking congestion (Anderson and Gordon, 2016; Gordon, 2012), and even on the next influenza epidemic (Roberts and Elbe, 2016). Practices of computational securities are not new but fall into distinct histories of collating data into tables, forms, and datasets as Foucault highlights in *Discipline and Punish* (1991). These analogue records can be equally, if not more, violent than the things they seek to secure against²⁰, and while computationally-mediated securities (that include cybersecurity) are not necessarily anymore violent or intrusive than previous epochs, their reach and extent into everyday patterns; and its transferability into contexts and ecologies unknown have grown (Amoore and Raley, 2017). Today, the ‘enemy’ is often not clearly defined, is distributed, and must be reconstituted through the tracing of suspicious actions under surveillant practices that allow for a recombination of attributes – such as travel, financial transactions, and phone records (Amoore, 2011; de Goede, 2018). Rather than *being* a threat, one emerges as a threat. Amoore’s work on data derivatives (2011) demonstrates how different components of the life of an individual are constructed as suspicious attributes that allow for security decisions to be made. In connecting dots and tracing movements and

¹⁹ The Millennium bug, as it quickly became known, was a potential issue in how computers dealt with dates, with the upcoming change of the millennial year from 1999 to 2000. This led to refrains to the collapse of computing infrastructure. In an interview I did for another project, on cyber-medical systems, one interviewee noted the issue that he named ‘Y2K’, “We got packs handed out with green stickers in that said Y2K compatible and stuck those on everything, doors and everything, Y2K non-compatible [laughter] it was just seen as a bit of a farce” (Interview, 20 August 2015).

²⁰ See the Australian ‘certificates of exemption’ that granted citizenship rights to Aboriginal and Torres Strait Islander peoples only by denying their heritage (McGrath, 1993).

behaviours, and crucially, in drawing these together, big data allows security infrastructures to formulate anomalous attributes and identify threats. This abstraction and reformulation of information can make, however, the construction of threats appear apolitical, mechanistic, and a rationalisation of a ‘truth,’ which obscures the extensive labour that go into their formation.

Cyberspace and Hackers

As I described in chapter two, there is a perception that the world can be objectively studied, that has had implications for how cybersecurity has come to be understood. This objectivity can be seen to underpin assertions in contemporary security dynamics that data can ‘reveal’ a certainty (Amoore and Hall, 2009). The prefix *cyber* itself derives from the Greek, κυβερνήτης (Oxford English Dictionary, 2016) with a translation as *steersman*, that underpins a gendered, masculinist dominance and control. This flows into and from *cybernetic* discourses that determines, at least in its first order, a commonality secured on understanding communication through abstracted information. That is, humans alone can control, map, and render knowable the world. The use of ‘cyber’ has a lineage that extends beyond security communities; which has important implications for two key words I engage with; *cybersecurity*, and its associated spatialities, *cyberspace*. It is not possible to separate some of the discourses that emerge around these words, and the implications for how cybersecurity, and thus malware, have come to be understood.

In the literary genre, cyberpunk, there is a fear of all-controlling technology – challenging a cybernetic discourse based on control, feedback loops, and order. Fears are dramatised in these literatures, represented in early films such as the Terminator with vistas of post-apocalyptic futures (Hayles, 1993, p.82). The novelist William Gibson popularises ‘cyberspace’ in his novel *Neuromancer* (1984), where it has been commonly interpreted as a non-stratified variable (Kneale, 1999; Bingham, 1999) reinforced by an accompanying computer game (Miles, 1988) loosely based on the novel that built on visions of ‘entering’ cyberspace. Yet, I argue there should be an alternative reading drawing on minor narratives already-present in these novels to query cyberspace’s origin stories and provide an alternative understanding of its histories. Cyberspace has always exhibited a stratified, complex, and intermingled spatiality. Developing on computational cognition, the different environments

in which malware interprets signs through cybersemiosis, mean that it is possible to see cybersecurity not only operating on a flat, unstratified plane but full of fissure, friction and *cæsura*. This alternative reading can be seen in an earlier 1982 novel by Gibson, *Burning Chrome*, where “glitch systems are [full of] cybernetic virus analogs, self-replicating and voracious. They mutate constantly, in unison, subverting and absorbing Chrome’s defences²¹” (1988, p.202) where viruses – the ‘cybernetic virus analogs’ – allow for something else to be present. That is, malware in various forms have always been a disruptor of cyberspace since its conception, so though cyberspace may have been thought of in an ethereal sense, there are alternative ways to perceive of it.

The apparently ‘jacked-in’ world of these literatures, of robots and androids, along with cybernetics, share a broad conceptualisation of information as the defining feature of existence and (in)security. The individuals who are meant to be able to wield that information, data, and code – hackers – are then positioned as central to malware activity. They are frequently tagged as the bored, intelligent, or simply those who seek ‘freedom’. As Sherry Turkle (1984) writes in her ethnography of hackers in the late 1970s and early 1980s, they were seen as broadly harmless, but that the growth of personal computing led to senses of threats entering the home. From this period, as different malware forms became more proliferate, and as societies became variously more dependent on computation for their functioning, hackers posed a growing risk to their security. Two decades after Turkle’s ethnography, Galloway (2004, p.157) told us that “the current debate on hackers is helplessly throttled by the discourse of contemporary liberalism: should one respect data as private property, or should one cultivate individual freedom and leave computer users well enough alone?” In this sense, hackers became both the liberators and detractors of cyberspace, crafting code with their special, unique skills. Computer hackers are deemed to be, at least in earlier periods, concerned with cracking systems – and Hruska (1990) distinguishes these from *fraks*, who wished to cause malicious harm – and refers to school computer clubs as being places that create both viruses and hackers. But, unlike this early fluidity, where

²¹ This comes from Gibson’s first book of the series, *Burning Chrome* (first published 1982), that was the first time that he used the word ‘cyberspace.’ Chrome here refers not to the popular browser ‘Google Chrome’ but to the all-powerful Chrome – where the “virus analogs” attempt to dismantle its geometric power.

hackers were seen mainly as liberators, experimentalists, or even artists, they were slowly twisted as a threat against, state, capital and property.

As computation and hackers moved out from research laboratories and ‘tested’ the systems of large corporations, militaries, and increasingly the home, problems emerged. This had implications for how cyberspace, and malware, have been seen since – where cyberspace’s apparent expanses became vulnerable to the tools of human control, malware. Hackers became increasingly associated with crime and deviance throughout the 1990s and 2000s. The movement from individual hackers, who often displayed their names as a product of their work such as the (in)famous hacker, Spanksa (Szőr, 2005, p.63), had morphed into anonymous productions that differently aimed for stealth, revenue-generation, spying, and extraction. Many contemporary hackers are now part of sophisticated large-scale, industrious organisations that participate in cybercrime (Jordan and Taylor, 2008; Holt et al., 2012; Lusthaus, 2013; Lusthaus and Varese, 2017; Skibell, 2002) or work for states – such as, but not limited to, China, the Netherlands, North Korea, Russia, the UK, and the USA. The way malware are produced has significantly changed too. On one side are specialised, state-level malware authors and sophisticated gangs and on the other hand ‘kits’ that have been produced to sell malware variants according to purchaser preference (Ollmann, 2008). The emergence, and deployment of state-crafted exploit kits however, that have been incorporated into criminal kits and malware forms, have led to a confluence and confusion of groups of individuals, where distinctions between different actors have become blurred and has challenged attribution, for instance (Maurer, 2018).

The Changing Spatiality of Cybersecurity

Whereas cyberspace may have given us certain resonance around control, flatless terrains and how hackers have been seen as central, there has also been a changing spatiality that has made cybersecurity become central to contemporary concerns of computation. Within international relations, and some work in geography, the Estonian ‘Bronze Night’ has been invoked as a turning point in cybersecurity (Kaiser, 2012, 2015; Kello, 2017), and even as the world’s first ‘cyberwar.’ I think this can be attributed to new forms of spatiality that transformed cybersecurity itself. In 2007, the relocation of a Soviet World War II memorial, the Bronze Soldier, from the Tõnismägi, in the centre of the Estonia capital, Tallinn, to the Defence Forces cemetery outside of its centre provoked a “riot” (Kaiser, 2012, p.11) that

garnered significant support from the large Russian-speaking minority. This was shortly followed by a Distributed Denial of Service (DDoS) attack (along with other spamming and the defacement of the Estonian Reform Party's website) targeted against government services, banks, and other institutions, leading to outages across the country. DDoS attacks involve sending requests to servers to overwhelm them to block 'legitimate' requests or to stop a website responding, for instance. This Estonian DDoS is cited as a seminal moment that sedimented the importance of cybersecurity within international relations (Kaiser, 2015; Kello, 2017; Robinson and Martin, 2017), centring Estonia as a principal site of expertise²².

But we may ask why did Estonia become so seminal? The attack exposed how reliant states and their citizens are upon computational infrastructures that introduce new vulnerabilities and risks – where interconnectedness and essential services dependent on computation now exposed to these new forms of vulnerability. The Estonian 'Bronze Night,' however, should not be deemed as some form of rupture. There were prior attacks against state institutions, just not at the intensity experienced in Estonia, which quickly coalesced into a discourse as the world's 'first' cyberwar as exemplified in a New York Times piece a month after the start of the attacks (Landler and Markoff, 2007). That this DDoS attack was (human-directed) against a state and its core private institutions meant it rightly garnered attention. But it is only one actualisation of a malware politics. It should not one be considered something new, but a continuation of expanded computational capacity, where more devices and critical services moving online has led to greater computational dependency. Hence the changing spatialities of life itself through computation have transformed cybersecurity, and where malware can be elevated to that of international relations.

As the role of inter-net(works) expanded from military sites, to corporations, to PCs, to running critical national infrastructure, to traffic lights, and home heating systems as part of the 'Internet of Things' (IoT), cybersecurity has had to respond to new spatialities, where everyday objects have become sites of interest. In early computer security, concepts such as the *Confidentiality, Integrity and Availability* (CIA) on isolated computational systems, as documented in the 1983 'Orange Book' by the US Department Defense, have been extremely

²² The NATO Cooperative Cyber Defence Centre of Excellence (CCDCOE) was created in 2008 at the instigation of Estonia and six other states (Germany, Italy, Latvia, Lithuania, Slovak Republic, and Spain). Today it is the centre of expertise in cybersecurity with the NATO alliance, and works on developing 'cyber' norms in international warfare and developing strategy guidance and has situated Estonia as a global leader in the area. For more information see CCDCOE (2019).

influential on contemporary practice (Smith and Marchesini, 2007). And the US Defense Advanced Research Projects Agency (DARPA) was instrumental to the funding and formation of the predecessor to today's internet, ARPANET, that was designed for resilience in a networked world (Leiner et al., 2009). Yet, conventional techniques of computer security, that used to focus on individual PCs and servers, developed borders for protection that do not work today. That is, rather than seeing computation as isolated units in which borders could be established – such as through firewalls – there is now a requirement to monitor *within* computing systems as computers have become more integrated and the sharing of data is seen as essential in the modern economy.

Everyday lives that were not previously considered part of cybersecurity are now dependent upon it – and thus malware are not only present on PCs, but on our phones in our pockets, and the webcams used to monitor and observe our children. Not only do objects we engage with everyday become security threats, they can also combine together in certain formations to produce *emergent* effects that an individualised security response alone cannot address. One example includes supply-chains that introduce vulnerabilities across different components that make final configurations insecure. For example, the 2016 Mirai botnet targeted weak authentication controls using default factory-installed usernames and passwords, leading to one of the largest internet attacks to that date (Krebs, 2016). The Mirai DDoS attack was based on accessing routers, printers, and security cameras that had factory-installed defaults that tied them into a *net* of controlled *bots* (hence the name *botnet*). In this case, this targeted Dyn, a domain name system provider that resolves human-readable domain names to IP addresses. The effect of the attack was that end users could no longer access a website simply by typing its human-readable name through Dyn unless one had the IP address. Supply-chain concerns have recently surfaced in relation to the Chinese telecommunications provider, *Huawei*, who has been accused of being able to undermine the security of '5G' infrastructures through the integration of their devices. A 2019 UK report at a specialist centre analysing the code of Huawei products has said that it can only provide 'limited assurance' of security (Huawei Cyber Security Evaluation Centre (HCSEC) Oversight Board, 2019). Thus, the posture of cybersecurity has had to change, according to the new spatialities of computation – that are increasingly becoming part of an embedded fabric of life through the IoT – meaning that a focus on individual devices is no longer deemed satisfactory. It is about understanding risks and vulnerabilities at least contextually,

not unlike other security regimes that use big data to ‘connect the dots.’ Therefore, the changing spatiality of computation, the growth in devices, and their interconnectivity has in turn transformed how cybersecurity now uses greater realms of contextualisation but within a certain patho-logic.

From the Tallinn ‘Bronze Night’ to the emergence of the *dangerous* hacker, through to how the military has attempted to secure their networks, has led to multifaceted and complex arrangements we see today. Whether that be in *who* and *what* is secured – to who ‘acts’ and has ‘intention.’ The creation of the UK National Cyber Security Centre has been such a response to the importance attached by states to cybersecurity beyond the conventional confines of protection of critical national infrastructure or the military. In becoming a ‘public-facing’ centre for the UK Government, it now distributes advice and assists in developing broader cybersecurity responses for the UK – demonstrating how the state is addressing a changing role of the state and private enterprise, that previously left cybersecurity to individuals and organisations, and how much ‘critical national infrastructure’ is now either in private hands or as part of public-private partnerships (Carr, 2016; Cavelti, 2015). The growth of cybersecurity beyond the military, has had profound changes in the spatiality of security. This means there are a mixture of actors, that include endpoint protection providers, that are responding to the growth of cybersecurity and the changing spatiality of computation – transforming attention from isolated devices and their boundary protection through firewalls, for example. Yet these are in turn influenced by how malware have been understood, where their architectures have been constructed as malicious alongside the fears and vulnerabilities of computation.

Architectures of Maliciousness

Perhaps surprisingly, there was no formal definition of the computational virus until Fred Cohen’s work in 1983. Though the word *virus* is attributed to his supervisor, Len Adleman (Zandi, 2014), by 1987 there was already a decade of referring to computer worms and viruses in technical papers (Highland, 1987). Cohen claimed “a “virus” may be loosely defined as a sequence of symbols which, upon *interpretation*, causes other sequences of symbols to contain (possibly evolved) virus(es)” (1989, p.325, emphasis added). In the same paper, he considers how a virus may ‘evolve’ as a program, and offers proof that viruses are at least as powerful as

other forms of computation. Though I do not centre this thesis on viruses, their history provides an important reminder that malicious software have a historical emergence that is conditioned by the viral and worming architectures of the mid-to-late-twentieth century, that intermingle with concerns of the role of computation in society (Turkle, 1984; Parikka, 2016; Parikka and Sampson, 2009). Furthermore, they are shaped by that history – where cybernetics, control, and feedback feed into senses of malware as being the perpetrators of potential systemic breakdown.

Malware vary according to the specific spatio-temporal dimensions in which they work – they must work in the ecologies they perform in. As Cohen outlined in pseudo-code – an abstract representation of program code – in 1989:

```

program virus:=
{1234567;
subroutine infect-executable:=
  {loop:file=get=random-executable-file;
  if first-line-of-file=1234567
  then goto loop;
  prepend virus to file;
  }
subroutine do-damage:=
  {whatever damage is to be done}
subroutine trigger-pulled:=
  {return true if some condition holds}
main=program:=
  {infect executable;
  if trigger-pulled then do-damage;
  goto next;
  }
next}
(I989, p.326)

```

This pseudocode presents the place to locate a virus, in this case at the end of a file. This is then added, with subsequent damage, to the ‘trigger’ that then goes on to infect another executable file. Though this is only one high-level way to think about how a virus operates (over time, viruses had many different entry points to infect a program), it sets out a logical architecture. This allows the virus to be “a program that can “infect” other programs by modifying them to include a, possibly evolved, copy of itself” (1989, p.325). From the viral ‘infection,’ *and all subsequent malware*, a *patho*-logic has been embedded within it. As Ferbrache noted in 1992, “a feature of the anti-virus community has been the adoption of a wide range of (often conflicting) terminology, based mainly on the analogy between biological and computer viruses” (1992, p.3). This early connection, though not necessarily equivalent, through a borrowing and deployment of analogy, has held a pervasive and strong grip on contemporary practice.

How software come to be categorised, formed, and catalogued as malicious is not a neutral activity. It is infused with lineages that structure how we even think such *things* as virus, worm, and Trojan horse. Across a multiplicity of terms, they are constructed through a desire to describe and allocate the ‘infectious other’ assisted by the equivalence of the natural and the mechanical in organicism.

The ‘virus,’ the most popularly-known malware architecture, has an evocative, seductive biological signifier. This has enabled the ‘virus’ to be relatable, especially when compared to the broader catch-all term, *malicious software*, that is commonly used today, and that I investigate. The virus, and its dominant representations in popular culture (Parikka and Sampson, 2009; Parikka, 2016; Helmreich, 2000) have then become the encompassing point of entry that assume maliciousness straight away. But this ‘evil’ of the virus has not always been present. In early computational viral research, viruses were not yet codified as being solely malicious. This allowed for the calls for the *benevolent virus* (even if this was pathologised) by Cohen (1991). That is, because it is an architecture (rather than a particular actualisation), intent can be ‘good’ or ‘evil.’ This is in contrast to the contemporary use of malware which can only be nefarious, we are told. Intentionality is folded immediately upon malware, closing down the possibility for *action* and *performance* in ways that move beyond intent.

Early research into viral and worm architectures were therefore not invariably imbued with senses of maliciousness. Creeper, created in late 1970 by Bob Thomas (Dalakov, 2018; Parikka, 2016, p.19), moved independently across systems and inspired Xerox Palo Alto Labs (presented in the paper by Shoch and Hupp, 1982) to develop worms in an effort to improve distributed computing to handle tasks more efficiently. As Cobb and Lee note (2014, p.74), “the allure of using self-replicating computer code to perform beneficial tasks dates back at least as far as the 1980s.” These benevolent viruses (Aycock, 2006) show these were not architectures of maliciousness, but certain computational practices that mutated later. Cohen (1991) was particularly instrumental in stating how viral architectures could be used alternatively. Yet whilst viruses and other ‘malware’ architectures were considered potentially benevolent in the early 1990s, this soon fell away as they were increasingly associated with malicious individuals – hackers (Parikka, 2016). This suggests that there are

no forms of computation, *viral or not*, that are inherently malicious²³. Even the virus, that was so commonly associated with nefarious acts and hackers, was not ‘bad,’ but actively became constructed as such.

Though benevolent viruses may have disappeared, their architectures have been developed in alternative areas such as in ‘digital life.’ The emergence of the (now defunct) Digital Life Lab (2009) at the California Institute of Technology, and the Digital Evolution Laboratory (2018) at Michigan State University in the early 2000s, demonstrate how certain behaviours now associated with malicious activity; such as virality, self-propagation, and mutation, became one means to test *artificial life* through the comparison between biological and computational viruses (Helmreich, 2004). ‘Artificial’ life has an earlier provenance in the 1984 game *Core War* inspired by the ‘Creeper’ worm and the 1961 *Darwin* game created by Victor Vyssotsky at Bell Labs. Vyssotsky described the game (1972, p.93) as;

“a game between computer programs as computer programs. The objective is survival; programs may ‘kill’ one another, and may create copies of themselves.”

Core War, inspired by early gaming, shows how the materiality of today’s malicious architectures (Dewdney, 1985, 1984) were important contributions to early conceptions of computer gaming that were connected to war, masculinities, and control (Kline, Dyer-Witford and De Peuter, 2003; Ross, 1990). As Johnston (2009a, p.24) writes, in a weaving of computer games and malware – “we find computer worms or viruses, computer games, and *Alife* – Artificial Life – three things that would later be separated out and sealed off from one another, but here jostling together.” In the development of artificial life, in the analogies between biological and computational viruses (and other malware) in an organistic manner, viruses and worms have a history that is warped in *how* one uses software forms (good or evil – or even for play!), encouraged by slippages between biology, medicine, and computation. This, in turn, is conditioned by the intermingling with narratives of war and supremacy that are seen in *Core War* – which is for killing, winning, and defeating the enemy. However, these early architectures quickly became separated from benevolent uses and rendered dangerous, as ‘control’ was lost, and they became associated with the exclusive control of the

²³ Some may use a common example of ransomware to argue that this particular architecture is already imbued with maliciousness. However, the function of opening and closing files whilst encrypting these may only happen in what we deem malicious – the process of remote encryption most definitely does not.

hacker. As these architectures became constructed as malicious as computational viruses and worms, there was a requirement to tackle these vulnerabilities, that leaned heavily on pathology as a way to deal with this nefarious other.

Pathological Imaginations

As Donna Haraway argued in *Simians, Cyborgs, and Women* (1991), “the immune system is an elaborate icon for principle systems of symbolic and material ‘difference’ in late capitalism” (1991, p.204). The broad sweep of biological metaphors throughout ‘late capitalism’ have been particularly strong in understanding security and, in particular, how it has been spatialised. As Parikka (2016, p.119) says, “the clean body of modernization found its imaginary ideal in the computer organism. Just as the body biologic (and politic) was, and from the end of the nineteenth century, the object of constant attacks by minuscule viruses and bacteria, so the computer soon has its own share of dirt.” Cleanliness, part of medicalisation and sterilisation, inoculation and immunity, become the mirror opposite of malicious software. In this section, I explore how pathology – that is conditioned by the history of organicism and cybernetics – has informed collective understandings of the architectures I explained and how this is generative of particular renditions, sensitivities and understandings of maliciousness and malware.

Pathological imaginations of computer viruses are apparent in early discussions, such as in volumes of the popular *Byte Magazine* (published 1975 – 2013), and in texts such as Fred Cohen’s *It’s Alive! The New Breed of Living Computer Programs* (1994). Some, such as *A Pathology of Computer Viruses*, acknowledge an analogical tension between biological and computational viruses (Ferbrache, 1992, p.3) but still continue to use these. By speaking in a biological register, computer science and malware analysis were able to lean on the authority of natural science, borrowing language from immunology where computers were “self-contained bodies that must be protected from an outside threat” (Helmreich, 2000, p.473). This equivalence between biology and computing, drawing on histories of organicism and cybernetics, intersected with a medical pathological lens, so that malware was treated in the same way as a biological virus either through examining inside them or by monitoring their activity and interaction with their environments. A tight connection to biology meant that computation’s associated security adopted a biopolitical register, infused by the concept of

the immunological (Bassaganya-Riera, 2015), where biology, according to Forbes (2004, p.103) “provides a remarkably good model for creating a computing immune system.” Therefore, in comparison to equivalence and analogisation of computation that I was concerned with in chapter two, here I am interested in its actualisation in malware analysis and detection and how it forged with biology and medicine through the 1980s to today.

Biopolitics

Biopolitics, as a pathologisation of politics, has a long, complex history that does not only apply to cybersecurity. As outlined by Michel Foucault across several works; notably *Discipline and Punish* (1991), and several of his lectures at the *Collège de France*; the 1975 *Abnormal* (2016) and 1976 *Society Must be Defended* lecture series (2003a); technologies have been central to security techniques. As Elden (2016, p.40) notes on *Society Must be Defended*, “politics becomes increasingly mathematical as it becomes medical.” This mathematical, medical, and biological connection is not necessarily new as I explored with organicism in Descartes and Leibniz. However, Foucault describes how security deviates from prior instantiations of power and political control (in a transition from sovereign power to disciplinary power), in its increased emphasis on surveillance and data collection which enabled the comparison (by states and by individuals themselves) of each member of the population to an idealised norm. These forms of power are then exercised through the regulation of bodily activity, self-regulation, and controlling and observing the movement of bodies to assess their deviation from those norms. Most critically, Foucault (2008) terms biopolitics as a *dispositif de sécurité*, that is a rendering of administrative, institutional, and knowledge structures to control the body. I think this is applicable not only to humans but computers as well. This has important implications when contemporary malware analysis and detection attempts to draw out maliciousness through (ab)normality and anomaly.

There is a link between what Foucault termed as biopolitics, the pathological, and the *viral* that Parikka (2016) sees as a structuring of the networks of contemporary society. Foucault argues that there has been a historic movement between sovereign power, of the power or right to kill, to disciplinary power, as the explicit control of the body through institutionalisation, to the biopolitical in that it is a politics to “make ‘live’ and ‘let’ die” (2003a, p.241), that is more interested in self-control, circulation, and monitoring. The latter’s focus on circulation, or on ‘viral’ network cultures, is prevalent in understanding

malware. But as Povinelli (2016) questions, the distinction in Foucault's work between sovereign, disciplinary, and biopolitical societies, should be challenged, as these work in tandem in different times and places. Likewise, Agamben (2005) sees sovereign power, or at least the limit of juridical order and biopolitics, as mutually reinforcing in the production of 'bare life.' Hence, when observing that biopolitics is also *pathological* in its monitoring and circulations to find the deviant, it is also reliant on both sovereign and disciplinary power. Dillon and Lobo-Guerrero (2008, p.276) also note that there is "no geopolitics that does not imply a correlate biopolitics, and no biopolitics without its corresponding geopolitics." Through this, there is an ability to draw connections between different strands of power, where the performances of biopolitics, securities and sovereignties are ever-present together, albeit in different configurations, that permit a politics that transcends typical hierarchical separation between the pathological treatment of malware forms and geopolitics, for instance. As I shall argue in chapters five and six, this permits for there to be different flavours of patho-logic in malware analysis and detection, one that is 'static' and about containment, whilst another that is 'contextual' and about behaviour. Yet, in endpoint detection, the sovereign power of the 'right to kill' malware is essential.

In biopolitical security regimes, the body is taken apart and risk management becomes central too. As Amoore and Hall (2009, p.451) write, at contemporary border regimes through the airport body scanner, a dissection takes place in the conviction that if the body is partitioned, it will somehow find and reveal a certainty. Yet rather than calculating risk through the analysis of population-level data (as in disciplinary regimes), it is increasingly around the "individual management of the genetic risks peculiar to one's own body" (Braun, 2007, p.6) for instance. There are two operations; one between the greater collection of data to comprehend the norms of broader data environments (Parisi, 2017), and a second, to drive to manage those risks in an individualistic fashion, so that there are tailored attention to bodies and outputs. Thus, biopolitics is not simply about the control of populations (to generate the norm) but also how to regulate individual bodies (or environments) in order to generate modulating and moving norms that allow prediction and anticipation of the future.

The biopolitical imagination has had a significant impact on how malware are treated – whether through an increasing medicalisation as things become mathematical, to how this intertwines with disciplinary practices and threats to sovereignty (of not only the human, but

also the state), and how this permits the dissection and observation of malware in the contemporary MAL. However, in pursuing a more-than-human reading of malware, this is not a biopolitics that is necessarily directed in one way, but one in which malware ‘speak back.’ Malware are constructive of a biopolitics in that they make choices, so are not necessarily neutral, but an actor in its negotiation. In security, and in MALs, these logics of both dissection and observation are all to attain a sense of (ab)normality, to work out what the difference is between different things, and then act upon this knowledge. But in order for a biopolitics to be effective, and for its impact on malware to be effective, it is imperative to consider how the normal and abnormal come together to delineate lines, demarcations, and differences that permit such a pathologisation of politics.

The Normal and the Pathological

“[D]isease has a land, a mappable territory, a subterranean, but secure place where its kinships and its consequences are formed; local values define its forms.”

(Foucault, 2003b, p.183, (1963))

Foucault reminds us that normality, abnormality, and the pathological must emerge, or coagulate, at certain sites and locales – and this is no different for malware. Georges Canguilhem’s *The Normal and the Pathological* (1989, (1966)) can be twinned with this to explore the crucial distinctions between normality, abnormality, pathology, and anomaly. These distinctions present in medical-biological-security discourses infused the Sophos MAL, so it is helpful to return to how these are understood within medicine. Canguilhem argues (as Foucault does later) that pathology can never be scientific due to its complicated relationship to the normal; it is socially constructed. By this, normality itself is incessantly moving, variant, and there must be deviance from a statistical normal in order for people to be healthy – a wholly statistically normal body is a fallacy. Thus, to be pathological, deviances must be constructed, have a *pathos*, whereby suffering, impotence and abnormality must be ascertained.

These deviances from norms that are not yet deemed abnormal are anomalous. The anomalous derives from a different etymological root in Greek to (ab)normal, it is instead “uneven, rough, irregular, in the sense given these words when speaking of a terrain” (1989, p.131). I argue this sense of *terrain*, as a spatial differentiation is to be given particular

attention. What is it about the spatiality of the anomaly which does not conform to the sense of the abnormal? I argue it is due to the uneven, rough, irregularity that emerges due to a deviance from a statistical norm, i.e. the dissimilarity from the average. This is on a different logic to the normal / abnormal distinction which reflects social value. Abnormality “implies reference to a value, and is an evaluative, normative term” (*ibid*). Thus, the normal and abnormal become evaluations through which societal meaning must then be translated from the anomaly. This is not always aligned as not all anomalies are signals for abnormality. As Canguilhem notes, “an anomaly can shade into disease but does not in itself constitute one. It is not easy to determine at what moment an anomaly turns into a disease” (1989, p.140). Therefore, something interesting is at work between what can be considered *normal* and *abnormal* and its relation to the anomaly. The former refers to acts of social judgement, whereas the latter is based on the mathematical. Thus, as I discussed in the previous section, the anomaly, as a mathematical and calculative product, becomes part of a politics that is already pathological.

The reason I turn to Canguilhem – and in part Foucault – is due to how the normal and abnormal in both security and medicine become intertwined. For example, in the *Abnormal* lecture series (2016), the abnormal and the connection between anomalies change over time – such as the figure of the hermaphrodite in Foucault’s descriptions in France. In cybersecurity, this is important where anomalies and their connection to abnormality is not straightforward, but change over time dependent on context. The environment, the anomalous *terrain*, is not some neutral backdrop on which malware operates. Ecologies actively (re)construct malware, which requires incessant work by a collective of humans and more-than-humans. When explicitly turning to the MAL, Canguilhem’s insights on the laboratory become vital. He says that the “laboratory’s conditions for examination place the living being in a pathological situation from which, paradoxically, one claims to draw conclusions on the weight of the norm” (1989, p.146). Therefore, the spaces of the laboratory, are already *patho*-logical – seeking to distinguish abnormality based on norms. This is no different in malware analysis. Computational ecologies are incessantly being (re)formulated, new programs, new translations, and choices being made. An attentiveness to *a politics* of malware requires looking at how ecologies are essential in forming norms. Even if I do return to biology, *to pathology*, for thinking through abnormality and anomaly; it is crucial to how malware analysis and detection have come to deal with an increasing use of the anomaly as it

turned to mathematical solutions, whilst also maintaining that abnormal on which maliciousness comes to be defined. This is not a simple process as I identify in chapter six – where there is no easy method to understand malware from human expectation, based on normative assumptions - that are themselves pathological.

Defining Malware

Pathology, the normal and the abnormal, have had a distinct impact on how malware have thus come to be defined. Unlike the architectures I explored earlier, software deemed malicious became defined, categorised, and rendered knowable in order to derive order. The broad definitions in Table 2 were published in a 2009 advert in the New York Times, based on neatly defined behaviours, which, over time, have fallen away as emergent, ‘blended’ threats have taken centre stage (Hughes and DeLone, 2007). As the spatiality of computation has changed, innovative malware architectures have been written by hackers to avoid detection, making easy-to-write summaries near to impossible, not unlike our more extensively-researched human societies. For example, WannaCry and (Not)Petya somewhat incorporate two ‘distinct’ architectures – the worm and ransomware²⁴ – that were combined to produce extremely effective propagation mechanisms that encrypted files on computers – causing mass disruption in the spring and summer of 2017. Yet, even as neatly defined behaviours and categorisations have fallen away, the virus and worm still remain central to popular understandings of malware.

Table 2: Taken from Machrone (2009), in an advert posted in the Wall Street Journal by the anti-virus company Trend Micro

| Malware Form | Description |
|---------------------|--|
| <i>Malware</i> | Umbrella term for any malicious software intended to do your computer harm or subvert it for use by others. Includes viruses, Trojan horses, keyloggers, adware, spyware, scripts, worms, etc. |
| <i>Trojan horse</i> | Software that pretends to be one thing, perhaps a useful download or utility program, or even an amusing picture, but with a malware payload. |
| <i>Virus</i> | Malicious software designed to spread itself to other machines. Viruses may simply destroy data or may install adware/spyware or bot software. Typically spread by email attachments or downloads. |

²⁴ A form that encrypts files or other important information for computing functionality.

Worm

Malicious software designed to spread itself to other machines. May simply destroy data or may install adware/spyware or bot software. Typically spread by Internet/network connections. Worms can spread with amazing speed.

These innovations have frustrated attempts to develop common naming conventions. As Mundie and McIntire lament, “nowhere in the cybersecurity community is this lack of a common vocabulary, and the problems it causes, more apparent [than] in malware analysis” (2013, p.556). Though attempting to classify malware according to broad-based categories has now been broadly abandoned, this still is part of the popular imagination, and is influential once malware encounters knowledges beyond the MAL, look at any media story on malware to find this. But this is not necessarily exactly a recent phenomenon. This joke was played out by Spanska, in questioning what his malware was:

“Is it a virus, a worm, a trojan? MOUT-MOUT Hybrid © Spanska 1999.”

(Szőr, 2005, p.63)

Malware are complex, multifaceted forms – one where behaviours and structures can intertwine. To speak of only virus, worm, or Trojan horse neglects this complexity. Others, in an attempt to bypass such reductions, such as the Tallinn Manual, have not attempted to define malware *per se*, but to step further away to talk of malicious *logic*²⁵. Yet popular culture and texts such as by Parikka (2016), that focus on the ‘viral’ do not attend explicitly to the materiality of contemporary malware. This supports Hughes and DeLone’s broad critique that, “discussions of computer viruses and other types of malicious software, which often do no more than rehash basic definitions” (2007, p.80). In this thesis, I acknowledge that these definitions are important for understanding and communicating. Yet, I also expose the complex relationships that contemporary analysis and detection have with malware – that attempt to maintain order through families and *genes* rather than high-level descriptions. Sometimes these lead to names that, through ‘automated’ machine learning detections, assign random names that do not correspond to conventional human-directed analysis, identification, and classification.

²⁵ “Instructions and data that may be stored in software, firmware or hardware that is designed or intended adversely to affect the performance of a computer system. The term ‘logic’ refers to any set of instructions by they on hardware, firmware, or software, executed by computing device. Examples of malicious logic include Trojan Horses, rootkits, computer viruses, and computer worms. Firmware comprises a layer between software (i.e. applications and operating systems) and hardware and consists of low-level drivers that act as an interface between hardware & software” (Schmitt, 2013, p.260).

Families and Genes

In MALs, classification is provided by the grouping of attributes into *families*. This no doubt draws on taxonomy to derive order and delineate how malware are related to one another. As Bowker and Star note, these classificatory practices in computation are likely to have begun with Leibniz (Wiener’s inspiration for cybernetics) and are connected to how medical classification operates and thus we see the influences of early thought percolating contemporary practice. Yet these focus on behaviour, similar code structures, and intent. As Lucas from the Sophos MAL said:

“Yeah, I think probably the strongest, you know, [reason] for making a new family is what it does at the end, what does it do behaviourally? What is the goal or target of this, what is the end goal?”

(Interview, 20 July 2017)

In this, the end goal is not the human actor, but the analyst’s work is to assess malware and its variants to find what its behaviours are. This is not true of all malware analysis, with some businesses actively conducting ‘threat research’ to develop assessments of authorship – but in an endpoint protection business such as Sophos that focuses on providing protection, this was not the case. Families are used to stitch different forms together as broad coalitions of similar behaviours and structures. For example, a malware form may behaviourally infect a file or computer with a particular set of techniques – lending from, and mixing, many of the old conventional architectures together. Then analysts will explore their (dis)similarities and attribute these to an existing family or create a new one (often extensively aided by contextual guides in MALs).

This is complimented by a genetic, ‘DNA’-inspired model of medicine, similar to *bioinformation* (Parry and Greenhough, 2018), *the gene*. During my time at the MAL, the gene was used to piece together detections – which could be simple snippets of code, certain frameworks, connections, or on what packer²⁶ a piece of software used. These genes search for certain attributes that suggest a difference, and thus, valuable information – some were better indicators than others for anomaly – i.e. a certain packer, or a technique that is rarely used in *clean* or *legitimate* software. These genes collated phenotype behaviours – where

²⁶ Packers stitch a program together so that it can become an executable and can be used for encrypting and reducing the size of a program.

collections of genes were brought together in detections by analysts to determine whether something was likely malicious or not. This phenotype based on the ‘genetic’ coupling mean that the analyst could infer some knowledge about the software under their gaze. Genes became operationalised in pathological practice to detect, not unlike in medicine to transform anomalies into abnormalities.

Connections between families and genes became interesting during a discussion with Daniel:

“So, I asked what makes a family? He responded that these are behavioural detections that assign families according to what it is doing – is it targeting certain things, using similar techniques to gain entry, etc. Hence families are not ‘gene’tic, but are based on behaviour.”

(Research Diary, 24 January 2017)

They may not be ‘genetic’ in the sense of looking at code like DNA, but the inspiration to develop phenotypes to construct detections²⁷ in turn delineates a pathological classification to families by weaving genes together. This is not to say that there is only one gene or phenotype for each family – in fact there were frequently multiple, varying detections over time. Yet, as terms such as *genes*, *phenotypes*, and *families* are used for the practice of malware analysis and detection, it demonstrates how biological and pathological traits *live* within malware analysis. As I explore in chapters five and six, these pathological comparisons to biological categorisation, in different orders, to phenotypes, genes, and families develop a biological image of malware as akin to animals and bacteria that can be defined, delineated, and categorised in order to render them knowable and thus condense the immense complexity of malware.

Malware do not have a simple history. Their architectures have been seen as potentially beneficial, as inspirations to military-themed gaming, and artificial life, to them being pathologised as part of a broader movement in the late twentieth century, as security was reconfigured as the spatial dynamics of computation moved from limited installations to the home. In this process, malware became categorised in order to render it knowable, to be compared and contrasted, and to categorise this malicious other – initially through the figures of the virus, the worm, and the Trojan horse, among others, and then through the more complex contemporary practices of delineating families and genes. This complexity, of

²⁷ These were used on static signature detections, certain forms of behavioural detection, and used for grouping and visualising malware detections.

attempting to categorise and define malware, feeds into broader discourses around the ability to (biopolitically) govern this political actor. However, along with the growth and changing spatiality of cybersecurity, anomalies have attained greater importance as calculative techniques have been increasingly adopted to deal with the growth in malware.

Anomalies and Algorithms

The detection of anomalies has also become the distinctive (algorithmic) calculative technique in the age of big data and machine learning. Louise Amoore (2014, 2011, 2009) has written extensively around how data, algorithms, and the translations of these into anomaly and norm operate. This is brought out in *The Politics of Possibility* (2013) that explores how risk and security have become conditioned through the use of algorithms, which transform the present based on the possibility of a future threat. This is extended in *Cloud Ethics* (2019, forthcoming) where algorithms are productive of an ethics *extended away from the human*. That is, there is a more-than-human agency to algorithms; choices are made by computational systems that are imperceptible (and arguably unattributable) to their human authors. This has implications, that I discussed in chapter two, for understanding who and what ‘acts,’ which in turn, is in part conditioned by the treatment of malware as a pathological threat; where if only more data could be gathered, then we could reveal the ‘secrets’ of malware’s existence.

As others have written, machine learning and algorithms are part of the broader security techniques of pre-emption (Adey and Anderson, 2012; de Goede and Randalls, 2009; Adey and Anderson, 2011; de Goede, 2014). These pre-emptive techniques and assessments of the norm are no longer based on conventional statistics, instead, anomalies have become enrolled “not [as] moments of lack or error, but a teeming plenitude of doubtfulness” (Amoore, 2018, p.3). By this, anomalies can be considered a currency upon which decisions become and are made known. A doubtfulness where anomalies become a form of *insecurity* – having to tie together various derivatives of data. The unknown, the potential of anomalies to reveal something becomes part of a pathological and therefore security potential. This is based upon the pathologisation of security – as Foucault demonstrates – that permits, through a movement of a biopolitical practice, a transformation of society where anomalies become central to attend to the new formations of cybersecurity.

This focus on anomalies reconfigures how we derive normality and abnormality, and thus how we even define malware. There has been a focus on averages and deviation in conventional statistics, towards defining the normal based on similarity and dissimilarity in contemporary algorithmic forms such as neural networks (Aradau and Blanke, 2018, pp.11–12). This has importance for the relationship between space and security. This is the difference between focusing on differences of kind and differences of degree – it is a different spatialisation – perhaps Canguilhem’s *terrain*, where malware becomes disaggregated and stitched back together. For example, in neural network algorithms in malware detection, there is the formation of ‘feature spaces’ through the abstraction and contextualisation of information about software. These feature spaces are where neural networks perform calculations which identify features between different ‘layers.’ Through this process, the algorithm identifies ‘closeness’ in these feature spaces – that through similarity or dissimilarity (Aradau and Blanke, 2018) – can identify the likelihood of something as malicious or not. It is this use of the anomaly, through increasing use of big data and algorithmic intervention, that turns on its head conventional statistics that Canguilhem based his discussions on the analysis of deviance.

Here, the difference is where information is gained and enables for distortions that are not necessarily based on a normal/abnormal binary – but how similar or dissimilar this software is to which the neural network has ‘learnt.’ This is common today in contemporary practice, that follows on other areas of security – such as in finance (Morris, 2018) or through the ability to use anomalies to construct a new (ab)normality (Amoore and de Goede, 2008). Therefore, we have an alternative pathological imagination, another path, to that of Canguilhem, which is about disease in place, a mappable quality that Foucault identifies, to one that is moving, abstracted, and is more interested in (biopolitical) circulations rather than whether something *is* malicious in itself – is it similar enough to previous malware or not? Yet, this still requires an actualisation in certain places, and the MAL is one of these places. Thus, we have multiple patho-logics in operation: first on a basis that could be considered taxonomic, cataloguing malware into groups, to one that is ‘genetic’ and behavioural – understanding software in action and collating anomalies to delineate maliciousness. These, over time, coalesced and actualised through anti-virus that is a quintessential expression of a pathologisation of security and politics.

Anti-Virus

During the 1980s in response to the growing, perceived threat of harmful viruses and worms (including logic bombs²⁸ among others that have dropped out of modern usage), anti-virus became the tool to secure ourselves. Its pathological imaginations emphasised the maintenance of cleanliness and containment that resonates with how we have tended to treat human health and hygiene (there are frequent references to maintaining one's 'cyber-hygiene') to eradicate 'dangerous' bacteria in a prophylactic desire to prevent infection. Anti-virus came in many forms, but ultimately, they were all designed to detect, and frequently *disinfect* or *inoculate* computers from the proliferation of new malware forms. As Rob Rosenberger reflected:

“The media’s perception of viruses took a dramatic turn in late-1988, when a college student named Robert T. Morris unleashed the infamous “Internet Worm.” (Some trivia: Morris’s father had a hand in the original Core Wars games.) Reporters grew infatuated with the idea of a tiny piece of software knocking out big mainframe computers worldwide. The rest, as they say, is history.”

(Scientific American, 1997)

The fear of the virus, and later malware, is essential to understand why anti-virus grew as both a product and as an industry. Contemporary endpoint detection engines, that developed out of anti-virus, work in the background of a computer (checking, monitoring, matching and so on) as part of a privatised, entrepreneurial, response as computers entered corporations and homes outside of those in the military and early computational industries, as new spaces required protection. As Jessica Johnston writes:

“The antivirus industry, as both a knowledge industry and technological profession, produces a commodity to manage and regulate one aspect of network vulnerability. Network vulnerability is explained through embodied metaphors of biological viruses and threats of widespread pandemics, grounding the threat in a larger culture of fear, health scares, and biological virus outbreaks.”

(Johnston, 2009: 43)

²⁸ Hruska (1990, p.16) says “a logic bomb is a program which causes damage when triggered by some condition such as time, or the presence or absence of data such as a name.” For example, the 2007-2008 Conficker worm had an ‘activation’ date of 1 April 2008 (April Fool’s Day). This caused much worry in the cybersecurity community but there is no evidence that anything was ‘activated.’

Therefore, one cannot situate anti-virus, today's endpoint protection, outside the pathological – it is imbricated and productive of such senses of vulnerability of pandemic, outbreak, and infection. These notions feed into the formalisation of how malware comes to be understood, categorised, and defined. Unlike earlier experimentation with malware architectures, those who came “to define the legitimacy of software was not restricted to security professionals, administrative personnel, and antivirus researchers, as it came to be at the end of the 1980s” (Parikka, 2016, p.11). Pathologisation is what increasingly took over as a dominant paradigm of response (say, compared to experiments on benevolent viruses and artificial life) that, through technologies to detect, led to a need to classify and render knowable the growing number of malware forms from the late 1980s onwards as I detailed earlier. Anti-virus engines and analysts then became, through the concentration of capital and knowledge in malware analysis laboratories, arbiters of negotiation, analysis, and curation of how software are malicious – that not only define these more-than-human actors, but also act as arbiters of *what is malicious* at all, shaping the contours of ‘normal’ software, practices, and discourses.

As an early book on anti-virus by Jan Hruska²⁹ notes, there was an increasing realisation of how the threat turned *serious* from the ‘fun’ of previous ‘vintage’ malware. “Unlike older viruses (1986/1987) vintage) which would place a silly message or bouncing ball on the screen, new viruses are highly destructive, programmed to format hard disks, destroy and corrupt data” (Hruska, 1990, p.13). Anti-virus developed at a point where the risks from malware became apparent. Rather than something playful, experimental or *harmless*, they become intentional, linear, and *harmful*. When Hruska's book was completed in spring 1989, he said that there was a (comparatively large) increase of separate viruses from seven to seventeen (Hruska, 1990, p.14). Though by contemporary standards, these numbers are miniscule, it signals the increasing worry that accompanied this growth. From this period many popular, and well-known, anti-virus programs came to the fore as computation became securitised – where these programs and their associated modes of production provided protection from these new threats. As Parikka (2016, p.21) explains, they were connected to the “increasing importance [of] software and network computing played in a post-Fordist culture. As networks and cybernetic machines started to grow into prime infrastructures of global

²⁹ Hruska holds a PhD in medical engineering from the University of Oxford and was a founder of Sophos, where the fieldwork for this thesis was conducted.

capitalism from the 1970s on, the need to control such environments became a key concern.” Thus, anti-virus became part of a response to *controlling environments* by expulsion and the sovereign right to kill. Whether this be a corporate network, a home PC, or today’s IoT devices.

Table 3: An overview of different antivirus products available, based on Hruska (1990).

| DETECTION TYPE | DESCRIPTION |
|--------------------------------|---|
| <i>Checksumming</i> | Uses a cryptographic algorithm to give a unique file identifier. |
| <i>Scanning</i> | A ‘virus pattern’ or signature created on attributes within a piece of software. Patterns are then searched for through scanning a computer. |
| <i>Monitoring</i> | Searches for access to certain parts of a computer and <i>determines</i> whether this is likely to be a malicious activity. |
| <i>Integrity Shells</i> | A ‘look-up’ of a software checksum to check it matches with the ‘real’ value, if it does not then the software will not execute. |
| <i>Access-Control Products</i> | These provide an ability to only allow certain programs to run on a computer, this is sometimes also referred to as <i>whitelisting</i> . |
| <i>Virus Removal</i> | This can include the <i>capture, quarantine</i> , or removal of malware. |
| <i>‘Inoculation’ Software</i> | The antivirus program maintains a database of attributes of files (such as size of the file and the date of its creation). Software is then scanned to check if it matches the information in the database. |

There have been multiple changes over time in anti-virus engines – that run on endpoints, typically a computer – but there remain some fundamental techniques. Although these have changed in focus with the availability of big data and computational capacity, there are two main forms of anti-virus in the traditional sense: non-specific (generic) and virus-specific. Hruska identifies seven main forms under these high-level categories of engagement (see Table 3). Although these uses have gone in and out of fashion, scanning and monitoring continued to be an important part of contemporary endpoint protection. In the next chapter, I explore ‘static’ and ‘contextual’ strategies of analysis and detection that broadly relate to ‘virus-specific’ and ‘non-specific’ respectively. Each have a particular form of pathologic; whether by dissection (static) or observing and monitoring the attributes of software (contextual).

Contemporary endpoint protection incorporates technologies such as reputation scoring, use machine learning (typically referred to as ‘nextgen’ or next generation products), and perform different forms of dynamic analyses and detections. Yet, one thing that has been a constant feature is the tension between specific and non-specific detections. That is, the more specific a detection becomes, the more likely it is to be *accurate* but lose the ability to ‘catch’ more malware.

“The length of the virus pattern is also of crucial importance. If a short pattern is used, the chances are that the scanning software will produce a number of false positives, finding the pattern in completely innocent software. If a long pattern is used, false positives will be reduced, but on the other hand the incidence of false negatives will be increased since any virus mutation will have a better chance of not matching the pattern, and hence slipping through the net.”

(Hruska, 1990, p.75)

Within virus-specific, static ways of conducting analysis and detection, *polymorphic viruses*³⁰ introduced issues for conventional anti-virus signature ‘virus pattern’ detections in the early 1990s. Viral *mutation* made it difficult to write new detections at each iteration – at the speed and quantity required to capture new variants. However, these difficulties spurred innovation that blended a range of techniques including emulation³¹ developing after Hruska’s categorisations. Yet this movement to more contextual approaches introduced anomalies that do not mean abnormality – and greater difficulties in ‘accurately’ ascribing maliciousness to software. Though new methods and techniques may have emerged, endpoint detection still centres around the central issue of the specific *vis-à-vis* generic approach.

The growth of malware have transformed how anti-virus is now regarded as endpoint security, due to the growing rate of suspect software and multiple techniques *beyond* static conventional anti-virus methods that they typically denote. This has encouraged the growth of contextual strategies of behavioural monitoring and observations that in turn has transformed the practices of engagement where the *content* of malware becomes less

³⁰ Polymorphic viruses are a category of viruses that can modified versions of itself, and use a different encryption keys on each iteration, making a previous copy of itself indistinguishable to a scanner – making it far harder to detect based on signatures.

³¹ Emulation is a complex process of running through the software code which, in my reading of most common emulation does not actively execute but walks through each step. Some may claim that emulators are not distinct from sandbox environments, but I disagree. Emulators are usually much more limited and do not seek to fully execute malware frequently (often to just a point to decrypt software, or find other information for analysis by signatures, for example).

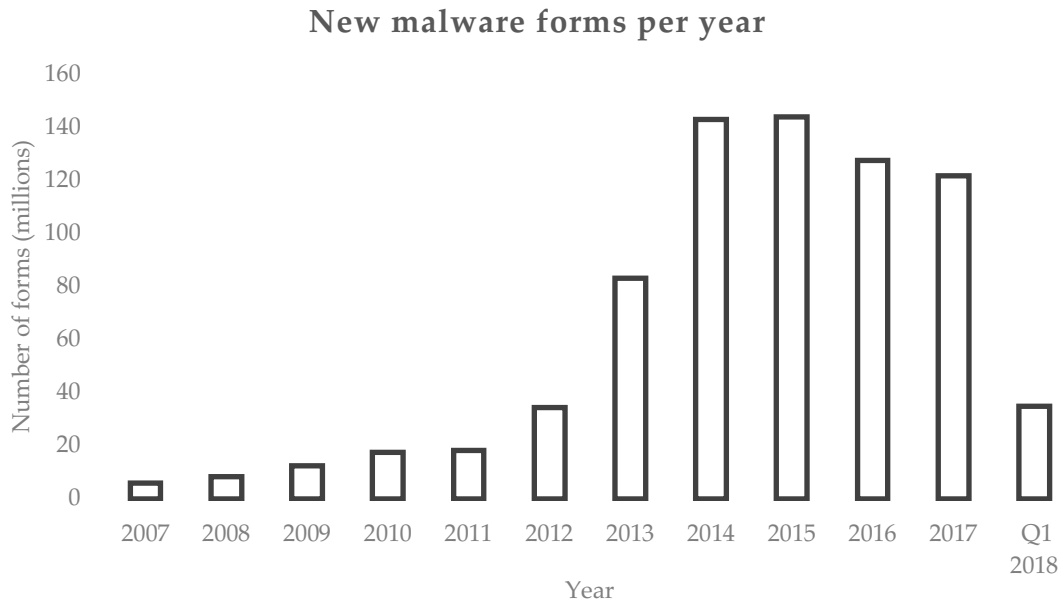


Figure 4: New malware forms over the past decade using data from independent anti-virus product testers, *AV-Test* (2018, p.2).

important as less detailed analysis is performed, not unlike in other security practices, with movements from disciplinary power to biopolitics. Anomalies have permitted a different engagement and capacity to deal with malware, as the speed and volume of malware have increased. As Joe, an analyst at Sophos, told me:

Joe: “Historically it used to be the case you’d see the files; you’d do a quite deep analysis on that file and make some decision and classifications and then that information became public through detections. But that amount of analysis that we do today is a lot less than what we did 10 years ago.”

Andrew: “Hmm.”

Joe: “Just because there’s too much stuff and we’re looking at the different stages in the attack and so what the payload looks like might not even matter for the types of protection we would write.”

(Interview, 21 July 2017)

The growth of malware forms (see Figure 4 for an overview in the decade up to 2018) can be disputed over how one *classifies*, however the general trend is clear (e.g. are they unique ‘families’ or simply variants?). Contemporary endpoint security has to respond with ever-greater speed and response to the growing number of malware forms and spatialities of computation in order for it to be a viable product – changing what pathological strategies should be pursued, incorporating both static approaches to the observation and monitoring of behaviour and movement. However, a movement away from “deep analysis” has brought

out an inability to sometimes consider where problems may arise and the interdependencies that exist in modern malware detection.

Anti-virus, endpoint protection, and their associated businesses are extremely important parts of contemporary cybersecurity spaces. In this, the categorisations and classifications they produce are essential to understand how malware are responded to. As an interviewee from the European Union Agency for Network and Information Security (ENISA) said:

“If I was to compare the AV industry to traditional media, I would say that AV companies in the past and even now they play a huge role in hard facts: they tell you this sample is infected, this sample is not infected. It’s suspicious, it’s not suspicious. So, they mostly focus on hard facts, right now because of this new environment... you have something similar to the new media explosion, in the sense that you have a young independent analyst promoting their report, their views and so on.”

(Interview, 18 October 2017)

In this interview, there was further discussion about a destabilisation of knowledge as individuals publish unverified blogs, suggest ideas on Twitter, and how ‘true’ information was in the 2017 WannaCry/(Not)Petya malware attacks;

“There was the mixture, I don't know how to call it, I would call it even tabloid feeds, or borderline tabloid, in some cases. We had two individuals in the taskforce that were mostly looking for facts, and they were checking all the information being shared by various groups of people that we were following on Twitter and blogs and so on. And sometimes finding the original source of a particular claim was extremely, extremely tedious.”

(Interview, 18 October 2017)

There is a value placed on the productions of MALs and their endpoint protection engines; on what and how they define malware and how this becomes translated, and curated, for consumption beyond these spaces. This is a negotiated practice, where malware frequently intertwines with geopolitics, and balancing a likelihood of probabilities, which is similar to work by Tanczer, Brass, and Carr (2018) on understanding the formation and negotiation of politics of cybersecurity in different organisations. There are histories of the endpoint security industry that intersect with capital, pathology, and various techniques for analysing, classifying, and detecting malware. This is not simply a given, but endpoint security actively co-constructs what is a threat alongside software, environments, and hackers. In working with malware, new forms of security discourse emerge, dependent on the ecologies of its emergence, making endpoint security, or anti-virus, a central infrastructure of (cyber)security

– one that translates and works with a variety of lineages in order to detect and render software as malicious.

Contemporary Approaches

This chapter has focused on the lineages that inform how we understand and respond to malware today. This means that in understanding *forms* of malware, it is not only about their technical nor discursive practices, but their interplay. In the next chapters I explore how these lineages become activated in the contemporary MAL. This requires exploring how more-than-human (cyber)semiotic regimes intersect with the changing spatialities of computation and thus practices of cybersecurity, through to different regimes of politics and security as well as the pathologies that underpin malware analysis and detection. How we structure our responses based on senses of the biological virus, through wishing to make sure our computers are *clean*, to the containment and extermination all affect how we see malware and its politics.

Hence, as I shall turn to, how has Sophos developed endpoint security protection and practices, that have developed on the conventional ‘anti-virus’ products that emerged during the 1980s? Sophos is in the business of creating endpoint detection engines that are programs that ‘hook’ into a computer’s processes to monitor, scan, analyse and detect malware³². Yet unlike with the ‘bugginess’ sometimes associated with early anti-virus, contemporary endpoint engines often do little to make the user aware of their presence. They sit in the background, *quietly* analysing computational environments, software, and code. Even I use the Sophos endpoint detection engine – and in Figure 5 is a screenshot of my MacOS ‘Menu Bar.’ In fact, we are all encouraged to use endpoint protection, making it one of the most pervasive security architectures globally. The white shield with an ‘S’ inside is the Sophos Anti-Virus running on my computer. This simple white icon that I barely recognise, or in Sophos’ program with its blue hues, suggest that this is an innocuous, banal thing. What I intend to do is offer a reading of the complex more-than-human politics and lineages that permit such a banality. In this process, I also register how the practices of endpoint security

³² Not all engines perform the same functions and each endpoint detection business or vendor provides alternative forms of detection techniques.



Figure 5: A screenshot of the Menu Bar on MacOS with 'Sophos Anti-Virus' presented as the small icon shaped like a shield with an 'S' that persistently remains whether the program is 'open' or not. Author's own image.

inform the patho-logic of contemporary relations to malware, and the implications for neglecting the role of malware as an actor in our politics.

As cybersecurity and pathology are interwoven and imbricated with broader movements in security; whether it be on the circulations, observations, analyses, detections, or to ensure a *normal, healthy* computer, they all require a certain attentiveness. Not only to the social relations of pathological currents in cybersecurity and our understanding of malware, but in how malware's geographies are presented on cartesian maps that typically show whizzing lines of light (information) with a dark, suspicious background, that are resonant to how militaries attempt to use light as tool of power (Thornton, 2015; Gregory, 2011). Pathology is not only about critically engaging with the lineages of biology, how machines have been thought of, or how the military has influenced how we perceive malware, but also how these have generated certain perspectives towards cyberspace as 'man-made' and thus can be *known* and therefore mouldable by us. By exploring how lineages of disciplinarity, biopolitics, capital, and senses of normality mix under a pathological desire to detect malware; it is possible to see how there is a politics of cybersecurity that neglects a malware politics through its comparison to a lower-status of 'life,' and its control by human authors. Experimentation with alternative perspectives on benevolent architectures may have been apparently lost in this process as it became securitised— but this should not be forgotten as it opens the potential for a re-evaluation through an eco-logic.

Chapter Five || The Anomalous and the Malicious

I spent seven months undertaking (auto)ethnographic fieldwork, detailing the strategies that are employed by malware analysts, the technologies they use, and the logics that construct how software becomes malicious. Through the spaces of SophosLabs, and its malware analysis laboratory (MAL), I explored how pathologies underpin and tie together malware analysis and detection, where medical-biological-military narratives saturate the need to secure and defend. The MAL, in order to detect and sell endpoint protection, has developed strategies, techniques and technologies to thus identify software as malicious. This is a complex process in which data is produced through computation in detection (machine learning), through testing (internal laboratory data), with other MALs (data feeds), and from customer environments (telemetry), that come together through big data to inform both computational choice and human decision. Anomalies – that are tied together to identify maliciousness – are blended with human senses of maliciousness to construct a pathological, *abnormal software*, malware. The spaces of the MAL then become a point of interchange that draw upon noises and signals of data, computation, and human affect that intertwine to sift, sort, and delineate software. This constructs a body of information that allows for distinctions between ‘normal’ (clean) and ‘abnormal’ (malicious) to be drawn.

In this chapter I identify, and expand on, two broad strategies within which these distinctions were made. First is the static strategy that renders malware outside of its environment, separating out a malware ‘object’ to be analysed and classified *post hoc*. Second is the contextual strategy that draws extensively on environments, contextualisation, and software attributes to understand how malware may be identified, at one step removed, through its malicious affects and traces, in an attempt to speed-up, and pre-emptively detect. To investigate these two strategies, I look at the computational ‘tools’ that underpinned these ways of approaching malware, and examine the extent to which these draw on and reinforce pathologies within malware analysis. The transition from the software sample to detected malware, from the anomalous to the malicious, are generative of particular geographies and senses of security that are critical to comprehending our contemporary relationship to malware. In the first part of this chapter, I work through the Sophos MAL and explore how data construct a ‘noisiness’ that feeds signals of maliciousness which are

then picked up and interpreted by analysts and other computational tools. This is then supplemented with a discussion of the importance of the environment to the MAL, and how these are understood with specific assumptions, in order to delineate the signal to craft the anomalous (and the abnormal/malicious). The rest of the chapter is then devoted to a detailed trip around the MAL's different approaches for both analysing and detecting malware through the lens of static and contextual strategies of responding to the threat of malware. All strategies are somewhat dependent on data, that tie these strategies together and are part of a movement to using big data to further delineate (ab)normality through greater dependence on anomalisation. I then provide some concluding thoughts on how both strategies, albeit different, are part of a patho-logic to analysing, detecting, and eventually informing broader societal interpretations of malware.

The Lab

It was -4°C on a crisp, January morning when I caught a train to Radley, a rural train station close to the Oxfordshire town of Abingdon in the UK. I was entering Sophos' headquarters for the third time (I had been three times before; once on a university-organised event, a one-day shadowing with Alex, and once to meet Daniel and Kyra from HR), I stepped into reception, where a photograph of me was taken on a background of a hastily-arranged Sophos-blue material secured to the wall with magnets. I was issued with a temporary pass and greeted by Elliott. He escorted me to the MAL and set-up my computer. My trepidation in starting could not have been higher. I was placing myself in the middle of a highly-specialised laboratory. I had no serious experience of the variety of knowledges, practices, and techniques that were required: of file formats, computer architectures, coding, using the command line, or of low-level 'assembly language' (a human-readable per-instruction command used by software to instruct hardware). This all had to be learnt 'on-the-job.'

My desk was next to the large glass façade that surrounds the headquarters. The MAL was spacious and full of light, with computers whirring in the background. It was not only full of malware analysts in the generic detection team that I joined. There were a whole host of different people in other teams, and in 'quality assurance,' who maintained and controlled the testing of systems to ensure that the lab functioned that could be compared to technicians in other laboratories. Computational tools were essential, including a host of

software; simulated operating system ‘environments’ such as Windows (XP, Vista, 8, 10), Apple MacOS, and mobile environments such as Android and iOS. Then there was the hardware; supporting systems such as air conditioning that kept the servers cool; customer telemetry data produced through analysis performed by the endpoint detection engines; data feeds from third parties; the videoconferencing with other Sophos sites in Australia, Hungary and Canada; Twitter; and a multitude of other more-than-humans that sustained the laboratory.

Much of my experience as becoming-analyst was condensed at my desk. I had three screens, a keyboard, mouse, and my headphones. It was not all that different to how Kim Zetter, the technology journalist, describes visiting a ‘threat intelligence lab’ where “the term... conjures a sterile workshop with scientists in white coats bent over microscopes and Petri dishes. But Symantec’s lab was just a nondescript office space filled with mostly empty cubicles and a handful of workers who stared intently at their monitors all day, mostly in silence, doing methodical and seemingly tedious work” (2014, p.55). It was in this type of space where my becoming with malware happened. The screens on my desk displayed different aspects of my analysis and detection spaces (see Figure 6 for a graphic overview):

- On the left; corporate services (the *green* network, a ‘virtual machine’ (VM)) provided email and ‘safe’ internet access.
- On the right; contextual information (the *yellow* network) contained multiple data feeds, twisted with algorithmic processing from other MALs and organisations; internal ‘wiki’ information pages; search tools for software samples; and data from customer telemetry.
- In the centre; the analysis and detection-writing space (a *yellow* VM) layered on ‘top’ of the yellow network that allowed the VM to access Sophos’ tools.




Figure 6: An overview of how two networks, yellow and green, operated via a representation of three screens. The solid black lines represent virtual machines with the dotted line being based on the underlying yellow network. Author’s own image.

As I soon realised, the Virtual Machine (VM) is a technology of confinement in the spaces of the analysis. The VM is a virtualisation, or simulation, of ‘real’ computation. The VM environment restricts access to the ‘physical’ computer. By this, it prevents, or at least limits the potential of, malware being able to infect an analyst’s computer. This limiting of malicious agency, as malware can sometimes perform beyond the expectation of the analyst, i.e. moving outside a VM, becomes an essential technology in order to segregate environments. The MAL maintained the separation of networks (or access) to enhance this containment (see Figure 6). The yellow network is the specialist analysis and detection network with limited connections beyond the MAL’s internal systems and tools. The right-hand screen, depicted through a dotted line, was not in a VM and was able to access internal laboratory resources, with some limited ability to connect to the internet³³. The central screen, where I conducted analysis, was within a yellow network VM layered on ‘top’ of the yellow network. Processing ‘live’ malware in a VM – that replicated access to tools, but not to the MAL’s internal networks themselves – allowed me to conduct analysis while containing the ability of malware to access my ‘real’ computer host. The left screen was a *green* network VM, connected to corporate services. This was contained and separated to keep corporate networks ‘safe’ from malware that may have moved beyond the yellow network VM onto my computer. This demarcation, where borders, and flows were tightly controlled demonstrate the lengths to which the MAL is concerned about malicious agency.

The Allure of Environment

The MAL demarcated space, separated, and brought together certain materialities at certain space-times. When assessing environments, assumptions were made about the kinds of environments malware *executes* in the wild, *out* on customer endpoints (where endpoint detection engines are); for example with the presence of computing ‘patches,’ the assumption that a (human) user will perform certain actions, and that there will certain logical steps. Not only was there a *normalisation* of software, but also the environments in which they perform. I presented the two broad strategies for analysing and detecting malware that resonated with Daniel during a ‘feedback’ meeting over a year later at Sophos (Research Diary, 15 November 2018): the static and the contextual. These both treat environments differently. Static

³³ This was possible but was warned against due to the potential for malware to be able to exploit this.

strategies analyse and detect malware through ‘non-performative’ action. That is, it is analysed *in stasis*, without executing in an environment. Whereas contextual strategies monitor behaviour or attributes of software. Thus, normality in the former is broadly context-less with abnormality defined by the analyst’s affectual experiences, whereas in the latter, the environment and contextuality itself become enrolled in practices of (ab)normalisation. Both strategies assume a pathology, something abnormal to be found. The result is a border that is incessantly moving. The two strategies provide alternative zones where maliciousness is identified.

The MAL takes the environment as something to be measured, distilled, and sensed not only in contextual strategies, but also in the environments of analysis – the performativity of their ‘instruments’ so to speak. Analysts frequently saw each instance of an environment as ‘fresh,’ as if wiping the slate clean, without stratification. It became a background upon which (malicious) execution takes place. As considerable work in geography and beyond has shown, environments are not simply shaped, but are fragmented, exhibiting agencies and shaping our worlds (Massey, 2005). This is reflected in the MAL’s constant preoccupation with (re)working environments to enable malware to ‘work’ and behave as expected, or not, as it often turned out. The environment was then an essential matter of attention – even if it was not recognised as so by analysts. This was most clearly expressed through the separation and compartmentalisation of the analyst’s computer – with its virtual machines and different networks. To prevent malware ‘escaping’ *sandboxes* – an ‘automated’ VM tool for analysis and detection that I explore at greater depth later – environments must be actively crafted, brought into being through tedious work to ensure they were not *interpreted* by malware as simulations (many malware perform anti-VM and anti-analysis checks before they execute). To see environments and malware as co-produced is essential, otherwise understanding a malware politics is impossible – as each environment is productive of different signs which enable different choices to be made. But environments or malware that did not behave as expected, were frequently seen as errors, something that the analyst (or malware author!) had done incorrectly, rather than environments as being performative in their own right. This is an allure of the environment and ecologies that may be ‘clean’ or ‘neutral’ – not unlike in popular imaginations of cyberspace – but in practice is one of the MAL’s preoccupations – upon which signals of the normal and abnormal can be distilled. Thus, environments in a patho-logical rendering of the laboratory must be kept under control as much as possible, to

be a background upon which malicious activity takes place, and rendered as fully-knowable by analysts to permit not only contextual strategies, but also static strategies that I now work through.

Static Strategies

Static strategies are primarily a human-directed activity (with extensive technological support), and are the conventional form of analysis and detection from the early days of anti-virus, frequently performed on the most common executable, the Windows PE³⁴ file with its extension ‘.exe’. Static strategies can be applied to different programming languages including JavaScript, Python, PowerShell³⁵, and to embedded executables in PDFs³⁶ and Microsoft OLE³⁷ files (commonly Microsoft Word, Excel and PowerPoint), among others. Static techniques have prioritised the Windows PE format, and before this, the dominant MS-DOS³⁸, due to the greater prevalence of malware written for, and found on, Microsoft systems³⁹. In this section, I walk through and analyse both the different tools of static analysis and the common static detection method, the anti-virus ‘signature’ with which I had significant exposure.

Static Analysis

The most common tool used for static analysis of Windows PE files was IDA (the Interactive Disassembler). This is what is termed a ‘heuristic’-based program which reverse engineers software and stitches it back together. So, in a sense, we could say this is not *static* – in that it does not just render software without significant reconstitution. However, it does

³⁴ The Portable Extension (PE) file is the Windows Operating System’s executable. This means it lets programs be run, such as those that end ‘.exe’, or in dynamic link libraries (which deal with functions called by PE files) ‘.dll’. For a more detailed, technical explanation of the PE format, please refer to Pietrek (1994).

³⁵ PowerShell is built upon the .NET framework that allows full access to a variety of features within systems and was designed initially for the Windows Operating System. It is based in the console and allows for commands to be run remotely.

³⁶ The Portable Document Format (PDF) is constructed from a series of different ‘objects’ that are threaded together in the format allowing for the embedding of files.

³⁷ Microsoft OLE (Object Linking & Embedding) files are typically referred to as the ‘Office’ suite of programs. They allow for a sharing of different kinds of data into a document form, such as a picture or to embed another file.

³⁸ An early operating system by Microsoft (Microsoft Disk Operating System) prior to the development of today’s popular Windows operating system.

³⁹ 69.96% of malware in 2016, though with substantial growth for MacOS, iOS, and Android platforms (AV-Test, 2017, p.2)

not execute software, and thus was considered by analysts as static. Some of the earliest versions of these disassemblers, such as *Sourcer*, appeared in 1988, just two years after the first ‘PC’ virus, the 1986 *Brain Virus*⁴⁰. The first version of IDA [v0.1] was released as a MS-DOS console application in 1991 (Guilfanov, 2014) and what was key was the ability to distinguish between data and code, which is no easy task at the lower layers of computation, as both can appear similar, and can be incorrectly allocated. Figure 7 includes a screenshot of the freeware (v.5.0.) version of IDA. The disassembler (de)constructs binary, or digital, representations (which are often expressed to humans, even at a low level, as opcode⁴¹ hex⁴²). During analysis, I infrequently had access to source code which was not obfuscated⁴³ in some way. This is due to the compilation of software and code by compilers⁴⁴ that transform the software into the binary representation so that they are executable. Hence, IDA was required to translate these programs into something I could *read*.

Full of slippage, there are assumptions of code execution, leading to some (un)intended execution paths. For example, where does the file header – a PE format’s ‘contents page’ – end and the executable code begin? IDA constructed execution maps according to the software binary it was examining. This produces a symbolic representation, assembly language⁴⁵, as in Figure 7. This closely follows machine instructions that are based on the most common chip architecture, Intel x86⁴⁶. IDA represents data storage, memory

⁴⁰ Interestingly, the authors of the Brain Virus – otherwise known as the Pakistani Computer Virus (Webster, 1989) – the brothers Amjad Farooq Alvi and Basit Farooq Alvi claim to not have intentionally created malware, as in an interview with F-Secure’s Mikko Hypponen (Brain: Searching for the first PC virus in Pakistan, 2011).

⁴¹ Opcodes are the machine-level ‘operation codes’ that specifies the particular instruction to be performed. For example, the ‘07’ opcode instruction would ‘pop’ (or return) the ES register from the stack based on the backwards-compatible x86 Intel 8086 CPU (Central Processing Unit).

⁴² Hex is the colloquial of hexadecimal, which is a base-16 number system. This is written as 0-F, with each hexadecimal equivalent to 4 bits (a nibble), with two hexadecimal (e.g. 1Bh = 27 in decimal) being half a byte (8 bits, or two nibbles). Four hexadecimal are usually grouped together to form a byte when doing analysis.

⁴³ To reduce the size and as a part of compilation as mentioned earlier, both “clean” and “malicious” files use packing as a way to distribute. This is often to prevent stealing binaries (to recompile to use a problem for free and therefore miss out on revenues) or to prevent analysis. There are many commercial widely-used packers, but also specialist packers that can be both “clean” and “malicious”. However, they both slow down analysis due to a layered process of going through encryption routines which is a slow, and often ‘boring,’ process.

⁴⁴ Compilers are used to ‘package’ software to a binary representation that can operate in different environments. This helps to enable interoperability so that it is not dependent on the host processors that the software is created upon.

⁴⁵ Assembly Language is typically referred to as a ‘low-level’ assembly language compared to ‘high-level’ languages that do not often represent the machine code instructions. Hence, there is a strong connection between the symbolic representation assembly to the machine code instructions. For further reading on the particulars of assembly code, please refer to Chapter 4: “A Crash Course in x86 Disassembly”, pp.65 - 85 in Practical Malware Analysis by Sikorski and Honig (2012).

⁴⁶ For more information read Chapter 1 of Practical Reverse Engineering (Dang, Gazet and Bachaalany, 2014).

locations, and registers that are parts of computing processing units (chips). IDA was able to produce such representative maps due to the logical outputs of computing at the low level of the chip. It was where cuts are made, based on file headers, where execution flow, in loops and ‘jumps,’ determined whether the mapping resembled something akin to the original (as it had to translate how signs were interpreted through the layers of computation). Yet these are always unique iterations. IDA could not cut, for instance, if there was a customised packer that could not be decrypted. If IDA failed to make a cut that followed a particular decryption logic, for example, the machine instructions in assembly code would make no (logical) sense, as the cut could not be *logically* enacted, and required alternative practices to be deployed by an analyst.

The process of ‘stitching’ that IDA performed required making assumptions about software execution and follows this. However, this does not equate to the software sample being readily available to access without this technology. There is a dynamic process with an active performance of computational processes into an appropriate cut. If the cut was incorrect, there would be mistakes and slippages in the representation, that lead to incomplete renderings of the software that I wished to analyse. This is an all too frequent issue within static strategies, causing long, tedious work in teasing out layers of (de)encryption to arrive at an appropriate assumption of the ‘original’ software for instance. What PEView and PESTudio (Figure 7), along with IDA, perform is akin to ‘mapping’ a more-than-human software body. This can be compared to the dissection of the human body, rendering it open to analysis by the gaze (Foucault, 2003b). Security practices at the border have produced scanners that have dismembered and reconstituted bodies for the expert viewer (Amoore and Hall, 2009), providing an ability to render the suspect body visible. As Amoore and Hall (2009, p.446) eloquently put this, “the image viewed by the screening operative at the border checkpoint, then, is not a ‘copy’, but an abstraction and a recomposition of the dimensions and densities of the body.” Taking bodies in their broadest sense, to incorporate software, malware analysis equally tears, stitches, and queries the suspect sample.

Static strategies actively construct software as something to dissect. When working with different samples, I realised that these were only ever representations, cuts in the material arrangements at a specific point in space and time. I was not simply working with a copy either: (malicious) software are always fleeting, escaping our grasp, a cut was made in space

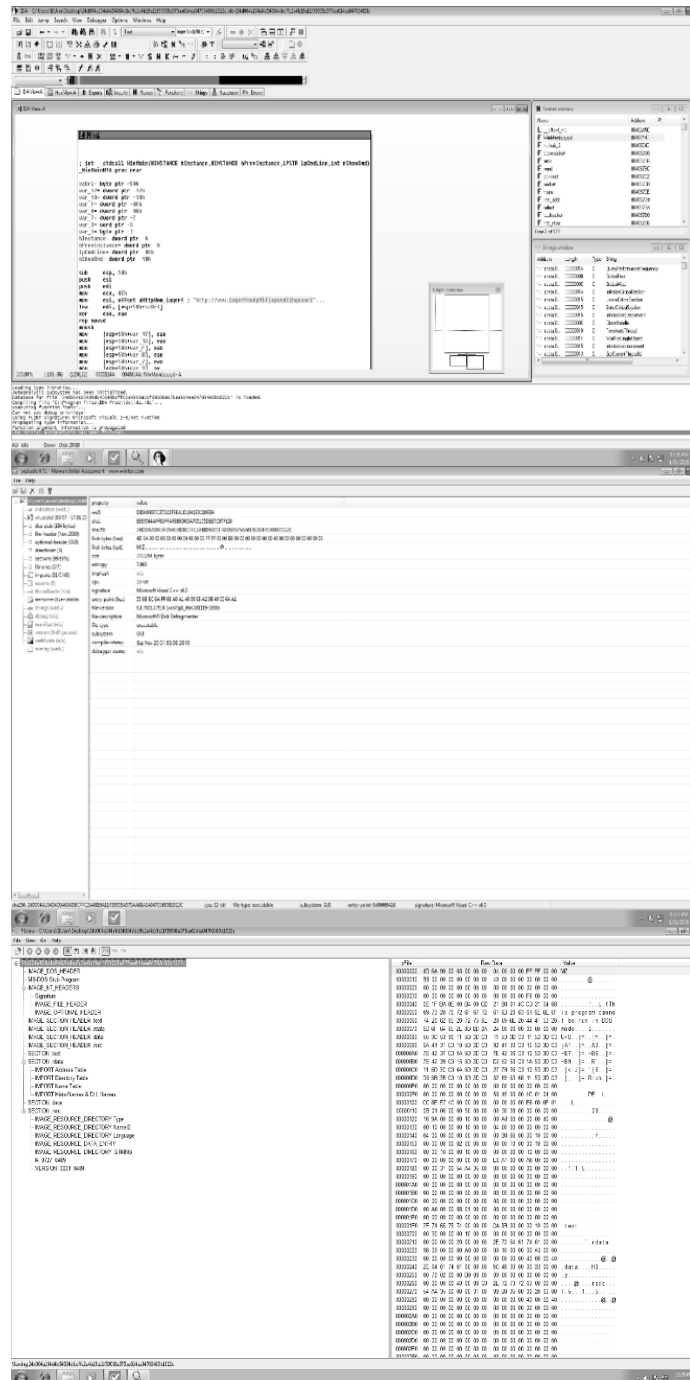


Figure 7: Overview of some 'static' tools. Top-left: IDA Pro Disassembler; Top-right: PEView; Bottom: PEStudio. Author's own image.

and time when it was rendered through IDA. If I had software to view, it *became with* its environment, transforming through its reconstruction in a decompiler, through separation by its file header. The tools I used highlighted certain areas for my attention, with execution paths to follow, dynamic mappings of logical flow, machine operations that jump to other machine instructions, with suspicious areas highlighted red. I engaged with tools not only

built to assess for anomalous attributes but developed an already-present connection to the malicious. Software anomalies were folded with the search for the abnormal to delimitate what should attract the analyst's attention. It is exceptionally troublesome in the contemporary MAL to differentiate between what is *anomalous* and *abnormal*, with maliciousness becoming normalised as existing *a priori*. As analysts created tools, these overlapped with previous experiences of other tools themselves, folding experience of (malicious) analysis. Malicious abnormality became equivalent to the anomalous. When the malware analyst found new anomalies, if they were not found to be abnormal, they are commonly deemed legitimate and therefore became non-anomalous. They were disregarded, and tools must be found for the 'correct' anomalies to provide new metrics or tools to aid with analysis.

Yet, there are significant limits to an analyst performing individualised static software analyses. Growth in new samples and the speed in which software can propagate, change, and duplicate exceed the capacity of human-centred analysis and detection. Significant barriers make static analysis difficult; many samples were packed and obfuscated to hinder the speed and possibility of me being able to dissect and find the software's 'essence.' In the past, analysts had to wearily 'unpick' each sample they received. As one malware analyst recalls as part of his training:

"We spent like the first two years or something in the lab just writing packer detection, which helped me learn assembly and stuff, and get through packers and like understanding detection. Quite but, like, yeah like, it was unbelievably futile so there would be days where I would update like a specific GPK [a packer detection type] like two, three times a day."

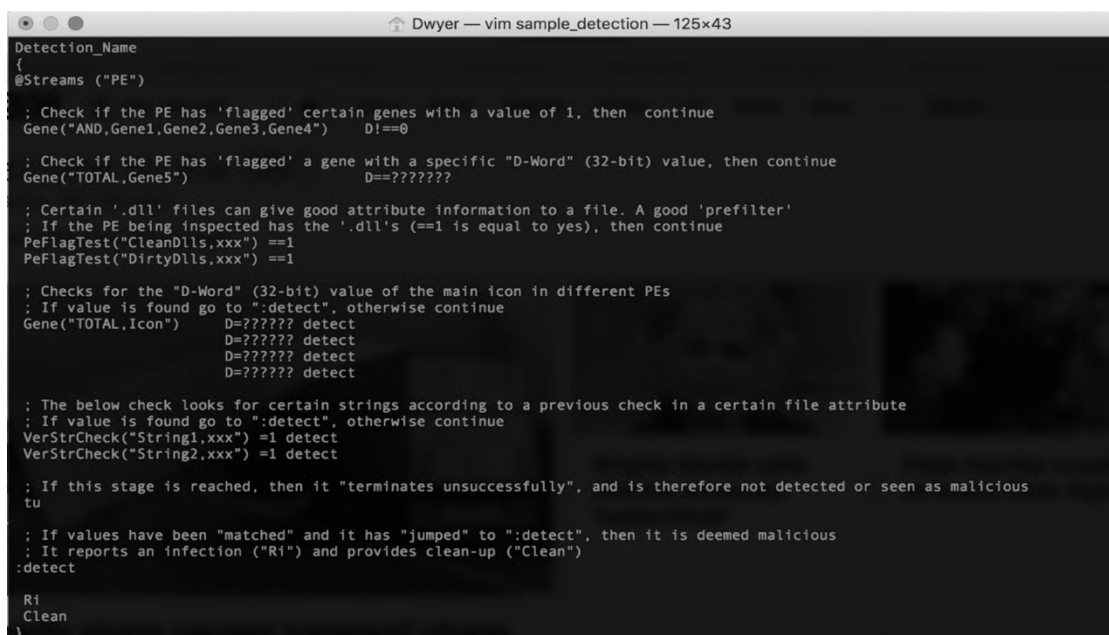
(Nathan, Malware Analyst)

Unique packer algorithms were dismantled, each layer examined in order to *unpack* and *reveal* what had been 'packed' inside. The time to analyse the minute packer transformations over the course of a day (often without the use of contextual information), meant it was "unbelievably futile" as the volume of new software to analyse continued to grow. Nathan and Mason highlighted over a cup of tea in the kitchen how intimate these practices were; where they felt they were fighting hackers, as they tweaked the packers each day, they had a 'relationship' with the human on the other side, recognising their techniques and practices, their *signatures* (Research Diary, 20 April 2017). So, though there may have been a monocity to these older packer analyses, the way their faces and bodies excited, tensed and expressed

joy at telling these stories, it suggests that something is integral to the experience of static analysis – the relation of malware analyst to malware and hacker – as a form of affectual, embodied, capacity to develop senses of (ab)normality. The changes in computational capacity, speed of required analysis, and the greater ability of malware to make ‘choices’ have disturbed this relation, when in the past malware were tightly bound to their authors. This boundedness permitted a greater sense of closeness that has transformed along with the new political economies of malware production – that has increased volumes of malware that are now authored by criminal gangs and nation states.

Static Detection

The relationship the analyst develops during analysis, if found to be malicious, must then be detected through the dominant static detection form: the conventional ‘anti-virus’ signature. A signature is a sequence of instructions that match the software to be detected through endpoint detection engines. Signatures are crafted according to two main methods: one on file attributes (such as the size of a particular section, or on the presence of an image) through genes that I introduced in the previous chapter, and second on specific components of the file (such as a sequence of bytes, assembly structures, and certain network connections). I heavily engaged in the production of signatures as part of the generic detection team, whose central aim was the production of signatures. This blended the methods, knowledges and affects that I had built during the initial four months of training. This included finding



```

Detection_Name
{
  @Streams ("PE")

  ; Check if the PE has 'flagged' certain genes with a value of 1, then continue
  Gene("AND, Gene1, Gene2, Gene3, Gene4")  D1==0

  ; Check if the PE has 'flagged' a gene with a specific "D-Word" (32-bit) value, then continue
  Gene("TOTAL, Gene5")  D=???????

  ; Certain '.dll' files can give good attribute information to a file. A good 'prefilter'
  ; If the PE being inspected has the '.dll's (==1 is equal to yes), then continue
  PeFlagTest("CleanDlls,xxx") ==1
  PeFlagTest("DirtyDlls,xxx") ==1

  ; Checks for the "D-Word" (32-bit) value of the main icon in different PEs
  ; If value is found go to ":detect", otherwise continue
  Gene("TOTAL, Icon")  D=?????? detect
  D=?????? detect
  D=?????? detect
  D=?????? detect

  ; The below check looks for certain strings according to a previous check in a certain file attribute
  ; If value is found go to ":detect", otherwise continue
  VerStrCheck("String1,xxx") =1 detect
  VerStrCheck("String2,xxx") =1 detect

  ; If this stage is reached, then it "terminates unsuccessfully", and is therefore not detected or seen as malicious
  tu

  ; If values have been "matched" and it has "jumped" to ":detect", then it is deemed malicious
  ; It reports an infection ("R1") and provides clean-up ("Clean")
  :detect
  R1
  Clean
}

```

Figure 8: An abstracted signature detection in a proprietary stack-based language, Virus Description Language (VDL). Abstraction taken from author-written detection (Research Diary, 9 June 2017). Author's own image.

structures, strings⁴⁷, decryption routines, certain file attributes, and using and producing genes. Genes identify anomalous attributes and appeared as contextual information, ways to group, connect, and intersect the vast volumes of samples held by Sophos. These would appear as pieces of information if I tested a sample in the lab, being ‘flagged’ by the detection engine. These allowed me to ‘tighten’ my detections (to exclude other *unrelated* software) through combining these genes as ‘prefilters’ to speed a detection up as the commercial drive demands that users should not notice the engine’s running.

In Figure 8, I present an abstracted signature detection I wrote for a malware form that I have adapted to remove any information that could be replicated and used to avoid detection. Here, I outline how this detection, based on *file attributes and genes*, performs. I wrote a sequence of checks that search for certain values, either Boolean (yes or no, 1 or 0) or “D-Word” (32-bit) responses that match a specific cryptographic value. The signature proceeds sequentially, line by line, each check after another. Let me walk through this:

| Step | Action |
|-------------|--|
| 1 | “D!==o” acts as a prefilter to check certain features are present. |
| 2 | “D==?????” checks for specific cryptographic values on an attribute produced by a gene. |
| 3 | I searched for “.dlls” that are libraries of information that had been previously identified as ‘clean’ or ‘dirty’ by analysts in the MAL. |
| 4 | Most files had similar icons, so I searched for four particular cryptographic hashes ⁴⁸ of these icons that, if found, ‘jump’ to “:detect” and ‘fire,’ positively generating a match, and thus, a detection. |
| 5 | If no matching icon was found, I attempted to find one of two strings. This check is only reached if, <i>after</i> filtering, there are files that do not match a specific icon. If neither strings were found, then the signature terminates unsuccessfully “tu” and did not produce a detection. |

This signature demonstrates the folding of contextual information into the contemporary static detection signature. In early, anti-virus signatures, it would have been code snippets rather than genes. This is what distinguishes anti-virus from endpoint detection – the

⁴⁷ Strings are concatenated alphanumeric symbols of any length. For example, one detection I wrote included a check for the string, “old McDonald had a farm” (Research Diary, 25 and 26 July 2017).

⁴⁸ Cryptographic hashes are unique ‘signatures’ of a form of software or other attribute.

availability of data and contextuality. In this, multiple analysts' knowledges and labours are threaded together to produce new detections through the gene. It was rare for a detection to be written completely anew for particular malware families, often developing on attributes by calling upon other genes. Genes thus helped to speed-up detection as it excluded possibilities quickly through filtering software. I needed to *execute and terminate a detection as quickly as possible* to ensure that there was no performance lag on consumer endpoints. Knowing *which* malware to write detections for has thus transformed in contemporary practice. Instead of analysing all malware using static strategies, contextual information provided by genes have folded static analyses and detections on themselves. What comes to the attention of the analyst is mediated through the technologies of the MAL, that blend past static strategies to highlight what may be missed. The signature no longer stands alone like in anti-virus but is tied to other processes of the MAL – this is why I refrain simply delineating contemporary signatures to 'anti-virus' as there are whole new practices present as contextuality takes hold.

Contextual Strategies

Sophos' acquisition of the machine learning business, Invincea, in February 2017, was an attempt to attend to the temporal constraints of static-based malware analysis and detection – as well as a struggle to detect the 'unknown' malware as part of a pre-emptive security logic. As one manager, Oliver, said in a MAL-wide meeting, "actually most malware can be traced back to [MS-]DOS in the early days. However, the things that have changed are the breadth of the attack landscape, and the scale with which to deal with new malware" (Research Diary, 8 February 2017). The difficulties in 'hand'-adjusting malware, its volume and the anti-analysis techniques developed (such as packing and obfuscation) by malware authors explain the reason for the movement to contextuality. This can be thought of as analogous to the growth of bioinformatics in preventative medicine, which use big data to generate normals against which to compare individual data and predict future (ill)health. An assessment is made of the *contexts* of execution and its attributes, in order to establish 'just-in-time' what forms of software may be malicious. Static strategies (still) form a backbone of malware analysis but are deemed too slow and are now only reserved for specialist, in-depth, analysis.

Analysts do not become devalued during a process of contextualisation and computational analysis and detection, but are elevated, however. Not unlike in other labour markets (Agrawal, Goswami and Chatterjee, 2010), the banal (mass) labour of processing the growing volume of malware samples has been outsourced to ‘threat research’ in India as part of a ‘frontline,’ primarily dealing with customer submissions and spam (Leyden, 2014). Hence, the changing political economies and spatialities of malware and their authorship are transforming responses by MALs as part of global neoliberal labour markets, colonialist legacies, and the modularisation of work (Vira and James, 2011). Contextual strategies draw on big data to provide a more-than-human sensing of malware’s attributes and performances. Contextual strategies develop a probabilistic confidence in what *may be* malicious through a new processing of maliciousness based on anomalies. This is a process at greater risk of the ‘wrong’ identification of software as malicious. In exploring the contextual, I outline what I found to be significant approaches – that bring together a variety of analysis and detection techniques that consider execution, data, and attributes produced by malware as central to a speeded-up, pre-emptive pathologic.

Reputation

SXL⁴⁹, the Sophos Extensible List, is a system for retrieving information about software on endpoints that can include ‘looking-up’ in Sophos’ cloud architecture for detections, including to help identify *anomalous* software and code through customer telemetry. The ‘reputation’ system is built through condensing knowledge through analyst-written rules on what would *typically* be malicious. The reputation system performs a series of checks on the attributes of files; such as its provenance, age, and packer type.

“The reputation-based system works by allocating a score based on the likelihood of a file to be malicious... Although this evidently creates some false positives, an attribute-driven form of detection means that one can detect malware that one does not even know about.” (Research Diary, 31 January 2017)

⁴⁹ SXL (Sophos Extensible List) is a method for the retrieval of data from a remote computer and for also performing a ‘look-up’ to check for detection information. For more information refer to Sophos (2017).

The core concern was the development of appropriate metrics to define what is *anomalous*, and thus can be emulsified as malicious through a bringing together of select anomalous attributes.

A movement to attributes, not unlike how an analyst draws together genes in static detections, assesses the value of malware attributes. In producing the static signature, I assembled a range of genes I thought best captured a group of software samples in order to create a unique detection – one that would not ‘clash’ with others. But reputation works differently, it is not based on the affectual, sensed, (ab)normality. It produces scores based on what is deemed to be suspicious, anomalies, as indicators of potential maliciousness, based on rules written by analysts in the MAL. A scoring initially allowed for a binary flag on software anomaly. This requires a contextualisation of what is *normal*: what a *normal* piece of software looks like, what a *normal* computing environment may be. When I questioned this, Daniel scoffed, “who knows what an ordinary computer looks like?” (Research Diary, 21 June 2017). There is no normal computer or environment. It is incessantly modulating and cannot be determined in a singular way. Yet the attributes that have been built by analysts in the MAL are forever playing with environments to tease out anomalies (based on their experiences of malicious abnormality) in order to undertake reputation scoring.

Preference for a scaled reputation scoring system emerged, however. Binary scoring: where one attribute was malicious, and another not, were too restrictive, too coarse to provide useful information that respected that software was contextual in its forms. This is at least a partial appreciation that malicious software may exhibit both malicious and ‘clean’ attributes, as part of contextuality. The scalar scores provided ratings for a ‘risk score’ to be applied to singular software forms. This is not unlike other contemporary security techniques, such as at the border (Amoore, 2011; Pöttsch, 2015; Vukov and Sheller, 2013) that produce, and powerfully and politically, sort and list (de Goede, Leander and Sullivan, 2016). As Pöttsch (2015, pp. 107–108) argues with reference to bodies at the border, “a tracking of movement and association, combined with a mapping of behaviour and affective responses, serves as the basis for an identification of algorithmically determined deviations from a calculated norm that triggers increasingly automated mechanisms of inclusion and exclusion.” Software experiences a similar fate. A politics emerges at this moment – a more-than-human one – found in the relationship between previous decisions made by analysts

and how environments perform in reputation scoring. The MAL envelopes the previous affectual becomings that develop between software and analyst to blend different scores together to produce the colours and figures that appeared on my screens; but by the time they reached me they were stripped of this politics, simply presented as a probability.

A transition to a modulating, moving normal has occurred through the use of the reputation scoring algorithm. As Sophos (2016) describes on its website, “a team of researchers in SophosLabs is assigned to monitoring and ensuring the continual efficacy of the scoring algorithm and the data that affects it.” Attributes produce scores where its affects are emergent, according to a combination of variables, the wrapping of customers’ ‘wild’ environments, and the weights applied in the scoring algorithm. When I looked at my right screen, I often had a cursory, almost impulsive, glance toward this score to consider whether it was *malicious enough* to investigate (I rather fell into the trap that the anomalies collated *were the abnormality*). Software packers, where they originate, and the age of the file require telemetries and feeds. These data and scoring algorithms enabled me to look to contextual information on reputation before engaging in static analysis – presenting a looping of static and contextual information that shaped where my attention was directed towards.

Behaviours

Nathan: “I looked at [the software] every day, so I knew how they coded, so I get... get that, and I get this and whoever it was, it could have been a teen, but I knew what they were doing. If you look at the same packer every single bloody day. It’s like, oh, it’s *him* again! I wish he would put his name in or something, I could have said hi or something. It, it was interesting, but yeah, in the end if I had done a [behavioural detection] I would have, which I don’t think we had at that point.”

Andrew: “Yeah.”

Nathan: “And it would have killed it every single time and I wouldn’t have had to bother.”

(Interview, 21 July 2017)

As Nathan expressed, static strategies that focused on an individual analyst continually reviewing detections have been transformed through the addition of contextual strategies, and especially through detections based on behaviour. Nathan is clear that behavioural detections could “have killed it every single time,” if behaviours were tracked rather than the emphasis being placed on malware *in stasis* – that is, instead of reviewing packer algorithms

each day – one detection type would have meant he “wouldn’t have had to bother.” Hence, the MAL became interested in behaviour. Two dominant techniques for behavioural analysis and detection emerge; i) for both analysis and detection, the automated sandbox environment (Sophos called this *Sandstorm*), and ii) the behavioural detection, the HPMal. Other forms of behavioural monitoring exist, such as tracking common ‘malicious’ behaviours in the endpoint detection engine that are more generic. One example comes from a commercial desire to detect ransomware that tracks and stops the behaviour of sequentially opening, encrypting, and saving files; one that is *normally* only associated with ransomware. Although the knowledge and practice of this behaviour may have developed in the MAL, these techniques were not built within it, being part of the detection engine, but not developed by malware analysts, so I do not focus on them here.

Behavioural techniques tackle issues the MAL faced with respect to the need for greater human labour as a result of the increasing volume of malware. Whereas static signatures were applicable for a certain relationship between materiality and temporality – of lower volume and speed of production – they are not well adapted to the increasingly various software materialities, or envelopes to riff off Ash’s (2015) use of the term. Intermingling with human analysts’ knowledge and affective capacities to the abnormal with behavioural detections, and working with more-than-human behavioural monitoring in endpoint detection engines, have established new forms of ensnarement; new ways of anticipating and identifying what may be malicious in ‘real-time.’

The Automated Sandbox – ‘Sandstorm’

Sophos’ sandbox, Sandstorm⁵⁰ (an ‘automated’ VM) let software execute, enabling for both analysis and detection. Within Sandstorm, a range of tools are used to understand how malware interacts, in order to draw out anomalies and combine these into malicious traits. In tracing software’s execution and their becoming, automated sandboxes are used as a way to monitor performance according to rules written by analysts that are based on behavioural assumptions developed in the MAL. These assumptions could include, for example, that certain access to a particular part of a computer through a particular method is only *normally*

⁵⁰ This is the project name given to the sandbox, which has a different internal MAL name that cannot be revealed. For more information see Sophos (2018).

used by malware, or that ‘dropping’ certain file types and sizes may be anomalous, but not necessarily limited to malicious software. Software could be both uploaded from consumer endpoints for analysis in a cloud-based instantiation of Sandstorm (and possible detection) or used by analysts ‘locally’ (i.e. on their computer) for quick analysis, detection development, and improving and writing the rules for the sandbox.

Sandstorm was well-used. When shadowing Mason during his analysis and detection of Dridex he employed process monitoring tools and sniffers⁵¹. There were a series of automated tools that collected data in behavioural analysis; processes initiated and terminated, registries targeted, dropped files, files modified, and outbound connections to shared folders or to the internet. Not only was the process of detection contextualised, so was malware analysis. This transformation came with an increased risk and (for me, a seemingly) reduced controllability. For example, when I used Sandstorm during analysis;

“I put [the software] in [Sandstorm] locally [on my computer] and left it to run - however it came back as a clean file. So, I asked [Charlie] whether it was usual for [Sandstorm] to do this and he said that it was very unreliable - sometimes working, sometimes not... [Charlie] suggested it may have exceeded it, may have exceeded [Sandstorm’s] limit of x ⁵² seconds in the sleep [Windows API⁵³] function. The malware had exactly x seconds of sleep - and was probably the reason Sandstorm did not pick it up... though [it was detected as malicious] elsewhere and didn’t flag as being particularly malicious?”

(Research Diary, 19 May 2017)

Sandboxes often slightly deviated according to different environmental parameters, frequently exceeding an analyst’s explicit knowledge why in one case software may flag as malicious, and in another case not. There was a fragility to these contextual analyses. Observing and manipulating behaviours introduced resistant more-than-human choices and agencies into the process of analysis. Critically, this worked within similar pathological notions to those explored by Foucault’s (2016) *Collège de France 1974-1975* lectures on the “Abnormal.” In Foucault’s characterisation of the plague town, there is the persistent monitoring of key points in the city. He describes this as “a kind of pyramid of uninterrupted power... in its exercise, since surveillance had to be exercised uninterruptedly” (2016, p.45).

⁵¹ Sniffers look at ‘traffic’ between the host (a computer) and external connection, this can reveal incoming and outgoing information, websites contacted, command and control servers and so on.

⁵² The x represents a certain number of seconds that cannot be revealed due to its sensitive nature.

⁵³ The Windows Operating System has an API (Application Program Interface) which allows access for software to different versions and tools.

Yet the MAL treated malware in a far more insidious way. It was more similar to the ‘societies of control’ outlined by Deleuze (1992), where movement was tracked to tease out the anomalous. The MAL let software execute that meant malware unpacked itself and one could observe malicious agency. Yet once *something* breached the rules of abnormality written by analysts, it was deemed suspicious, and had to be stopped and prevented.

Behavioural Detection – the HPMal

Unlike Sandstorm that operated on the cloud or in the MAL *away* from endpoint detection engines, the behavioural detection (HPMal), is based on HIPS⁵⁴ *within* the engine. In the same shadowing session with Mason, whilst building a detection for Dridex⁵⁵, a close link with static signatures emerged. Mason worked through various samples, in an effort to accommodate a change from ‘32-bit’ to ‘64-bit’ malware, that uses different registers (for storing and calling variables) on a computer’s central processing unit (CPU). Mason did this through his personalised sandbox, inserting samples of his detections through this system, often failing, getting cross, commanding we go for several cups of tea, and going out for a fag. Mason looked back and forth, analysing further before moving forward with a testing of the current detection and repeating the process again. This made sure the system ‘picked’ up and flagged certain attributes of Dridex.

The behavioural HPMal detection is an advanced kin to the conventional static signature. However, it is dynamic. It monitors software execution through the detection engine. The software is allowed to unpack itself, so the MAL does not have to build systems to do this, for example. As the software executes, the detection monitors the software, that in Dridex’s case, allowed it to reach a certain point before Mason had decided that it had reached a threshold where he was convinced it was Dridex, and therefore detected. Detection is reliant on execution, checking if the software performs certain jumps or accesses certain parts of the computer’s operating system. Software structure can vary significantly but have similar behavioural attributes. Under static signatures this can cause serious issues but HPMals can reduce the number of detections that have to be written for malware families and variants

⁵⁴ HIPS (Host Intrusion Prevention System) is the behavioural detection monitoring system that is used by Sophos to identify behaviours at the ‘run-time’ of the execution of software and code.

⁵⁵ This was for the new version 4 of Dridex, which used the ‘atom-bombing’ technique to gain access. See Baz (2017) for an analysis of the Dridex implementation, and Liberman (2016) for the original technique on a blog post.

by expecting certain behaviours (often with some similar structures, to ensure that it is the malware form you wish to detect). The point of the decision by Mason in the MAL's environment therefore assumed how he expected it to execute in the 'wild.' Decisions made here on what *is* malicious *become* defined as malicious elsewhere.

The growth in 'file-less' malware that do not come as executable software⁵⁶ has challenged the ability of static signatures to analyse and detect. Malware are often not 'software' at all but can be snippets of code that embed into a registry⁵⁷, equivalent to the 'keys' of the computer. This leaves no software or code to analyse post-event, self-deleting after they have performed certain tasks. Static strategies that deconstruct software in *stasis* post-event are sometimes no longer able to deal with these new materialities. Instead behaviours, such as observing access to and processes on PowerShell (a cross-operating system tool developed by Microsoft), can monitor for certain commands to prevent malicious code being used and are but one example of the movement beyond the static signature to the use of behavioural approaches. Suspicious access, combined with other contextual behaviours, can be denied before its full execution. Changes in the materiality and growth of malware have driven a commercial imperative to move to behavioural techniques; yet these still require extensive human interaction, senses, and affectual modes of maliciousness. As these senses are transferred to behavioural detections, the more important it is for the MAL to retain the analysts' sense of abnormality. This provides a stability to identifying what is malicious or not. In turn, this permits malware samples to be detected *behaviourally* and *contextually*, reducing the need for intense labour, so that it is taken away to more-than-human computational detections, so that *better*, human-directed broader detection methods can be developed by these analysts, whose affectual build-up with malware are increasingly valued and capitalised.

Machine Learning

It is inconceivable that the huge volumes of data that the MAL collects from contextual strategies could be processed by a singular (or even group of!) human(s) without extensive computational and algorithmic intervention. Machine learning techniques have emerged to

⁵⁶ This means that it cannot self-launch but requires other executable programs in order to perform.
⁵⁷ A registry is a store of 'keys' that allow for certain permissions and actions on a computer.

extract an understanding of software features, and how they may be considered anomalous, in addition to the reputation scoring algorithm. The acquisition of Invincea by Sophos brought in its extensive machine learning expertise, having previously received significant DARPA⁵⁸ funding (Invincea, 2017). At Sophos, this came in the use of neural network algorithms that have become extensively used in security techniques that Louise Amoore has explored with relation to facial recognition, natural language processing, and algorithmic ethics (2019). However, machine learning is not new in malware analysis, with Kaspersky, the endpoint protection provider, noting that ‘n-gram’ machine learning was in use from 2001.

“When it comes to cybersecurity, machine learning is nothing new either. Algorithms of this class [n-grams] were first implemented 10-12 years ago. At that time the amount of new malware was doubling every two years, and once it became clear that simple automation for virus analysts was not enough, a qualitative leap forward was required.”

(Malanov, 2016)

Neither are neural networks themselves new; they are mentioned in early malware analysis research as a potential way to detect malware (see Guinier, 1991, for an example of this in computer science) – and were part of early cybernetic thought and development. However, the recent development and impact of neural networks is crucial to demonstrate how contextual strategies aim to work with computational cognition through their ability to process and recognise anomalies (and as close as possible, abnormalities). In publicly available information from Invincea data scientists, Saxe and Berlin (2015) express how convolutional⁵⁹ neural networks have become operationalised in malware analysis and detection. These algorithmic forms are by no means the only algorithms used in the MAL (reputation scoring for example). But their significance comes from how human senses of abnormality are intertwined with machine learning algorithms, that in self-learning forms, select their own attributes for detection, permitting the development of their own spatialities, own dimensions, and making *choices* based on teasing out abnormality from learning data through the formation of the anomalous.

Figure 9 provides an overview of how Saxe and Berlin articulate how an older, non-self-learning convolutional neural network operated, where they pre-define the features to be

⁵⁸ This is the US Defense Advanced Research Projects Agency. For a history of DARPA and its influence see Jacobsen (2015).

⁵⁹ Convolutional neural network algorithms are a form of ‘deep-learning’ algorithm that are primarily used in visual analysis.

extracted. In (1), ‘feature extraction,’ the algorithm takes information such as PE headers (an executable’s ‘contents page’), strings, and other metadata that are also widely used in other static and behavioural strategies. Other feature extraction however take a radically different approach: such as byte distribution and string histograms (thereby being able to render the PE file as an image for the convolutional neural network to operate) – which are beyond human *cognitive capacity* but not for computation; in that the processing of big data can be done beyond the explicit capacity of a human to draw links and connections. Later work by the authors outlined how automated feature extraction can expose a whole raft of algorithmic agencies through ‘self-learning’ techniques that I cannot do justice to here (Saxe and Berlin, 2017; see Amoores, 2019 for examples of how self-learning neural network algorithms become operable)⁶⁰.

Attributes from the neural network’s feature extractions are added to layers of ‘deep learning’ (2). Figure 10 provides a graphical illustration of a neural network diagram to more clearly demonstrate how machine learning extends and amplifies computational cognition through higher levels of sign-processing. There are varying techniques of how to write neural networks, such as how to select an *appropriate* ‘weight’ (each line representing a weight, w) to ensure that neurons activate and process data. The neural network processes features, and by what it *learnt* through the training data provided to it, *recognises* what is malicious. As in (3) of Saxe and Berlin’s table (Figure 9), the *Bayesian estimation of P(malware)*, is only a

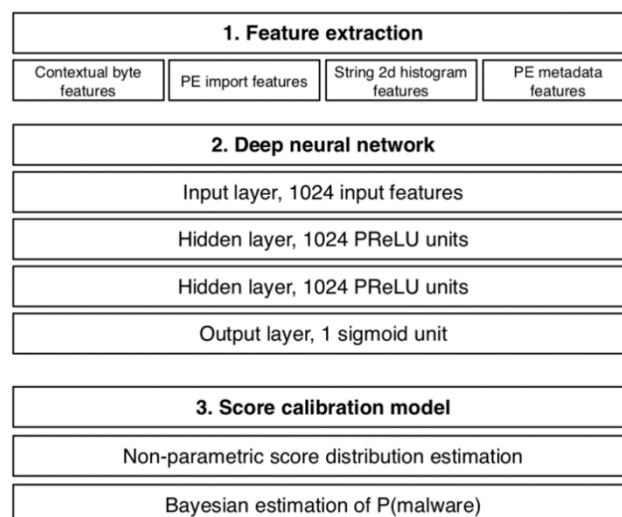


Figure 9: Overview provided by Saxe and Berlin (2015) in their construction of a form of neural network for malicious software detection, based on training data from Virus Total.

⁶⁰ Since the start of the writing this thesis, there has been a publication of a new book, Malware Data Science (Saxe and Sanders, 2018) by individuals from Sophos.




Figure 10: A graphic representation of a 'neural network' algorithm with different layers, neurons as circles, and the weight by each line, connecting each neuron in another layer. This does not mean that every neuron becomes 'activated' but shows possible movements. Author's own image.

probability. Each neuron, if activated, sends a *weight* to all neurons in the next layer of the neural network, according to what it has read and interpreted, on whether a particular feature is likely to be malicious. These features may have predefined that x is a malicious attribute, but in self-learning even attribute x is chosen by the algorithm. This becomes a non-linear process, where the neurons in the network bend and twist what should be *counted* in multiple dimensions as the neurons' weights cross one another, extending beyond typical human comprehensibility. These algorithmic structures are imbued with data science's neuroscientific analogies as much as in malware analysis, developing on mind and body distinctions as in early work by von Neumann (2012) in the 1940s and 1950s. Malware, once again, then become reduced to information, so that they can be understood through abstraction, without an environment but contextualised with other malware. Thus, in discussing malware ecologies, it is not only enough to think about malware as political actors, but how the process of detecting malware itself are now part of cognitive ecologies which emerge between neural networks, humans, and other computational technologies in the MAL.

The process of extracting features based on software structure is different to what I identified in both static and other contextual strategies. The variability in software bit-distribution now becomes part of a more-than-human focus. This distribution of bits *to mean something* is beyond human comprehension. Processing and working with the features in the

neural network are qualitatively different to an analysts', and my own, encounters with the strategies I have detailed thus far. This is a more-than-human cognitive process that is political beyond the capacity of explicit human intervention on why something is happening. Training neural networks with 'good' and 'bad' bins (data sets) is required to delineate between malicious and clean – even if there is the persistent worry of 'dirty data'⁶¹. The use of Virus Total data, a commercial repository of malicious samples, to train the neural network folds analyst knowledge, practices, and affects from across various endpoint detection businesses and their MALs. Hence, it was not only the Sophos MAL who came to matter in the production of the neural network in these 'nextgen' detection techniques. But they are formed through other MAL environments, customer telemetry, multiple human decisions, variable malicious performances, and how and what is deemed to be a priority to detect.

The capacity for neural networks to operate in new dimensions and layers came after a Sophos-wide information session;

“Each [neural net] layer focuses on a different form of feature extraction that will eventually 'learn' whether that file is malicious or not. They called it separating the good from the bad. The feature space is exceptionally interesting as the guy doing the presentation said the data “lives” in the space and can therefore be manipulated in ways that cannot be done in 2D or 3D.” (Research Diary, 2 March 2017)

Neural networks disrupt the differentiation of space – through a topological, non-linear piecing together of similarity and difference – to comprehend the anomalous. Working in multiple dimensions (through the interweaving of data and information between neurons, weights, and layers), neural networks do not simply sift through a simple linear line “separating the good from the bad,” but are entangled, twisting and turning according to the various features extracted in their complex contextual relationships. This is why cyberspace is important – as it processes signs in a different register to humans. In ways similar to discussions on 'ethical' artificial intelligence, this complex movement of 'features' makes it difficult to identify a point where an algorithm 'decides.' There is a movement to a condition of uncertainty that the algorithm makes easy to disguise, in what I argue is a new form of

⁶¹ This is where data is 'contaminated' through the incorrect allocation of 'clean' software as malicious. However, this is becoming less of a worry as the neural networks keep adding learning data. Indeed, in a MAL-wide meeting discussing the acquisition, the reason that the neural network was self-learning is because it was seen to be better at allocating what is a better malicious attribute than the analysts (Research Diary, 2 March 2017).

computational biopolitics; one less interested in an (ab)normal distinction but similarity and dissimilarity to discern the anomalous. As Amoore and Raley (2017, p.5) say, “to secure with algorithms, then, is to reorient the embodied relation to uncertainty, so that human and non-human beings are constantly attuned to novel events and features in their data environment, becoming perennially creatively alert.” When I was in the MAL, learning, extracting and looking at scores produced by the neural network, at reputational flags, and developing a sensitivity to the probability of maliciousness, all came together to create a disposition towards always anticipating malware.

As a neural network learns based on previous iterations of what is malicious, one may assume it is identifying maliciousness. This is not the case. They ‘learn’ or cognise through anomaly. I refrain from calling this malicious as a computer does not recognise the complex ecologies in which maliciousness *becomes* in our societies, infused as they are with the histories and trajectories of pathology that intertwine prior and through our current epoch. A folding of multiple human-sensed analyses from a variety of MALs are undeniably infused with human senses of abnormality. Neural networks are trained with data that are constructed through the human analyst body, its knowledges, its intensities, and how the human multiple reach a decision of maliciousness. However, the ways in which the neural network folds in multiple dimensions, learns in an alternative cognitive register to humans (such as through bit distribution), ensures that any inference that neural networks reflect analyst intentionality are wrong. As the learning data is only tagged as malicious, it only searches for differences or deviations away from a multidimensional norm (Aradau and Blanke, 2016, 2018). Therefore, as the norm is constructed *through developing senses of abnormality and through data as anomaly*, there is an intriguing relationship in neural networks that work broadly parallel, but on a different register, to human-directed analysis. By this, neural networks are able to construct new forms of anomaly, that frequently *exceed* a MAL’s sense of abnormality – exposing the ‘unknown unknown.’

Neural networks, due to their cognitive capacities, are then not always aligned with societal expectations of maliciousness, pathology. They can exceed analyst expectation, seem inaccurate according to human senses of maliciousness, producing what is known as a high false positive rate (the incorrect identification of software as malware). Machine learning and algorithms are generative of confusion, extending and moving abnormality through the

contours of maliciousness unknown. They are redefining pathology as they are re-enrolled with the analyst in the environments of the MAL. Their contextuality make them *risky* to deploy without reference to other contextual (human-directed) strategies. This means software can be sent to Sandstorm, the automated sandbox, for a behavioural detection or in ‘complex’ cases sent to analysts for further probing. Though contextual strategies allow for speed, scale, and identification of malware as yet unknown, static strategies are still more trusted due to their connection to the affectual capacities of the analyst – otherwise known as ‘experience.’ Percentages stared out of my contextual information interface when performing reviews and analyses of software samples, meaning everything came together on the three screens of my MAL set-up. Due to the complex cognitive ecologies and relations to abnormality, they were deemed risky and were too ‘wild’ to be deployed without reference to the human-directed static and behavioural detections, that are in turn imbued with many more-than-human agencies. What a mess of agency!

Contextualisation has transformed the pathologisation of malware analysis to one that is no longer one of stasis, of borderlines, but of borderlands, turning to Hinchliffe *et al*'s (2013) distinction between breaching walls to ‘tipping points’ in biosecurity paradigms. The MAL is aware that maliciousness emerges in *noisy* borderlands, coming from the fuzziness between clean and malicious attributes. There is no simple dividing line which demarcates ‘good’ from ‘bad’ software. Instead there is a score that renders the likelihood, or the reputation assessment, that software are malicious, derived from a proxy value or abnormality when compared to ever-changing, contextualised norms. This concern with borderlands informs distinctions between the normal, abnormal, and anomalous that feed back to the contemporary static ‘anti-virus’ signature itself, in how it is infused with complex, embedded, convoluted knowledges and practices of humans unknown, of more-than-humans unknown, as demonstrated through the ‘matching’ gene. By this, the reputation of software informed my sensibility to detect. As I searched for which attributes to select, what attributes made the software malicious, the distinction to cleave between malicious and clean samples was exceptionally difficult. Contemporary and conventional analyses and detections are by no means hierarchical but express alternative space-times of production and the complexity of the relationships between forms of analysis and detection. The coalescing of agency, environments, and distinctions between what are normal and abnormal are all about a

slippery transition to using context as borderlands for permitted performances, rather than borderlines that anti-virus presented.

Data

Data is at the core of contemporary MALs. Data is something that gives, must be performed into action and is co-produced between human and technological affects (what to search for, the *accuracy* of the data, where does the data come from?). Drawing on the etymological root of data from both Latin and Ancient Greek, *to give* (OED Online, 2018), one can play with this give of the world. This runs similar to the use of ‘capta’ by Kitchin and Dodge in *Code/Space* (2011, p.5) where they declare that “capta is what is selectively captured through measurement.” Capta explicitly refers to the *capture* of data from environments but does little to comprehend the complex *performance of* digital data and more-than-human political choices, which are always negotiated. Thinking beyond capta respects that computation does not simply produce data to be gathered, to be swept up by humans, but is an active participant in its negotiation.

Data *feeds* throughout contextual systems where three broad categories of data are available to the MAL: external data feeds and metadata, customer telemetry, and internal data that now form ‘big data’ that have been a recent addition to the endpoint detection industry. In this section, I outline the importance of each element of data production, what it collects, and how it becomes enrolled within the practices of developing analysis and detection techniques, which are then folded in both static and contextual static strategies. Whereas in the previous section, I looked at the two strategies, I now turn my attention to how data is what permits contextualities that now permeate the lab itself. The value of data was underscored by senior analyst, Joe:

“We base our decisions on the data we have available. As I said, 10 years ago we had no data. There was SXL⁶². We had no visibility, we had nothing.”

(Interview 6, Senior Malware Analyst)

⁶² SXL (Sophos Extensible List) is a method for the retrieval of data from a remote computer and for also performing a ‘look-up’ to check for detection information. For more information refer to Sophos (2017).

I was intrigued by this declaration of “no visibility.” It signals that data has somehow enlightened or brought *visibility*, whereas before big data, there was a darkness, a lack of ability to see. This reflects crucial elements identified in other analyses of contemporary security paradigms which share this concern with revealing information that may be ‘hidden’ (Amoore, 2011; Thornton, 2015). Before, Joe told me, endpoint detection and analysis were tightly threaded to the individual malware analyst’s affective senses of risk, to the *abnormalities* of software. Lack of data limited the ability to test detections in the ‘wild’ as there was little to no feedback on their success on endpoint detection engines; to construct knowledge of malware forms for analysis as there was little sharing of samples; and, that ‘good’ detections were not ‘scientifically’-grounded due to the lack of feedback on their success or comparison to other businesses.

External Feeds and Metadata

Within the MAL a proliferation of external feeds sent metadata. These (metadata) feeds are transferred and shared in the endpoint security industry and other cybersecurity organisations, frequently as best practice to ‘improve’ security – more data means a greater *certainty* (Ben-Asher and Gonzalez, 2015; Franssen, Smulders and Kerkdijk, 2015; Winkler and Gomes, 2017). One dominant, commercial feed, *Virus Total*, that was acquired by Google (now Alphabet) in 2012, is widely accepted as the staple repository of malicious software samples. Virus Total provides a dedicated API⁶³ (Application Processing Interface) to its repository, where these samples are in turn supplied by its subscribers (either by MALs or through its online deposit). This assists sample sharing among endpoint protection businesses, ensuring swift distribution. The free-to-use web search interface for individual software samples performs a check to see if a software form is detected, and by which endpoint detection engines. In other words, if a new malware form is found, this database is used to see if other engines have detections. This develops a collective indication of maliciousness between MALs and is used to construct an understanding of the likelihood, the probability, that something is malicious. The greater the number of detections for a software sample from other vendors, the more likely it was considered to be malicious; even without reviewing it, the analyst performing the search now knows many others deem it as

⁶³ These are used as an interface to both send and accept information and interact with Virus Total’s services.

Figure 11: An example of the Virus Total interface when searching using a file 'hash' (a unique cryptographic signature). In this case, I used a SHA-256 hash of a WannaCry sample. Screenshot taken 21/03/2018. Author's own image.

Search or scan a URL, IP address, domain, or file hash

62 engines detected this file

SHA-256: 01b628fa60560c0cb4a332818cb380a65d0616d19976c084e0c3eaa433288b88

File name: tasksche.exe

File size: 3.35 MB

Last analysis: 2017-12-28 06:12:11 UTC

Community score: -83

62 / 68

| Detection | Details | Relations | Community |
|---------------|-------------------------------------|--------------------|-----------------------------------|
| Ad-Aware | Trojan.Ransom.WannaCryptor.A | AegisLab | Troj.Ransom.W32lc |
| AhnLab-V3 | Trojan/Win32.WannaCryptor.R200571 | ALYac | Trojan.Ransom.WannaCryptor |
| Antiy-AVL | Trojan[Ransom]/Win32.Scatter | Arcabit | Trojan.Ransom.WannaCryptor.A |
| Avast | Win32:WanaCry-A [Trj] | AVG | Win32:WanaCry-A [Trj] |
| Avira | TR/Ransom.Gen | AVware | Win32:Malware!Drop |
| Baidu | Win32.Trojan.WannaCryc | BitDefender | Trojan.Ransom.WannaCryptor.A |
| CAT-QuickHeal | Trojan.Mauvaise.SL1 | ClamAV | Win.Ransomware.WannaCry-6313787-0 |
| Comodo | TrojWare.Win32.Ransom.WannaCrypt... | CrowdStrike Falcon | malicious_confidence_100% (W) |

such. It is rare to find malware samples on Virus Total that are detected by all endpoint protection, but if there were *certain, reputable* endpoint engines⁶⁴ reporting similar detections, this was a good indication.

Big data analytics have utilised feeds in two ways; i) as learning data for machine learning to produce automated detections as I detailed in the previous section, and ii) for malware analysts to tailor their detections through contextual information. In the MAL, this was collated into a database where dynamic searches were possible according to different variables selected by the analyst, such as the number of software samples identified by another endpoint protection detection engine. This enabled an examination of software samples that had not been detected by the MAL, assisting in the development of priorities for further human attention, to improve and refine detections. For example, at Sophos I was given detections to improve, prioritised according to information derived from feeds. One of these was a piece of software that was designed to install something – an installer⁶⁵ – which was used almost exclusively for malware. As I wrote in my diary “Terry had put in a ticket [used to allocate tasks in the Sophos team] as we were missing about 11,000 samples”

⁶⁴ Over time, I developed a sense that some detection engines were better than others at appropriate allocations as I used these to construct my own detections and realised some were incorrectly classed as one form or were not detected at all.

⁶⁵ An installer is a program that is used to load other programs onto a computer – such as when one downloads a new piece of software there is a piece of software which normally takes one through the installation process.

(Research Diary, 23 May 2017) in contrast to a combination of other endpoint detection engines (without knowing *what* exactly they had used to detect on). I did not know if the missing samples were likely malicious, but the fact other vendors had decided to produce detections for them lead me to assume they were malicious. With data from the external feeds, it became clear that there was a slightly defective filtering gene near the start of the Sophos detection which meant samples were no longer being flagged for continued analysis and detection. Along with Elliott, I was able to develop an alternative filter that then enabled these samples to be included, and thus detected. So, when Joe said he had “no visibility,” one component was a lack of the ability to compare the effectiveness of Sophos’ detections compared to other endpoint protection businesses. Without the feeds from other MALs, through Virus Total, it would have been impossible to identify the ‘missed’ samples and thereby to see the need to improve the detection.

Customer Telemetry

Yet, these external feeds were not the only source of information. Data was produced by customer endpoint detection engines through Sophos’ SXL that enabled for a contextualisation of the computing environments in which their engines work. In other words, if there was a detection engine on a computer, it would (with authorisation) feed data to the MAL. The MAL then used this data to develop normalised senses of environments in the ‘wild’ through data on reputation which, in turn, helped to define borderlands of maliciousness. This system reported back to the MAL a complex variety of data produced by the engine analysing software forms such as the age and provenance of a piece of software, as well as on detection rates. This was supported by a distributed cloud⁶⁶ infrastructure (or more accurately, a distributed server network) where ‘lookups’ retrieve information about the file being analysed on customer endpoint detection engines. For example, if I received a file, “thesis.docx,” then the endpoint detection engine would ‘lookup’ the attributes of the file, such as its age (both timestamp⁶⁷ and its record in the MAL’s system) and provenance, and this data would provide information on what the detection engine should do on my computer. The engine comes into contact with cloud infrastructures to determine

⁶⁶ See Amoores (2015) for an overview of the complex geographies of cloud computing; about its locales, politics and productions of ‘sovereignty.’

⁶⁷ A timestamp is the time in which a piece of software is compiled (i.e. transferred from code to a distributable format). This is often a (roughly) true indication of time of compilation but can be doctored and so cannot be trusted.

maliciousness according to a *corpus* of data given by the laboratory. A dual system was at work, providing additional contextual information about the attributes of software, whilst also constructing senses of the norm. By collecting contextual information from customers, Sophos increased its data sample (or background noise) against which it could look for signals and also be able to detect more malware in the process.

Laboratory-Produced Data

Data was also produced within the MAL through both analysis (through Sandstorm, for instance) and detection (in identifying and categorising software as malicious). Sandstorm's analyses and malware categorisations were not the only form of data in the MAL; they also came in reputational data, from conversations in meetings and over the *Jabber* communication tool. The MAL's data folded upon itself; and this is where the data from this thesis comes together, through how the laboratory twists, mutates, and renders visible data to affective states that then become actionable in the process of analysis and writing detections.

Tools written by analysts were in turn based on previous analysts' tools. That is, the senses of what were malicious were constructed by the tools used to highlight anomalies, that in turn, constructed over time what I thought were likely to be abnormal. This could have been a simple tool to look at what a Windows program 'calls' – e.g. information such as the computer's time on the operating system – or whether the timestamp on a piece of software looked 'abnormal.' Where senses of anomaly and abnormality emerge were not clear, convoluted and pushed to the limit the question of who curated malware, to whom should it be attributed, and in this, the production of security prior to its dissemination beyond the MAL. Data must be considered in its broadest sense; as emerging through the cumulative build-up of knowledge of the malware analyst, of the dimensions produced in a neural network algorithm, as well as the colours selected for my information screens. This allowed for a data crafted in a more-than-human register; in interactions within and beyond the spaces of the MAL, not necessarily being directed by humans, but in a confusing spiral of alternative contributions – where computational cognition, humans, and other technologies mix.

Crafting Signals

Yet data alone was not a simple transmission of knowledge, or affects, of an environment. But what data was starting to do, I think, is to show how important an ecological role for understanding malware is – even if it is warped through pathological logics. That is, there is a recognition that there is a need to know more context around environments, on the particularities of malware materiality and so on, but only to the point of detection. These data can then be considered as *noise*. Noise is that which is the condition of politics, we live in a noisy world, with its redundancies and space for choice. This is what enables change, enabling coagulants to form, and densities that generate visions of stability. These senses of stability could broadly be considered to be signals. I argue that the work, or the *curation*, of the MAL is to discern from this noise the signals of maliciousness. This is constructed with the data that are produced within and from outside the spaces of the MAL. (Human) curation is re-enrolled into technologies that perform analysis and detection of malware. Signals are curated through an extensive working of the malware analyst and technologies to identify what to observe, what attributes to zone in on, and develop more-than-human senses of what is malicious from a variety of anomalous activities.

Discerning anomalies came through this signal production. For example, if I was to look at 10 000 Windows executable (PE) files, their commonalities could be drawn together. Files may be obfuscated, exhibit similar-sized sections, and have no embedded files. Yet the noisiness of computation meant I could search for similar files and potentially find new attributes that fitted ‘better.’ The majority of files may exhibit structures and behaviours that are similar, enabling the drawing of anomalous signals for this group, that in turn helped to distinguish them as abnormalities. Though signals gravitate to a normal, they never quite achieve this. They are always fuzzy, complex, and forever escaping true verification. Noises can coalesce into a tune, a signal, where discordant, noisy anomalies can be brought out, defined, and analysed. There are two norms at work here then – one on the mathematical, logical difference between different attributes and then a norm based on societal expectations and sensibilities to the malicious – sometimes these two norms align, clearly so with ransomware, where structures lead to a clear alignment between the two – anomaly and abnormality. But this was frequently not the case, and this was the work of the MAL that I explore more closely in the next chapter.

In taking inspiration from Karen Barad's discussion of quantum physics, I understand signals of maliciousness relating to *both* mathematic and computational logics, but also to the situational and social recognition of abnormality;

“devices don't disclose preexisting values but rather that it is the specific material configuration that gives definition to the notion of the property in question, enacts a cut between the “object” and the “measuring instrument,” and produces determinate values for the corresponding measured quantity, leaving the complementary quantities indeterminate. Which is not to say that human observers determine the results, the data doesn't come out however we want, but rather the specific nature of the material arrangement of the apparatus is responsible for the specifics of the enactment of the cut” (Barad, 2007, p.264).


Demarcating the signal was not one that was purely human but criss-crossed with more-than-human agencies that enable for specific iterations, or cuts, of what the measured quantity is. The specific ecologies were an essential component of how I arrived at a point of being able to determine what abnormality was. So, the MAL was not only conditioned by the affectual capacities of environments but also the performances of human affectual decisions and computational choices. The tools and data that the MAL drew upon have material enactments that are so complex, folding, that the malware analyst must attune to these in a register beyond the conscious cognitive, to one that is affective, felt. This is what allowed Joe to say that though analysis is based on “gut feel,” decisions are now also informed by big data, where before there was “no visibility, we had nothing” (Interview 6, Senior Malware Analyst).

Blending Logics

Contextual strategies are not detached from static strategies; they are linked by a contextualising patho-logic. They reinforce one another in the contemporary MAL through big data, machine learning, reputation scoring, and behavioural techniques. I have outlined that contextual strategies are tied together through the treatment of software through monitoring environments, or its likelihood of being malicious through anomaly. These speed-up malware detection, but also increase risk due to this reliance on anomaly. Static strategies still provide the foundation for a multitude of other technologies and techniques but are themselves now informed by contextuality through data. The gene, containment of malware in the Sandstorm sandbox and in VMs, to using data in similar ways to the use of *bioinformation* (Parry and Greenhough, 2018), are all part of this broad pathology to malware.

This is not to say that it follows a particular pathology, but is complex, following different strands of its lineages; which has allowed for developments in static to contextual strategies. The latter has been accelerated by both the production, or performance, of (big) data as a *lubricant* to increase the speed, scale and *potential* to identify malware for commercial gain. This requires a movement from human-centred analysis to one that openly intertwines computational agencies with human senses of maliciousness that can lead to issues of the *recognition of maliciousness*.

A restrictive perspective of the MAL may conclude that static strategies are now ‘dead’ or old-fashioned and that they should move forward to contextuality as thus:



This is the sort of image some in the endpoint detection industry would have us believe – in an effort to sell ‘next generation’ products. That is, static detections are too slow, cannot deal with volume, and cannot identify ‘unknown’ new malware like in anti-virus. I disagree. In discussions with those in the MAL, at the *Virus Bulletin* conference, and when I returned to the MAL to meet Daniel to discuss these ideas, there was agreement on the distinctions I made between the static and contextual, and how these intertwine, even if Daniel did say these were rather “academic” (Research Diary, 15 November 2018). Perhaps they are, but I think they are significant – they show how malware analysis is itself developing a new politics of cybersecurity – one that is based on using expansive computational cognition such as through the use of convolutional neural network algorithms.

The three screens at my desk, the people around me, the server rooms, and the probable hundreds of cups of tea all actively constructed me, the becoming-analyst. Without these different factors, the MAL cannot easily be discerned or grasped beyond actually experiencing it. I could not become without the complex mixture of different signals, senses of anomalies and abnormalities. Combining reputation, behaviours, and machine learning into the contextual is therefore a way to have a balance between the ‘generic,’ through normalising strategies, the amount of software that is required to be seen by an analyst through ‘specific’ detections, whilst still ensuring that malware are appropriately identified as malicious. Data environments have become central to contemporary practices of malware

analysis and detection. In this sense, the endpoint detection industry is moving from borderlines to borderlands, from borderwalls to tipping points (Hinchliffe et al., 2013). Static detections are still a significant (active) human labour at work in the MAL. But more-than-human technologies, computational agencies, environments, and big data come together in differing temporalities in the contemporary MAL. Ecology is however one that recognises that labour in the MAL is not solely human, and the point of decision of maliciousness is also increasingly co-produced through tangled alignments of choice between mathematical, logical, and societal norms.

Machine learning is not a step-change in this process, but a new method of the contextual and not separate from it. As Saxe and Berlin (2015, p.18) state in their work, due to certain attributes, “static analysis [are] very amenable for machine learning approaches, which tends to perform better as data size increases.” This challenges linear *progress*, with a clear connection between static strategies and contextual developments. They are interconnected. Hence contextual strategies exist to change the spatio-temporalities of malware analysis and detection to speed up and increase the volumes of software to be pathologically interrogated. This is what introduces issues such as around the ‘false positive’ – as anomalies, as Canguilhem tells us, do not constitute an abnormality – so when a neural network ‘sees’ malware it is only working from statistical anomalies in order to find societal abnormality.

Eco-logical practice asks *how* human action is being changed, how a politics is being constructed that is more-than-human. I do not intend to leave pathological practices of malware analysis behind; they (broadly) work in the task of detecting software as malicious – and I wish that to continue. But, when computation expresses a politics beyond and through humans, then how we understand, and curate malware becomes essential. The MAL is a site of coalescing data where it is worked, reworked, and rendered as information to the malware analyst. Big data is a critical component of the laboratory, and surrounds, complicates and challenges the malware analyst to work with these others as well as the *nitty gritty* daily engagements I had with malware. When I became involved with these complicated alignments of information, it was crucial in shifting my attunement to software I was analysing, moving me in ways that are difficult to ascertain. Data is as crucial to comprehending the MAL as well as the agencies of the collective computational technologies that form the laboratory, of which malware are a part. This chapter then has delved into two

strategies – the static and the contextual – that structure pathologies of malware analysis and detection. These are based on various resonances with environments which are increasingly infused with data that bind together both strategies in the delineation of abnormality increasingly through anomaly. Yet these require an active curation to ensure that anomaly and abnormality align, and this is what I consider in the next chapter.

Chapter Six || Curating Maliciousness

The malware analysis laboratory (MAL) is in an incessant process of *curating* malware. Curation is used in its double meaning, that between both *poison* and *cure* in the writings of the Greek, Platonic and latterly, Derridean (1981; Stiegler, 2011) sense – between both the poison of (more-than-human) exclusions, and the potential cure of the malicious, and also in its more pragmatic sense of curation in a museum or art gallery. My claim is that maliciousness *becomes with* a variety of technologies, expectations, and discursive lineages in ways that fold and entwine static and contextual strategies – that I detailed in the previous chapter – in moments of more-than-human choice and human decision. This chapter examines the threshold of these logics to expose the difficulties in pulling together and maintaining the integrity of the malware form in the context of wider ecologies of agency. Working through the productivity of noise, the worry of the ‘false positive’, and the Potentially Unwanted Application (PUA), I explore how humans and computation come together in analysis and detection. The MAL labours and curates to maintain semblances of stability to whatever is identified as malicious; but this is always being (re)built to keep the fragile affects, technologies, and knowledges together.

Central to my discussion is the claim that a process of curation is *the* crucial labour in the MAL. Curation was an emergent communication between me, as becoming-analyst, and malicious software. This is part of the condensation of the experimentation of malware analysis into *factory-like*, industrial detections that the MAL produces to then distribute to customer endpoint detection engines and elsewhere. The laboratory is not simply contained in its knowledge, as we have seen in the incorporation of external feeds and telemetry feeds in the previous chapter. The use of Twitter, feeds, news stories, and communication with other MALs assists in establishing bounds of maliciousness and developing co-curated knowledges. This practice is full of fragility, has potential to go wrong, and often did. In the MAL’s pathologies, I argue that there is an experimental curation, adapting to the particular conditions of software forms and their environments, that is twinned with an industrious production of standardised visualisations, names, categories and families in order to render sensible the noise of software to distribute detections as a commercial product.

The commercial drive of the MAL meant that its core activity was to categorise, make knowable, and condense possibility to prevent malicious activity on customers' computers. Malware must be detected. In pursuing this goal, these curations produced dominant human authors and environments as inherently, and thus potentially, *knowable*. This limits malware to pathology, emergent in the practices of the MAL; that goes on to affect government response, international relations, and even media studies (in particular with reference to Parikka's *Digital Contagions* (Whitson, 2017; Parikka, 2016) that draws on *virological* approaches to understand society based on malware). I believe I can trace a different path, one that is ecological in its orientation. In order to do so, in this chapter, I draw heavily on the insights that emerged from the process of how, during my (auto)ethnography, I *became an analyst* in the MAL. I first detail how malware became embodied through my experience, then second, turn to the false positive as the active negotiation between anomaly and abnormality. Third, I turn to the resolutions that the MAL turns to such as the 'Potentially Unwanted Application' that is used to transfer the moment of decision for maliciousness to the user. In the second half of the chapter, I then consider how the MAL made sense of the performativity of malware and condenses it in visualisations and other outputs for dissemination. In concluding, I look at what moving beyond pathology looks like in an ecological turn before I move to narrating new stories of malware in chapter seven.

Embodying Malware

"I am quickly, now, starting to notice structures. It's a very strange thing, noticing these structures... they have an effect on how I perceive the malware and how I visualise it (in a rather topological way, that is probably due to the dominance of this representation⁶⁸)... Therefore, I can scan and get a good sense of structure reasonably quickly. I'm sure this will only increase in speed in the future."

(Research Diary, 6 March 2017)

In this section, unlike the introduction of SophosLabs as in the previous chapter, I trace how my becoming was not simply one of knowledge, but of an embodiment in order to appropriately curate abnormality, and work with anomalies in order to render them abnormal (or not) through this embodiment. Yet this affective connection did require me to

⁶⁸ It is important here to distinguish between what is meant by topology. Within the computing sciences, and what I refer to here, is a 'flat' diagram showing network interconnection. This is not the same as for those in the social sciences, and in particular geography, which refer to the complex, myriad focuses on energies, forces, and processes in a process of becoming.

read dense pages on Intel's computer chips, learn assembly language dominant in static analysis strategies, learn and write in Sophos' proprietary *Virus Description Language* (VDL), assess who I could go to for help, and shadowing other analysts. This led to incessant failure, frustration, thoughts of quitting, and thinking I was not 'up to the job.' The failure I experienced was crucial to the curation of malicious software. This meant recognising, and then being alert to, your bodily senses so that the structures presented through technologies such as the IDA disassembler (see from page 110 for more information), generated an awareness of the value of different contextual parameters, allowing a sensing of whether these were reliable indicators of abnormality. As I was coming to the end of my intensive training in March 2017, things started to gel. I was becoming aware of how the common Windows executable format, the PE file, was structured; what held it together as a piece of software, and what the common API calls to the operating system looked like. I was beginning "a very strange thing, noticing these structures." This was essential. These cumulative encounters with (malicious) software, other analysts, and the technologies in the space of the MAL generated a bodily engagement, feelings based on glancing, a heightened sense of anticipation according to colours and structures, in order to be successful at making a 'strong' detection. It was not only a conscious learning, but an affectual build-up of knowledges, engagements and interactions with the anomalous and the abnormal. As I reflected:

"I think this is more to do with now feeling malicious software as they come onto my screen and as I read the [contextual] information... I get a good sense of what is going on. If I have PE [Windows Portable Executable] files, then a quick look over the imports gives a sense of what is likely to be happening fairly quickly, which a more thorough investigation often agrees with. *That hunch has come to me without me even realising - through the repetitive observations of the static code.*"

(Research Diary, 18 April 2017)

As Adey (2014) writes in similar ways on security atmospheres, it is worth considering how analysts are part of the affectual intensities of the MAL: the various knowledges, practices, eyes, computer mice, silence, and meetings. The static and contextual strategies (chapter 5) that are underpinned by pathological discourses and practices are mediated through this collective. Yet, its start and end are not clear; it is convoluted and difficult to disentangle. From my first day in the MAL, looking at the command line, to working with other analysts to appropriately define the contours of a particular malware family, enable through

contextual strategies, an attunement that escapes linear explanation. I cannot fully articulate how the more-than-human computational technologies, textbooks, training materials, and working with a variety of Windows executable PE files assisted in my becoming-analyst. Yet they are crucial components of how I came to know, to feel, and how *good* I was at curating in a more-than-human register and subsequently my articulations of ecology. As Latour and others have noted over several decades, scientific knowledge is always in part affective (Myers, 2015), it's about gut feeling as much as formal knowledge.

I studied many of the computing processes that are exploited by malware. This required a constant appraisal of what the 'legitimate' program was. As I primarily only had access to malware and little access to legitimate, *clean* software, I had to learn this through negation. My awareness of these structures were constantly refined, and I learned the *legitimate* use of a process, primarily drawing on the Google search engine, and particularly Microsoft's resource for developers, MSDN⁶⁹, to slowly accumulate a sense of the anomalous. I was building upon previous analyst knowledges to come to the point of recognising that something was *odd*. A development of skill and knowledge was required, where my eyes scanned across screens and my body alerted me to something odd: it forced me to take a second look. I felt a slight increase in heartrate. My eyes zoned in, I read; I (re)cognised. This strikes similar to Mol's (2000) analysis of pathologists concerned with atherosclerosis of the leg vessels. They concerned themselves with structure to identify the thickness of cell walls, the calcification, and induce that a patient had atherosclerosis. As Mol says, they 'smelt' the metabolic disturbances of the body. In the MAL, there was an embodied sense of maliciousness that warped anomalies and brought them together in the site of the (human) body to determine maliciousness – albeit with a different materiality to Mol's laboratory. This was not a smell, but visual stimuli, that were pre-conscious, affective, my eyes zooming to sections of code before I could even reckon why. Embodiment defined the analyst; being good at the 'job' was to be affectively enmeshed with the technologies and knowledges of the MAL.

The malware analyst, similar to the pathologist, takes the splintered or *disassembled* software form (in static strategies), observes behaviour (contextual strategies) and puts it under their

⁶⁹ Microsoft Developer Network, see <https://web.archive.org/web/20180118015427/https://msdn.microsoft.com/en-us/dn308572.aspx>

metaphorical microscope. *Dyes* and colours are added to tools such as the IDA disassembler, to highlight certain structures. The analyst must be able to perform the functions of recognising structures, developing (with) technologies, for a purpose: detection. This is seen as a highly demanding task which takes many years to achieve. During a tea-break, chatter turned to my training, and how difficult I found the process. As Nathan informed me, in the past, a six-month training programme for new analysts was required (5 July 2017, Research Diary). Many individuals did not complete this process, with this high bar seen as an essential step to prove oneself capable. This requirement has since been dropped as ‘deep’ analysis – primarily conducting static analysis – is now less frequent, with a new preference for on-the-job training to makes analysts more productive quicker. Yet, there is still great time and dedication required for static analysis even if there is no longer an initial onerous hurdle before one conducts simpler analyses and writes easier detections.

Significant time, over the course of several months, I spent alone reading, playing with the training materials, all in an effort to hone my curative craft. This was supported with extensive interaction with analysts. This included shadowing in most of the different MAL ‘teams’ that specialised in different aspects of malware analysis; such as my generic detection (static) team, behavioural, web exploits, threat analysis, and Android. One shadowing session with Sam included a new technique to detect malware written in *.NET*⁷⁰, where I became the second person to ever utilise this in the MAL. These frequent conversations with other analysts, whether through shadowing their activities, in the morning stand-up meeting, or over a cup of tea, helped resolve some of the issues I had when attempting to refine my knowledge in order to curate better detections. This helped me to explore the limits, the zones, and borderlands of the MAL in defining maliciousness, in curating senses of knowledges, affects and practices that come to categorise malware. Embodying malware took time, effort, resistance, and slippage. In my time becoming-analyst, I did not reach anywhere near a complete knowledge or feel for malware. But I attained enough to understand and partially narrate its stories, and the importance of translation of malware capacity and choice through the false positive, the Potentially Unwanted Application, and how this generates particular relations to malware itself.

⁷⁰ The *.NET* framework was designed initially for use on Windows, but is now an open-source, cross-platform that allows for common functionality. More information is available at: <https://web.archive.org/web/20180222035150/https://www.microsoft.com/net/learn/what-is-dotnet>

The False Positive

The ‘false positive’ is the incorrect *positive* identification of software as malware. The laboratory’s prime purpose is the correct identification of software as malicious, which due to its commercial drive, meant keeping the false positive rate as low as possible. Customers expect a product which only detects malware otherwise it would become increasingly useless. Reducing the number of false positives is also key to ensuring that the endpoint product does not slow-down the computing device a detection engine is operating on. So, how do you exclude clean software as soon as possible, whilst also maintaining a low false positive rate? Comparisons between competitor endpoint protection businesses comes from ‘real-world’ testing by external auditors, and by companies such as Gartner that produce ‘magic quadrants’ that indicate whether an endpoint provider is a leader or a visionary in the industry. How the malware analyst works within the drive of the MAL to reduce the false positive is convoluted and is tied into the processes that I learnt in becoming-analyst. This required *accurately* curating software as malware; which as part of a complex variety of computational agencies, is no easy task. Studying this allows an understanding to develop on why malware ecologies are difficult to compact into a patho-logic, as something neutral, and imbued with intention that can linearly extend from its author.

There are certain assumptions made to accurately capture malicious forms, but adequately avoid ‘clean’ software being detected. Not all detections are deemed equal. Those that have extensive human intervention (i.e. the production of signatures and behavioural detections) are where the (human, societal) malicious is defined and considered more reliable. You may remember in chapter four, where we discussed how ‘anti-virus’ split between virus-specific (broadly static) and non-specific (broadly contextual) ways of detecting malware. This distinction is still in play today – where the former is seen as more reliable than the latter. The analyst writes static signatures and (HPMal) behavioural detections, as well as developing contextual tools. But the false positive rate is found in both static and contextual strategies including behavioural detections, reputation scoring, and machine learning. Extensive work goes into developing technologies based on big data to produce reputation scores and develop neural networks (machine learning) based on previous categorisations of maliciousness, but these were not (at least yet) seen to have a low-enough false positive rate and must therefore be aided by human-intensive detections. In the compression of spatio-

temporalities of detection aided by contextual strategies in the MAL (it is faster to use both contextual and static strategies than static strategies alone), the false positive became of great concern.

There is an interrelationship between the speed and spaces of detection. Being acutely aware of the condensation practices that each computational component performs, in ways that often exceed the conscious awareness of the analyst, is crucial. By this, if I looked at contextual information, how much could I *trust* a piece of information? Over time, and through interaction with other analysts, I was able to tease out which indicators were important (such as certain rival vendor detections) and those less so (a software timestamp that is produced when it is compiled). In asking how this was blended however, analysts were unable to cohesively identify *why* it was more reliable and were always full of caveats. Or as in some training on the MAL's wiki, I wrote that "it is clear that it [writing detections] 'comes from experience' – why does it come from experience and what does that mean?" (Research Diary, 9 May 2017). The assumptions that take place in the curation of software as malicious are about the condensation of the signal from various anomalies, where the different strategies have alternative techniques to limit the false positive. The clearer the signal, the lower the false positive. The clearer the signal, the more likely that some things the MAL would like to tag as malicious are missed. It's a complicated balancing act.

A balancing act pervaded my training in how I was taught to tighten the parameters of detections. This required performing (internal) testing to ensure not only that my detection captured software I wished to include, but also what I wished to exclude (clean samples and other malware families and variants). I was not only reducing the potential volume of software in my detection but refining the malicious signal. This interplay is not necessarily something to be avoided. The exchange of information, the construction of the signal, is something that drives *innovation*. In similar ways to how Stiegler (1998) talks about the relationship between technics and innovation, there is an incessant drive to greater efficiency. This is also present in Gilbert Simondon's work on *concretisation*, particularly in *On the Mode of Existence of Technical Objects* (2017). Concretisation refers to the process whereby a technical object gathers a more consistent and coherent form. From this perspective, technical objects are forever in a process to self-define themselves, to become ever-complete and arriving at some sort of better process. The drive of the false positive, to

adequately identify the signal and avoid clean software (in its complex becomings) is what moves the MAL to develop further strategies. Yet, due to complex, changing spaces and environments that construct the MAL, there is no stable state in which technology can somehow 'adapt' to its environment and does not permit (cyber)space for choice. Remember that a computational, malware ecology, may be cognising and have political agency but not to a human-constructed sense of 'better.' This allows for a playful approach to technology, one to condense the signal, and condense the false positive, which does not always move forward in a linear sense.

There are numerous other individuals who are essential to maintaining the MAL, and ensuring the false positive rate is maintained at a low level. Quality Assurance (QA) maintained the technologies underpinning detection engines to avoid 'clashes' with the infrastructures that the MAL incorporates. QA occupied roughly half of SophosLabs and provided maintenance of the various systems that I and the other malware analysts used. This required a lot of work, curating the false positive sets, ensuring that detections worked in different environments (think of all the different operating systems, tools, and patches there are) and providing assurance that what the analysts produced did not 'falsepos.' This included what should be in tests, how the analysts should test, and making sure that these systems stayed online. It involved transforming the detections written in the VDL proprietary language into a form which the endpoint detection engines could recognise and use to detect malware. A relationship between the experimentation of the malware analyst and the industrial-like processes of QA emerges; attempting to control environments, performing testing regimes, and establishing rigour to condense the activities of the more-than-human analysis.

The 'false positive' rate was brought *under control* through testing. Through becoming-analyst, I was brought under control through the testing framework. When writing detections during training, I was given a set of software that had been previously identified as a malware family. I ran detections against a 'local' set of malware samples (on my computer), to see whether these 'fit' (i.e. detected most, if not, all). What to select was not pre-ordained, and there were multiple paths which could be followed, making the writing of detections a curative act. I was attempting to understand how a malware form moved, how it was different, and using this difference as a way to ensnare and to detect. Every detection I wrote had to be reviewed

by another analyst, after it had been through the MAL's tests, to ensure that it did not cause issues in different environments or detect against a set of software identified as clean. This was called the *falsepos rig*, maintained by QA, that included inputs from internal laboratory data and external feeds. This was not always smooth as I found out after one lunch:

“When I got back from lunch there was a message from Elliot via Jabber linking the Troj [a type of static] detection to a falsepos test... Elliot said we may have to ‘Di’ [a command to stop] the identity [a detection] if it was FPing. Quickly after this Daniel messaged me on Jabber asking whether I had spoken to Elliot and that Daniel thought it was likely to be due to ‘polluted’ data in the FP set. This was probably one of the scariest moments of anxiety I’ve had in the lab, where I needed to resolve whether I did have these FPs, and quickly, for the alert to be released [to endpoint detection engines].”

(Research Diary, 22 June 2017)

A delay to the alert meant that protection information was not being sent to endpoint detection engines, and potentially computers were vulnerable to malware as no detection was present. ‘Pollution’ in the falsepos rig meant the detection had included some samples that had been incorrectly tagged as clean. This was due to another vendor categorising these files incorrectly, which had been used as inputs into the internal testing datasets. Though my detection was thankfully detecting *correctly*, it demonstrates a fragility and difficulty in maintaining and curating data feeds and multiple environments that converge in the MAL. The ‘wild’ environments outside the laboratory are forever changing, part of complex ecologies meaning that the MAL only ever experiences it partially, that in part contributes to the false positive. The work of QA is to maintain an awareness of these ‘wild’ environments and ensure that the MAL is *stable*, somewhat like technicians in other laboratories.

The Silent

As the false positive is affected by environments in which malware perform, certain formulations of curation must be maintained (as illustrated by the impact of ‘polluted’ data). When software does not explicitly fall within clear bounds of maliciousness, the detections and systems that support the MAL may incorrectly identify software as such. The use of SXL telemetry – data from ‘wild’ customer endpoint detection engines – inform the MAL’s contextual strategies for analysis and detection. This has particular resonance with the ‘silent’ detection. The *silent* is written to report whether a detection would potentially ‘fire’ (i.e.

detect) on a customer endpoint detection engine or not. This detection acts as if it has detected malware but only sends data to the MAL through the cloud from detection engines. This detection was used when it was deemed too risky to release it for immediate use on customer endpoints. The silent detection identity allows for (limited) contexts to be developed around a detection. As the detection identity runs in the cloud, it is activated whenever there is an SXL 'look-up' from an endpoint against the targeted malware. This meant if too many lookups were being made, I could stop the silent identity easily without an impact on customers. A tapered introduction was used to facilitate this, so that if there were too many cloud lookups (which could overload the system and cause outages), this could be cancelled before all endpoints would perform a cloud lookup. The data collected through SXL allowed me to derive the efficacy of my detection; such as the number of lookups and whether these were on a particular malware form (as I had defined it at least).

Riskier detections (in terms of the potential false positive) could then be tested based on multiple computing environments. In one case, I wrote a static signature using the string check (an alphanumeric concatenation), "setup.exe" (Research Diary, 13 June 2017). This string exists in a lot of software. As this was a latter check, coming after other genes, it was *likely* this was a 'tight-enough' detection. But, due to the selection of genes using mainly contextual features, it was deemed 'not strong enough' during review. "Terry said that my [string check] on 'setup.exe' had a small chance of having a FP rate which meant that I should look at putting in a Shh (silent) identity" (Research Diary, 13 June 2017). To reduce the likelihood of it *falsepos-ing*, the silent detection was a method to utilise data of the 'wild,' environments of customers, to reduce risk. When I came to review this, Terry had been correct, and though there were only 13 lookups, it was firing on another legitimate file – so I stopped the identity and went back to the drawing board (Research Diary, 19 June 2017). This blends what could be called a 'field trial' and a 'controlled experiment' – as Lorimer and Driessen (2014) have explored in the development of 'wild' experiments to explore rewilding at Oostvaardersplassen in the Netherlands - by working on 'field' data, that of computing environments, yet controlled in its silence and tapered growth.

The feeds and data given in contextual strategies play a significant role in reducing the false positive rate in the 'wild.' In the past, there was a limited ability to help consider how

detections operated and how successful the MAL was in *defining* and *curating* maliciousness.

As Joe put this:

Joe: “You basically got a detection, sat down and thought, I wonder if that will false? Test it as best you can on a false positive rig, and you know full well the limitations of a false pos rig. However big it is.”

Andrew: “Yeah.”

Joe: “Er, and then you released it and you crossed your fingers and if it’s detected loads of bad things, you never knew about it.”

(Interview with Joe, 21 July 2017)

Contextual strategies have blended with analyst embodied senses of maliciousness and intermingled with other detection technologies to assess the false positive. This is an important driver for the MAL. It drives the production of a more industrial-focused production of detection in endpoint detection engines, that through QA, combine and fold the experimental, curative detections that analysts wrote. This folding of static and contextual strategies in a patho-logic to attain a signal is consistently constricted by the spatio-temporal logics drive to define malware at *appropriate* speeds (i.e. the need for the contextual). Hence the false positive rate emerges, and must be controlled, in the balance to speed-up and increase detection release. This is the compromise of the MAL, to tackle greater volumes of software but also to maintain a low false positive rate that become more difficult to achieve as anomalies of contextual strategies are translated to abnormalities.

The Potentially Unwanted Application

A definition of maliciousness cannot be applied to all anomalies however; as they are a condition of software. The MAL, however, encountered software that did not achieve a threshold of maliciousness but were not necessarily benign either. Hence enters the Potentially Unwanted Application (PUA), otherwise known as the Potentially Unwanted Program (PUP). PUAs include software an individual may not choose to run in their environment, but at other times may wish to do so. Frequently PUAs ‘bundle’ other pieces of software without authorisation, such as; adware (that changes adverts on a computer and monetises this); dialers (software that rings phone lines and incurs charges for the user); certain forms of spyware; remote administration tools, and; hacking tools (toolkits with multiple vulnerability exploits). In different environments, and ecologies, these could be

determined to be useful to a user, but in others deemed unwanted, and to some even malicious.

I became interested in the construct of the PUA in the MAL as it is something of a cusp, something on the edge, something that seemed to bring together the difficulties of exposing maliciousness. The category of the PUA exposes the fallacy of malware as something that exists *a priori*. Here, we have a clear demonstration of the performance of the bounds of maliciousness by the MAL. Malware are curated. Software are curated into categories. It exposes frailty of signals, with noise everywhere, where there is *no clear maliciousness*. Here it becomes crystal that the *intent* of an author does not necessarily align with the MAL's interpretation of maliciousness. There was an odd bind at work in the MAL; where I (and the multiple computational analyses and detections) did not have access to the intent of the author(s). It required a second-guessing what the intent of the authors may be in order to work out if something was “fishy” as Lucas put it:

Sometimes you get into a little trouble you know with er DRM [Digital Rights Management] types of things where it looks fishy [laughter]. Erm, but ... you know some more uncommon build processes.”

(Interview, 20 July 2017)

Even ‘legitimate’ programs could be brought into a zone of indistinction. By this, there were judgements to be made, through a mixture of affects between malware analyst, computational agencies, and the ‘wild’ environments of endpoint engines. This produced a category that cannot be easily distilled in the *abnormal*. Software categorisation meanders according to human expectation and even though the PUA may be under a pathological lens, it is indeterminate if the anomaly translates into an abnormality.

A strong interest for me concerned how the MAL dealt with *falsepos-ing* in relation to PUAs. I have explored what may be malicious or not (that can often be clear – such as stealing credentials and encrypting files). However, the PUA sits at the edge of this distinction, and became part of a solution to satisfy a commercial drive to not ‘FP’ (false positive):

“You know we were scared of FPing because people did used to complain, because like ‘we [the customer] use Metasploit⁷¹ [a hacking tool] internally to test and you’re detecting it!’ So, we had to set things as PUA ... It is

⁷¹ Metasploit is an open-source penetration testing kit that has modules created to exploit vulnerabilities. This is used by both ‘white ‘ and ‘black’ hat hackers. It is available at: <https://www.metasploit.com>.

potentially unwanted, so that they could authorise it, but then other customers don't have PUAs on and then it's like, why did you get owned [attacked]? Because the customer kind of thought we were falsepos-ing so we couldn't make it a /Troj [a malware detection type]. So now we just kind of go bollocks to that, like detect everything.”

(Interview with Nathan, 21 July 2017)

Through Nathan we find the complex curative actions the MAL has to take in order to consider if software are a PUA or not. A complex relationship was thus maintained about the *expectation of maliciousness* beyond the MAL. It was constantly (re)negotiated with customer environments in which the work of the MAL operated and was balanced with the need to maintain low false positive rates *and* maintain varying expectations. In this case, Metasploit⁷¹ can be used by both legitimate ‘pen testers’ that test the security of an organisation, and by individuals with malicious intent. False positive rates however can be aided by PUAs for software that do not acquire a threshold of maliciousness. By this, customers can use the endpoint engine to decide whether to let a PUA run or not – passing the definition of maliciousness to users and enrolling them into the process. This process ensures that ‘legitimate’ software can be rendered as *potentially unwanted* rather than malicious, and thus at the same time dealing with the complexity of ecologies of expectation where there could be a false positive.

This has been complimented by movements by different industry bodies, such as AppEsteem⁷², to produce certifications for ‘clean’ applications and software. This is to encourage ‘good’ behaviour by software developers who may engage in questionable tactics – such as bundling unwanted applications (such as adware or other programs) – that exhibit *anomalous* characteristics but do not raise them to a threshold of maliciousness. Hence, we see a form of politics operating – where ‘legitimate’ actors are encouraged to subscribe to what MALs and certification bodies see as normal. Over time, curative processes among MALs have led to certain performances, certain structures to be deemed as unsatisfactory (in tandem with broader moves in understanding cybersecurity). As geographers and others have referred to thresholds in algorithmic decision making (Amoore and Raley, 2017; Aradau and Blanke, 2018; Amoore, 2011; Crampton, 2016; Noble, 2018; Roberts and Elbe, 2016), this is also at work in the MAL. This is what *the* curative exercise is about – delineating where

⁷² See <https://web.archive.org/web/20190115162816/https://www.appesteem.com/> for more on AppEsteem.

these thresholds of maliciousness are. Whether by the affectual capacities I developed as becoming-analyst, in the signals produced in computational and algorithmic processing, or in the PUA as an expression of the fuzzy boundary of this threshold.

Not all techniques used within the laboratory produced the same relationship between the false positive and the PUA. For example, machine learning techniques struggled to work with samples detected by analysts as PUA due to the neural network algorithm's complex, interweaving relationship between the anomalous and the abnormal. Due to the conflation between the clean and malicious, the PUA is frequently avoided in machine learning as the *distances* between software spatialities (features) are not distinct enough (Research Diary, 3 March 2017). This is due to the ecological conditions under which software operate. Neural networks are intricately tied to the folding of affects and knowledges of analysts, other technologies, and multiple MALs in how they come to form the training data for machine learning. The PUA is often closely tied to legitimate processes. For example, a remote administration tool may be used in a school to observe what students are doing on their screens or a university may remotely update computers to install patches (indeed, the University of Oxford does this), but we may not want this on our personal computers. This could be compared to older forms of whitelisting, as Alex described in an interview (21 July 2017). Certain organisations and individuals may want to run a piece of software but in others not – and are categorised as PUA as they can be used in the different ecological alignments for human sensibilities to maliciousness (based on intent). As PUAs are so dependent on a socially-constructed norm, it meant that basing decisions on the anomaly, as in computational modes of (re)cognition, became exceptionally difficult. This is because the threshold of maliciousness and abnormality are not as well aligned to distinct groupings of anomalies. Thus, if current computational methods for identifying and allocating the category of PUA were made without extensive human engagement, there would be *an exceptionally high false positive rate*.

In developing my becoming with PUAs, I spent several weeks, between other tasks, working with one particular software form that I will call 'Scary Install.' This is a software bundler that connected to a URL to download an additional file and gave itself enhanced permissions on a computer - which were not authorised by the user downloading the software. It claimed to be able to monetise traffic through using it to download files. Over the course of June and

July 2017, I updated the PUA detection for tens of thousands of variants which were being reported through a variety of external feeds. As I updated the detection, the closer I became to understanding and engaging with Scary Install and the tricks it was using to *avoid* detection as a PUA. In this – as I was writing static signature detections – it meant that I encountered new samples nearly every day which had tweaked some code or other features which ensured they avoided my detections (and no doubt the detections being produced in other MALs!). In so doing, they (authors and software) avoided detection at ever-greater lengths. There are two moments of slippage here. First, this taught me that while malware can be very much tied to the intent of its author – it is not directly tethered. The author is an important part of what makes software malicious or not, but it not the whole story. This makes the second moment – the boundaries of whether Scary Install was a ‘deceptor’ or whether it could be malicious – increasingly hazy. Terry and I had several conversations via Jabber about this:

“Terry messaged me on Jabber complaining how randomised the “ver info” [version information] is and that it’s a shame we only have fuzzed icons [to detect on]. Once I finished [detecting] these icons... I messaged back and said I think we’re about to reach the limit of this sort of detection due to the increasing frequency of the changes and also potentially the creators fuzzing the icons at greater frequency. Terry agreed and said he was thinking of upgrading this to a [malicious category] Troj/ due to its behaviour.”

(Research Diary, 13 July 2017)

The crucial part of this moment was trying to track the software’s behaviour, and from this trying to infer if the authors of this software were acting in a *legitimate* way. This is distinct from seeing malware without context. As I have outlined above, in detecting software as malicious, I was attempting to trace the logics of the code. This required being able to observe an *ecology of practice*: how it was understood to work, how the user interpreted the software’s use and permissions, and how the authors presented this all became important. Software’s instantiations are geographical, dependent, and constructed. The example of the PUA demonstrates that there is a politics to software designation when it does not fall directly and neatly into categories of normal/abnormal and benign/malicious; where an ecology of intersecting and co-dependent environments, people, and software performance all contribute to variable and changing categorisations.

Making Sense of Maliciousness

So far, I have been concerned with *how* the MAL defined maliciousness – but not with how these processes became ‘narrated’ or became known. Extensive labour goes into developing metrics and visualisations for consumption both within and outside the laboratory. This can be for internal awareness, providing information to customers, or for discussion with journalists and visitors to the MAL. These curations attempted to develop a sense of trends, malware propagation, and develop anticipatory actions to both justify investment and to sell, ultimately, Sophos’ endpoint detection products. Making sense of maliciousness is not only about *improving* detections, but also constitutes part of a broader set of anticipatory political practices in security (Adey and Anderson, 2012, 2011) which draw on neoliberal concerns for self-protection mediated by investing in privatised infrastructures.

Naming and Narrating

The process of narrating and curating malware condenses malware vibrancy. I did not provide a detection for each software form (i.e. one executable file) I deemed to be malicious. This would be an inappropriate action in most cases – a massive waste of my time and ineffective against the rapid growth in malware volume. In generating and representing detections which seek to identify a wide spectrum of malicious software variations, the code I wrote effectively ‘flattened out’ the distinctiveness and performativity of different software forms. Yet, there are exceptions. Sometimes a rapid, individual ‘blitz’ signature was required. For example, in May 2017 with WannaCry and in June 2017 with (Not)Petya, it was necessary to rapidly deploy a detection where detailed analysis was not possible. A main function of the MAL is to categorise, bring together and group different software forms through detections to families. This is not dissimilar to comprehending the order of animals, and their categorisation into a variety of levels. This often comes through examining similar traits and finding code comparisons, such as through ‘YARA’ rules⁷³. With the advent of DNA sequencing, the ability to differentiate between different animals and their categories has radically transformed. In malware analysis, a similar biological narrative can be seen. Familial connections and genes (with shared characteristics), identified through the use of dynamic

⁷³ YARA rules are another way to find similar traits across malware, and are similar to static signatures. For more information see Arntz (2017).

search tools and behavioural techniques, brought malware together in different groupings, as I discussed in chapter four.

It was possible to dynamically search for different attributes based on contextual information, such as other MAL family names, but also through internal genes. The gene, one may remember, is used to 'flag' on some particular indicator (such as a customised packer that is only known to be used with malicious software) or certain common traits (such as if there is an icon, and the icon hash value) that can be used in certain combinations to tease out connections among disparate software forms. The MAL uses genes (particularly in the formation of static signatures and behavioural detections, HPMals) in combination along with specific checks on strings, or sequences of code. This helps bring together familial connections. This draws in samples that initially would not be recognised as malicious by analysts nor a range of technologies, and assists malware analysts in constructing, and condensing simultaneously, the bounds of maliciousness.

The malware family was actively constructed through the process of writing and narrating detections. However, the families defined by the MAL were no guarantee of a shared genealogy. Though there were undoubtedly connections in those who partook in the authorship of malware this does not suggest that all families are accurately tied to their authors. After I analysed a new form of malware, I needed to construct or edit a detection. I usually filtered known variants in the dynamic search and observed what the most common genes were. Families could thus be created by using these previous gene categorisations, or what had been noted to be so from another vendor. Over time, I became better at identifying certain genes and vendor detections that seemed to detect certain families according to this experimentation. A variety of things could happen; to identify code sharing, potential flags for attribution, and indicators of how common a certain form was in the 'wild'. These condensations, or curations, were essential to categorise malware. However, it is not necessary to find a shared author for a piece of malware to produce a successful detection. As long as it was malicious and roughly looked like a particular family, it was more important to detect rather than find the most appropriate grouping. This would produce some structures or attribute information to work with and I would *play* with the different variants, experimenting with different combinations.

Consequently, the stories narrated by the MAL about the origins and associations of malware may or may not correspond to the actual sites or associations from which that malware emerged. Sometimes too, the external environment can intrude and shape the ways in which malware are narrated. For example, naming conventions have a history in terms of the broad categorisations of virus, worm, and Trojan Horse, however, the individual names of malware forms are much more haphazard. The way an analyst approaches the issue is dependent on a variety of constraints; namely temporality, other MAL naming, and evidence found during analysis. In the naming of (Not)Petya, there was a very hurried decision, as a blitz detection was needed to be distributed as soon as possible due to its widespread impact and pressure to protect customers:

“Daniel said out loud ‘does anyone have an objection to me naming this Petya?’ Which Mason said no, and Daniel mumbled that it looked like it anyway. At this point this was the determinant of what the malware was going to be in a highly-pressured environment that needed to produce a detection extremely quickly.”

(Research Diary, 27 June 2017)

During (Not)Petya’s emergence, Twitter already had several tweets stating that this was a Petya variant. It later turned out to be not so, hence why I write (Not)Petya (there are a variety of alternative names that revolve around the core of ‘Petya’). In the majority of cases, naming conventions came through external feeds, taken as an industry standard and henceforth used. In other cases, there may be some code that gives malware a name, such as a distinctive string that gave one family of malware the name *Tryca*⁷⁴ whereas another MAL had given the name ‘Faker,’ but which Sam disagreed with (Research Diary, 24 April 2017).

The condensation of maliciousness into families and groups, is not a solely human analyst activity, but is mediated through technologies, more-than-human senses, and computational (re)cognition. Machine learning detections often do not often have human-friendly names, generating automated strings, where forms and categorisations emerge beyond the direction of human-centred analysis in the multiple dimensions of neural networks. Information sharing is an essential part of the MAL, as I have hinted at, through the main communication tool, Jabber. Yet there are a variety of platforms, including Twitter, the internal wiki, and email that all took part in the condensing the knowledges, affects, and becomings of the

⁷⁴<https://web.archive.org/web/20180304173037/https://www.sophos.com/en-us/threat-center/threat-analyses/viruses-and-spyware/Mal-Tryca-A/detailed-analysis.aspx>

laboratory into something discernible. Jabber became a place for the me to ask little questions, share information, and develop relationships with other analysts – especially with those based in Hungary and Canada. These communication tools became the places where analyst knowledge was shared according to discussions (much more than email, which was rarely used in the MAL for non-official communication). Jabber was a place where notifications told me of new batches of detections released to customers, and how the Indian ‘Frontline’ dealt with emergent issues, customer interactions, and processing incoming software samples.

The internal ‘wiki’ was the only place to maintain and retain the fragmentary knowledges of the MAL together. Information here included the use of tools, procedures, training materials, and other information relevant for the functioning of the laboratory. I often flicked back to the pages of the wiki to find information and procedures (and it was frequently a way for me to avoid the awkwardness of, yet again, asking someone for help). It held a lot of information – comparable to a MAL ‘archive’ – with information dating from decades before, much information outdated. For me, however, it became a space of entrancement; I became fascinated by the way technical information was shared, the guides to the use of tools written by analysts, information on what was currently being detected, major software families, and how the systems operated. This repository, though always fragile and frequently out-of-date, was an invaluable resource. It provided insights into how experienced analysts had carefully narrated these more-than-human constructs. It was often technocratic, but there was no doubting that many, through reading through these wiki pages, had a clear dedication to the work. This is not to claim it is ‘exciting’ and indeed boredom, I would argue, is a crucial, and productive part (Anderson, 2015), of being a *good* analyst due to the slow, tedious picking through software, particularly in static strategies. However, the wiki, is in part emergent from boredom, of ‘extra’ time to write information in a form for others. Boredom, as a part of the practices of the MAL is essential in understand *what* is often retained in the repository rather than determined dedication (as evidenced in the multiple unfinished pages, technical details missing).

The wiki, however, was not simply a repository but became a lively space at certain times. This became most pertinent during the emergence of WannaCry and (Not)Petya. In both events, a wiki page was created as analysts pooled information to speed-up analysis. On the

afternoon of 27 June 2017, I added information about a shutdown ‘pipe’ during rapid co-analysis of (Not)Petya. Along with other methods to shut-down a computer, the collective knowledge was added, edited, and revised and soon suggested that (Not)Petya was not a conventional piece of ransomware but, as was later revealed, a wiper⁷⁵. The pooled information sharing on the internal wiki demonstrates the sometimes-fragmented conditions under which curation happens. Twitter was the source for the SHA-256 (Not)Petya hash in order to acquire a sample. It was processed through the sandbox, Sandstorm, for further information, statically and behaviourally tested in virtual machines by analysts, and even before this, exhibited behaviours similar to ransomware and was detected by behavioural detections and by neural network algorithms as part of Sophos’ ‘next generation’ product, *InterceptX* (Ducklin, 2017). The wiki does not stand alone, but connects with information from Twitter, Google searches, and others to ensure it is a dynamic space of interaction at certain moments, which then may fade, not be finished, and remain as archival material. Thus, the MAL must be part of a record-keeping, arriving at categorisations, and sharing knowledge internally to ensure there is some stability in order to use this in the future; to train analysts, to refer back to when working with an irregularly used technology, or to record what analysts were working on.

Visualisation

As the MAL condenses and narrates maliciousness, it must move from singular software samples to general trends in order to focus its labour. Using the detections created by Sophos and other MALs, it is possible to craft and curate trends, developing senses of where to direct attention next and to anticipate malware futures. Visualisation is a process of comprehending the collective malware categorisations that are crafted in and beyond the MAL. This requires significant labour, with new visualisations emerging during my time in the laboratory to track the effectiveness of certain detections and for identifying new trends. And there was always noise that was incessantly aiming to elude capture in the visualised form and distort it. During one shadowing session with Charlie, we worked through a day of new samples that were feeding into a trends visualisation.

⁷⁵ The article on Kaspersky’s *SecureList* website by Ivanov and Mamedov (2017) was one of the first to publicly comment on 28 June 2017 how (Not)Petya was not ransomware but a wiper after it was realised that a hard disk could not be decrypted even after a ransom payment had been made.

“We went back to malware which hadn’t been automatically classified by big data analysis. In the past day there had been 1500 new samples [of a particular form], of which around 300 had not be identified [automatically]... here Charlie had to manually change each bit with different tags according to that particular malware form.”

(Research Diary, 21 April 2017)

This ‘error’ of identification came from which malicious component of the form was the ‘true’ classification. In this case, the *misidentification* came from external feeds which had classified a PDF file as the main vector (of infection) in which the malicious activity had occurred. Yet, it was an ‘OLE’ document (the standard for Microsoft Office and open-source equivalents) that had been embedded within a PDF containing a malicious macro. To break this down, PDFs are formed of several ‘containers,’ which allow for other software to be embedded. In this case, a Microsoft PowerPoint file was within a PDF. The OLE structure enables the use of macros, which are a common method to automate tasks, but which are also used by (malicious) authors to enhance permissions and download other files. Charlie and I worked through each file to clarify whether each one had been categorised *correctly* (or as the case was, not). This process was part of big data’s condensation to feed the visual, to see and observe how malware trends were changing. This limited malware to singular categories, to a point where maliciousness was deemed to be, splitting the form, and reformulating it according to what was identified as *the* malicious component. These visualisations were not simply reflections of malware but were actively constructing what was and was not malicious. As Gillian Rose (2016) and Divya Tolia-Kelly (2016) have argued, visualisations are not simply reflections of a ‘true’ state but are constructed, and through their production, form new worlds of understanding. By using visualisation, malicious performance was reduced to colours and shapes, rendering malicious agency as something to be observed, in similar ways to the medical gaze (Foucault, 2003b) or in the representations of bodies on security scanners at airports (Amoore and Hall, 2009).

Technologies to aid in the production of visualisations included Datameer[®], a big data analysis tool. Datameer helped to collate information about different families based on SXL telemetry data, crafting a *proactive, anticipatory* monitoring of ‘missed’ samples in families, and new tickets for action based on this. The visual and active monitoring of ‘missed’ malware cannot be detached from the active contextualisation of the MAL where visualisation determines greater accuracy in big data (Kitchin, 2014), which are identified

through anomalies. As I discussed in the previous chapter, this production of knowledge through big data transformed the MAL from simply processing software it received, to a ‘data-driven’ approach that identified what the labour of the MAL should be. Curation has changed from simply encountering samples, analysing software and defining it as malicious, to developing a new anticipatory approach, with “visibility” (Interview with Joe, 21 July 2017) being part of an anticipatory commercial drive to find, target, and detect malware previously ‘unknown.’

Visualisation then became enrolled in contextual strategies that are not directly related to analysis nor detection. It is a curation, a condensation that came together in the feeds of SXL, the MAL’s patho-logic, and commercial drive to provide anticipatory detection. That is, in similar ways to how disease becomes mapped, compared, contrasted and abstracted from the environments and ecologies in which they emerge, malware becomes abstracted, something to mix, pull-apart, and reconstitute. These reconfigure the relations to anticipation (Anderson and Gordon, 2016; Anderson, 2015; Adey and Anderson, 2012) – on what should the analyst’s body be orientated? This abstraction from the specific does, however, allow the MAL to formulate and articulate responses on priorities, which malware should be detected, what tools should be developed, and where to invest (capital and labour). Visualisations provide a way to talk outwardly to others; whether that be to visitors who came on ‘lab tours’ where visualisations were presented on large screens at one corner of the MAL, or to share with journalists and other organisations to build ‘pictures’ of the malware landscape. It is at these points where malware becomes translated outside of the laboratory.

Translations

Translations processed through MALs are crucial to how *we* - ‘publics,’ governments, businesses - understand malicious software. What is published by Sophos or other vendors, such as Kaspersky and Symantec, are important and the strategies that I have detailed in the previous two chapters all contribute to how responses are formulated. Curation, in its dual meaning between a tethering together and the *poisons* and *exclusions* that occur with this, but also in its ‘cure’ – to detect malware – mean it is also pathological in its desire to render malware as something to be seen as an object of human intent. As is detailed in the work of Derrida on the *pharmakon*, the poison and cure are within itself. I think curation deals with malware that cannot be, in the conventional affective sense, touched, seen, or grasped. To

detect malware *at speed* within the limited capacities of MALs requires both abstracting software from environments and only reintroducing this factor in a limited sense through contextualisation through big data or in sandboxes when necessary. It is something that requires specific expertise (learned, for example, through seven months of ethnography in a MAL). Curation is then the method by which MALs make sense of malware. It is about communicating to the public that draws on popular tropes of pathology such as infection, viruses, and immunity, that still inform and perpetuate our current imaginations.

Malware analysis and detection is not a neutral practice. It actively interprets and (re)presents malware. Whether this be through the various blog posts by Sophos on its platform, *Naked Security*; through the hues of blue and grey on networked visualisations provided for customers; on the visualisations featuring charts and diagrams, and sometimes cartesian maps, of malware distributions, detection levels, and ‘infection’-rates; or in discussions had by Charlie, Nathan, and Joe with both customers and journalists – which led to quotes in leading UK newspapers – to advising individuals in government. This is *not* to say that the MAL is somehow actively involved in a conventional human politics. Nor do they seem themselves as actively determining something – rather they claim to be ‘objective.’ Nonetheless, I argue that a politics emerges from the banal, the everyday activities of boredom, searching, sensing that occur between the analyst and multiple computational technologies as they seek out, sense and curate *abnormality*. This process generates a social-material understanding of the normal that is political precisely because it makes a movement towards normalisation; in Foucault’s terms, it generates a discursive norm against which software (and those associated with it) are judged and politically located. This is no bad thing – searching, detecting, and stopping software identified as malicious. However, in accepting this as a ‘good’ does not mean that it is devoid of a politics (as much as good itself is a political negotiation!). Furthermore, this thesis argues that the way in which malware are detected, curated, and narrated through largely pathological imaginations, and in the practices, words, discourses, and even affects of how we are advised to ‘deal-with’ malware-related insecurities, is worrisome due to its lack of appreciation of computational agency.

Moving to Ecology

“The increasing use of large data volumes and distributive interactive systems in design has not only pointed to the limits of deductive logic (the

general includes the particular) but also diffused the use of inductive methods of heuristic thinking (starting from the particular and proceeding by trial and error to arrive at the general) in which the realm of physical contingencies and not of mathematical formulae are said to be central to computation.”

(Parisi, 2017, p.77)

Parisi argues that deductive logic has been questioned by increasing volumes of big data that have led to a growth in heuristic methods, which can be observed in the MAL in treating software statically and contextually. Contextuality introduces alternative temporalities to malware analysis and detection, and the pathological, with its imaginations of computation infected by an external viral agent directed and controlled by hackers. Increasingly analysts talk of malware forms or families, as opposed to adopting more generic terms such as virus or worm. Yet some aspects of a patho-logical imaginary remain, both the ‘deductive’ (static) and ‘inductive’ (contextual) models of malware analysis still approach malware as something that can be dissected, understood, grouped, and rendered knowable. I wish instead to move to an alternative curation, one which simultaneously repositions the human (as author and analyst), but which does not lose sight of the (important) mathematical logics upon which computing is built.

This requires exploring the choices and capacities of malware, computational cognition, environments, and humans together. This requires re-evaluating how space is understood in cybersecurity. I argue that part of the problem with the pathological approach is that it imagines cyberspace as a relatively flat geometric and Euclidean terrain across which it is possible to draw hard lines and borders, and to track malware back to its point of origin. Such a perspective fails to recognise how cyberspace is performative, challenging the notion that humans as the only political actor in security. Seeing malware as only the extension of an all-knowing human author plays directly into the ‘god-trick’ Donna Haraway (1988) warned of. As she states with reference to the eye, “vision in this technological feast becomes unregulated gluttony; all seems not just mythically about the god trick of seeing everything from nowhere, but to have put the myth into ordinary practice. And like the god trick, this eye fucks the world to make techno-monsters” (1988, p.581). This trick is often applied to all software, seeing them as something humans have built and therefore control, leading to anything unexpected being an ‘error.’ Malware are a very effective example for exposing the fallacy of this perspective. Governments spend huge amounts attempt to *control* it, including

funding large research programmes (not unlike the one I am funded by), all seeking to respond to this threat to humanly dominance. Yet, like Amoores and Hall's (2013) clown at the gates of the camp, malware acts as contemporary cybersecurity's joker; its agencies spilling out and eluding capture. I therefore advocate extending the pathological imagination with a new ecological perspective, one which radically repositions the human (in production, analysis, detection, curation, politics among others) in the story of malware.

The anomalous and the abnormal expose the intricate agencies and powers at work in the MAL. As Annemarie Mol has shown in a very different set of environments (and ecologies), around atherosclerosis of the leg vessels at hospitals, "medical practices perform their objects" (2000, p.82), where "the term performance nicely catches the coincidence in time between diagnostic activities and the disease they deal with" (2000, p.83). This comes through the conversation of the medical doctor, through the microscope, in the thick cell walls of the vessel. Multiple realities are exposed between those who practice pathology and those in clinical practice, that lead to multiplicities of reality (Mol, 2000, 1999). Software is performative – and the MAL's contextual activities I have identified expose the limits of pathologies in the face of the multiple realities and ecologies through which malware come to matter. While malware's particular performances depict certain and limited resonances (or affects) that are interpreted by us, at least by a specialised set of malware analysts, as something malicious, this should not be equated to 'execution' – which is a reductive understanding of computation logically processing *as if cyberspace and computational materiality are not performative*. Performance through choice then renders malware forms such as Stuxnet, the Dukes, and WannaCry/(Not)Petya alternatively. The pathologies of the MAL are a starting point, a departure, to not move beyond them, but complicate, contrast, and extend these.

Curation is the crucial activity of the MAL; performed through the alternative norms of abnormality and anomaly, through human and computational cognition. It becomes in a complicated mixture of different technologies, tools, conversations, and affects that a becoming-analyst must follow and engage in. In exploring the complications of the false positive and the PUA detection that drives the commercial logics of the MAL, there is an incessant negotiation of what maliciousness is, and what it could be, drawing on increasingly anticipatory methods for narrating and making malware present. These processes of curation

exceed and move beyond the MAL, to include the environments generated by and from customer telemetries and other data feeds from multiple ‘wild’ environments, other MALs, and neural networks. Analysts too are shaped and conditioned by their contexts. When I was becoming-analyst I was already conditioned in and saturated by the lineages that I partially explored in chapters three and four. They shaped static and contextual strategies for analysis and detection, and thus enact certain renderings of the world.

The lineages, narratives and logics that construct the MAL *also* construct our understanding of the agency of malicious software, and how this becomes shared outside of MALs; through reports, blogs, newspaper quotes, and in providing advice to governments. These productions are essential in understanding the wider ecology of how we understand malware – that are connected to the pathologies that infuse the MAL. Malware are not something intrinsically ‘bad’ but are performed in a more-than-human register to be so. Some of these are undoubtedly clearer than others, and I do not wish to say malware are *good*. I wish to respect the selection and decision that Stiegler (2017) identifies as questions for ecology to embrace; in what is included and excluded. The visualisations and condensations of malware vibrancy into families, categories, as allocations on visualisations, all emerge from a complex range of environments and agencies. Therefore, we must respect the productive curations of the MAL, but also acknowledge these are curations, exhibit a politics of normalisation and pathology. Therefore, I have explored the spaces of the MAL to interrogate how we have come to understand and work with malware today. This is in order to expose the limits of this approach in how we relate to malware and its impact on various politics that we tend to assume are human, but are in fact influenced by more-than-human actors. In moving away from pathology and towards ecology, I do not leave it behind, but complicate it as I do so in the next chapter.

Chapter Seven || Malware Politics

In chapters five and six, I presented how malware analysis laboratories (MALs) deal with malware through a patho-logic, and how this shapes the ways in which malware are detected and curated in the MAL. I instead use this chapter to examine three malware cases to expose different, intermingling, components of what I constitute as an ecological ‘malware politics.’ This includes the translation of the Sophos MAL’s work into an everyday politics that includes cybersecurity preparedness, the protection of ‘critical national infrastructures,’ and the security of the UK’s NHS. Malware, as I have argued, are performative and make choices that can exceed author intentionality, and here I offer perspectives to concretely thinking and approaching this. Furthermore, as in Barry’s (2010, 2013) depictions of the politics of metals and metallurgy, they impact politics, ethics, and publics in the (re)construction of ecologies. Yet, whereas Barry (2010, 2013) and Bennett (2010) focus on the *affective capacity* of things to spark political debate and generate matters of concern, I also consider computational choice to continually intervene in and to shape political matters. I thereby argue that (i) choices by malware are a key part of its ecology, and that (ii) an ecological understanding of malware is central to understanding its capacity to act politically. A more-than-human (malware and computational) politics can be glimpsed through re-reading some famous malware forms, that build upon my experience as becoming-analyst where I developed specialised malware knowledges to permit such as reading.

Through Stuxnet, the Dukes, and WannaCry/(Not)Petya – I intend to open-up what could be considered minor, alternative narratives in a new curation to understand how they have become to be named, known, and *influential*. In this, I acknowledge that all three malware cases have been attributed to state actors. I do not intend to add nor negate to these analyses, which are important and often convincing, but I question how and why human intention has become a central feature in a discussion of malware politics, and demonstrate how an ecological understanding of malware complicates the process of attribution – with extensive work on its difficulties (Buchanan, 2016; Lindsay, 2015; Lupovici, 2016) as well as its political impacts. Some of these malware forms have been destructive, causing great harm to many individuals, communities, businesses, and states. It is tempting in such ecologies to focus on the harm that such software causes, but here I want to offer a more-than-human approach

to the study of malware and its impact. This is different to the ‘diagnostics’ of pathology, searching or case hunting that comes in epidemiology, as I wish to suggest malware’s political effects extend beyond its pathological consequences. At the same time, I want to make a case for a rethinking of malware agency, questioning the common assumption that it is *solely* environmental variables that lead to divergence or unexpected ‘outputs.’ In place of an environmentally-determinist reading, I wish to draw attention to the way in which malware make choices, and in so doing, can be seen as a new kind of political actor.

It is no easy task to set up a ‘politics’ of malware. Indeed, there is limited scope for malware to exhibit a form of politics, *but that does not mean it does not exist or is unimportant*. To do so, I begin by outlining how choice and affectual capacity work in ecologies. I then turn to the three cases; on the effects on a uranium enrichment facility in Iran (Stuxnet), the tying of different malware forms to a Russian ‘Advanced Persistent Threat’ (APT) group (the Dukes), and to the release of the US National Security Agency’s tools that led to the creation of two of the most famous, and damaging, malware forms of the early twenty-first century (WannaCry and (Not)Petya). In concluding, I draw the examples back into a broader discussion of how an ecological perspective reconfigures our understanding of materiality and its implications for the political agencies of malware.

Ecological Curations

To develop a *more-than-human* politics, I draw on the work of Jane Bennett, N. Katherine Hayles and others in the pursuit of computational agency and affect that I initially explored in chapter two. This recognises, as it must do, that humans play a large role in interpreting, analysing, detecting, and translating the agencies of malware into a politics, even if that politics is not reducible to human intentions. As I discussed, there are particular ecologies within which computation become and (per)form. In this chapter, I move to a (re)cognition of a malware politics – through two strands; first in its power and affects, and second through its cognition and choices. The former belongs to all things, as examined in Jane Bennett’s work on the *political ecology of things* (2010). Whereas the latter belongs to those that cognate; defined here as the capacity to make a choice through the reading and interpretation of signs (a form of politics that belong to a broad range of political actors including us, worms, plants, and crucially computers). These two strands can be easily decomposed, but in advocating for

this division I wish to emphasise the latter (a capacity to make a choice) as the prerequisite of a thing to be political. I argue that (contra Bennett) talking *only* of affect and power diminishes the difference between determinism (things affecting politics) and choice (things crafting, making, generating, and intervening politically).

I aim to situate the three cases in similar ways to “the disruptive power of clowning [that] lies in its capacity to distract and confuse” (Amoore and Hall, 2013, p.99). That is, the ability to offer an ecological approach to the study of malware *which resists the temptation of always tying malware actions back to human intent*. As Bennett (2010, p.108) says, things “have different types and degrees of power... depending on the time, place, and density of the formulation.” I wish to build on this by suggesting that the equivalence that Bennett draws between leaves and worms for instance, under ‘power,’ does not express adequately what is going on. She does reckon that these have “different types and degrees” of power but these are not spelt out. No doubt many objects have great power, but often this power is non-directed, and deterministic, or only has an affectual capacity to influence choice by higher-order interpreters (by, for example, transforming the atmosphere of a room). But political actors, in my rendering, are those that interpret and make choices.

Computational choice, from Hayles’ discussions on cybersemiosis (2018, 2019, Forthcoming) and non-human cognition (2017), is generated through the interpretation of signs through the layers of computation in broader cognitive ecologies with humans. Drawing on this, I propose, articulate, and build two distinct forms of a more-than-human, malware politics, separated by an ability to make choices. Malware make choices based on such things as what variables to search for, accessing the date or time on a computer to trigger a certain action, to interact with the human user, or to choose where to propagate next. These are many different ways in which malware make choices – and these are not the same as we, as humans, make choice that is typically denoted on a liberal, reflexive cognition. This is about making choices that could be as simple as reading an environment and choosing (frequently based on human authorship) on a binary of whether to continue or not. Below I outline both agency and capacity as a way to think through a malware politics and its contribution to ecologies, and subsets to both:

- 1) **Political Agency expressed through choice:** Malware have different gradients of choice *available* according to computational logics and its relation to its author:

- i. **Output:** The author *writes* an instruction for a particular action, which would produce the same or very similar output at each iteration
 - ii. **Framework:** The author *writes* a framework in which environmental and/or random generation is an intended output. The author does not know what the exact output may be.
 - iii. **Surprise:** Malware interacts within certain ecologies and produces *surprise* beyond the author's intentions or knowledge, but it is within the *logical bounds* of computation.
- 2) **Capacity to affect politics:**
- i. **Material power:** The power of something – which can be deterministically-assessed. This is never complete, as the world always exceeds us. Malware as software forms do not necessary correspond to this category as separated from the infrastructures of computation. This involves the transmission of bits, the electricity that supports these infrastructures, and other components that are affected by the capacity of computation's varying materialities.
 - ii. **Disruptive power:** The capacity to affect something. For instance, the presence of malware on a computer, such as ransomware, which can transform how people, organisations, and states respond.

In my division of political agency, logical outputs and frameworks are what an author expects their malware to do, to an extent. In an output, the malware would be expected to follow a particular movement of memory or use a certain protocol the same way each time. In a framework, however there may be some metamorphic encryption or selection of IP addresses for propagation that use random numbers or environmental variables. This could even be a simple 'If, Else' statement which have been logically programmed by the author but for which the output cannot be *deterministically* known due to the ecologies it performs in. As Wendy Hui Kyong Chun (2005) has said – processed through Parikka and Sampson (2009, pp.16–17) that - “even a simple opening and closing of an applications, or rebooting of a system, can make changes to boot sector files, log files, system files, and Windows registry” that is supported by ‘virus researchers.’ This means that even these outputs and frameworks are under constant negotiation and performance, too, meaning that my distinctions are not necessarily steadfast, but a way of thinking through the complexity of *a* malware politics. What I define as ‘surprise’ constitutes a logical choice, but is unintended or unexpected when it was written by its human author – such as interacting in an environment that is beyond logical expectation (e.g. errors and glitches) or producing something new, interacting with a particular actualisation of an ecology that would always remain extended away from its author. These are the three broad forms of choice that I believe are available to malware (and indeed, other software forms) – output, framework, and surprise.

Political agencies then intertwine with the second form of malware politics, which is related to its *affectual capacities*. For example, the feelings generated by the presence of, and

interactions with, endpoint detection engines that alert us to malicious activities. After all, malware can encrypt, steal or covertly monitor us; we call it malicious for a reason. Malware's disruptive power to change emotions, through its impacts on financial lives of those affected by cybercrime, or through the *performative choices* of malware, can shape how states see one another, are all very important. Separating out these two broad forms of malware politics is no easy task – attempting to cleave these apart is only for clarity of argument, and indeed, how choices interact with affectual capacities is difficult to disentangle, but their difference is important. Choices are what make malware *powerful* and able to have such an affective capacity to influence us.

The Cases

Stuxnet, the Dukes, and WannaCry/(Not)Petya are frequently referred to in research on malware, allowing for a pragmatic comparison and contrast with what others have written. They have also reached a critical level of affective influence which made human politics more likely to take an interest in their actions. As Balzacq and Calvelty (2016, p.176) argue, “in recent years, highly publicised cyber-incidents with names like Stuxnet, Flame, or Duqu have solidified the impression among political actors that cyber-incidents are becoming more frequent, more organised and sophisticated, more costly, and altogether more dangerous, which has turned cyber-security into one of the top priorities on security political agendas worldwide.” Though I could instead focus on the multiple *minor* malware that I worked with in the MAL; to make the case for the ecological means reimagining and curating *with the trouble* of those historically pathological interpretations.

Instead of abstracting malware from its environment, as *an object of interest*, I centre my analysis on how malware emerge in distinct ecologies, in contrast to the (auto)ethnographic fieldwork in the MAL. Even though I present three cases, I do not consider them as *events* or *ruptures* as such. To do so would suggest that malware are an object that have a unique role, separate from their ecologies. This is significantly different to Alfred Whitehead's concept of event as being a radical rupture from the past, whereas I would rather think through Bennett's *thing-power* as a disruptor (Roberts, 2014, p.977), with events conditioned by their pasts. So, though my time in the MAL may have been in those 'controlled laboratory'

conditions, this chapter threads how we can apply my pathological becoming in the MAL to the ecological in the ‘wild.’

The majority of information about malware come from reports and blogs of MALs; and we have seen in previous chapters how knowledge may come to be formed, and some of the issues of doing this for an ecological approach to the study of malware. However, I use this as I have little access to events at the time, or to the malware. Reports, and other outputs of mainly endpoint detection businesses, are often produced with both an eye to great technical detail but also to the profile of the commercial imperatives that they serve. Indeed, the reports have a pathological sensibility, that suggests that malware can be analysed and attributed, can be broadly abstracted from their environments, and that intention is something that can be imbibed within them. This requires reading against the grain. In each case I provide a brief overview of their emergence, how they come to be political, and what an ecological reading can provide.

Case 1 | Stuxnet

Stuxnet has been frequently referred to as one of the most complex, multi-modular forms of malware to have ever been created, and has been touted as the world’s first ‘cyberweapon’ (Barnard-Wills and Ashenden, 2012; Kaiser, 2015; Kello, 2017; Langner, 2011; Rid, 2013; Sanger, 2012; Warf, 2015; Zetter, 2014). In this section, I outline how Stuxnet propagated, was ‘found,’ and became an object of international relations on broader discourses on the cyberweapon – through the premise of ‘kinetic’ action. Lindsay (2013, p.368) notes that “most accounts of Stuxnet have focused on its unprecedented technical wizardry rather than evaluation of its strategic consequences.” But I wish to travel in another direction that does not sit in this binary, one that neither attempts to talk of the “unprecedented technical wizardry” nor purely of its “strategic consequences” in the way that they centre on the human, but of the emergence of malware and expressions of political agency. This is not to say that those who created this malware form did not have *intent* behind the creation of Stuxnet; as has been comprehensively explored by Kim Zetter in *Countdown to Zero* (2014). However, I consider and rethink how Stuxnet has been thought of as an object of war or as a mere automation of intent to expose how, even in highly-specialised malware, there is a politics distinct from human authorship.

Stuxnet emerged in the ‘wild’ during 2009 and 2010 and became quickly regarded as a weapon targeted against an Iranian uranium enrichment site, Natanz. Purportedly created by the USA and Israel (as part of a broader campaign called the ‘Olympic Games’ (Zetter, 2014; Sanger, 2012)), it is thought to have been intended to halt, or at least disrupt the production of uranium that was suspected to be enriched for nuclear arms. Stuxnet was first identified by Sergey Ulasen, from the Belarusian anti-virus company, VirusBlokAda, as *Rootkit.TmpHider* on 10 July 2010, initially through the VirusBlokAda (2010) website and a post by Ulasen (2010) on *Wilders Security Forum* on 12 July 2010, followed a day later by more information in a PDF file following requests by forum members (Oleg and Ulasen, 2010). Reading through the initial comments on the forum suggests that, although there was an awareness of its sophistication, there was little idea of what the malware was doing: “we have added a new records to the anti-virus bases we are admitting [sic] a lot of detections of *Rootkit.TmpHider* and *SScope.Rookit.TmpHider.2* all over the world” (VirusBlokAda, 2010). Though later evidence points to Stuxnet being targeted at Natanz, at that time *Rootkit.TmpHider* was a new, unknown malware form.

Stuxnet quick became a common name, with Symantec originally naming it *W32.Temphid*, only later changing to Stuxnet from 19 July 2010, following the convention from other vendors (Shearer, 2010). With ever-greater numbers of individuals and endpoint protection businesses looking at Stuxnet, information quickly grew. For example, it was found that digital signatures from *Realtek Semiconductor Corp.* were used to sign drivers that came from a USB stick (*mrxnet.sys* and *mrxcls.sys* – with those drivers providing its name, Stuxnet). A whole array of other factors came to the fore, including the use of a vulnerability in Microsoft Windows shortcut files (.lnk) which were able to execute automatically as reported by the popular ‘Krebs on Security’ blog run by cybersecurity expert Brian Krebs on the 15th July (Krebs, 2010). This is a vulnerability typically known as a *zero day*. Zero days are vulnerabilities that have not been previously identified, and no Windows computers had any protection to stop the shortcut file executing, as no fix or patch had been written. Research located requests to a Siemens SCADA WinCC + S7 database (Boldewin, 2010), suggesting this was not a ‘typical’ piece of malware. This was quickly dubbed by Frank Boldewin as “industrial espionage or even espionage in the government area” due to its specificity in targeting industrial control systems in an email to *heise Security*, a German security business on 15 July 2010, which was obtained by the website, *the H* (2010).

Symantec produced the most complete overview to Stuxnet in its dossiers by Falliere, O'Murchu, and Chien (2010, (version 1.0), 2011, (version 1.4)), including a paper presented at the Virus Bulletin conference (O'Murchu, 2010) prior to their release. These reports go into deep detail about how Stuxnet operates. I use these details about how Stuxnet operated, to question whether its actions reflect not the intentions of its human authors, but its conditioned potential to do something different, something else. In particular, I question the assertion that it was deliberately limited to stop propagation (Stallings and Brown, 2015). Stuxnet was a highly specialised, and meticulously written piece of software. It was tailored to such an extent that it only searched for two forms of Siemens Programmable Logic Controllers (PLCs) that program Supervisory Control and Data Acquisition (SCADA) systems used in industrial control systems in order to disrupt centrifuges⁷⁶ used for uranium enrichment. Stuxnet searched for S7-315 and S-417 PLCs in order to further execute and had different attack sequences for two S7 databases if it found these types of PLCs. It would not 'do' anything unless these PLCs were found, and thus most of the code was left apparently redundant. It propagated through three main methods (Falliere, O'Murchu and Chien, 2011, pp.25-35):

- 1) Network Shares: This includes such methods as 'Peer to Peer' (P2P) networking⁷⁷ to search *within* networks for the latest versions of Stuxnet, exploiting a vulnerability in WinCC software, utilising network shares, through a Server Message Block (SMB) vulnerability and a 'zero day' printer spooler exploit that had a date of exploitation until 1 June 2011.
- 2) Removable (USB) Drives: Mainly through exploiting a zero-day vulnerability in .LNK files, where it checks for computers with certain conditions to load onto the computer in which the drive has been inserted with earlier versions using the 'autorun' feature.
- 3) Siemens 'Step 7' Project: Stuxnet copies itself to Step 7 projects using three methods that contain, again, certain environmental parameters that were searched for before it copied.

There are different materialities of Stuxnet; not only in propagation but elsewhere. Some of its logical outputs, such as searching for a particular age of a file, or a 'hard' end date such as 1 June 2011 set for the use of the zero day spooler exploit, are those that must be logically followed (unless the computational environment itself has an incorrect date or properties). However others, such as the way it propagates through using removable USB drives mean

⁷⁶ In particular, the technical analysis revealed that Stuxnet set centrifuges back to 1064Hz frequency after manipulating them, which is the speed at which IR-1 centrifuges at Natanz work best (Zetter, 2014, p.245).

⁷⁷ Peer to Peer networking is where computers share material between themselves without the need for a centralised authority.

that, although there may have been checks on the environment, as well as the .LNK zero day deleting itself after infecting three computers, the ability for the author to know *exactly* where Stuxnet went, due to the ecologies of interaction and logical frameworks associated with the use of USB drives, was unclear. Or in a more concrete example of this type of political agency, the use of Windows Management Instrumentation (WMI) enabled sharing across networks - Stuxnet tried to copy itself to all users on that computer and computers on the domain. This was an attempt to encourage the virus to move *laterally* through organisations, but gave Stuxnet an ability to move through ecologies without the direct control of its authors – it was left to the malware to read and interpret signs, and move according to the logical frameworks given to it.

Rebooting Computers, a ‘surprise’

As Kim Zetter says, “mapping digital code to a real-world environment is more an art than a science” (2014, p.175) and in an example of pathological readings, Lindsay (2013, p.387) says with that, “it would be imperative to find and stamp out bugs that could compromise the whole operation (as one eventually did, in fact, by causing an Iranian machine to get stuck in a reboot loop).” The reason Stuxnet was ‘discovered’ came from rebooting computers in Iran, that used the VirusAdBlok anti-virus, meaning Ulasen and others were called to investigate – without this, it is unlikely Stuxnet would have been revealed at that time. Though one could say that this was due to *human error*, I offer a different reading of this event. Stuxnet *chose* as part of computational ecologies of sign-exchange – whether to propagate, install or execute. Signs were read by Stuxnet and by the broader environments of computation interacting with it, which lead to new, surprising *choices* on those signs.

It was initially thought that the rebooting machines had an incompatibility between the malware’s drivers and VirusBlokAda’s anti-virus software. These rebooting machines are a sign of Stuxnet and computation performing in different ecologies that led to the rebooting – not as an intention, but as a result of a particular alignment and choice – one that led to an output *we would not expect*. Indeed, “researchers at Kaspersky Lab later tried to reproduce the problem but got inconsistent results – sometimes the machine crashed, sometimes it didn’t” (Zetter, 2014, p.9). As the UK newspaper, *the Guardian*, reported nearly ten months after Stuxnet had been “captured” by VirusBlokAda:

““These computers were constantly turning off and restarting,” Ulasen told the Guardian. “It was very strange. At first, we thought maybe it was just a problem with the hardware. But when they said that several computers were affected, not just one, we understood that it was a problem with the software the computers were running.”

Ulasen was given remote access to one of the malfunctioning machines, but he soon realised he needed help. He roped in a colleague, Oleg Kupreyev, the firm’s senior analyst, and they spent a week unravelling samples of the computer virus they had “captured” which was affecting the Iranian machines.”

(Hopkins, 2011)

The assertion that a bug, glitch, or error was why Stuxnet was ‘found’ is to ignore the choices of malware in computational ecologies, based on a pathological rendering of computation and malware. Instead of thinking that this *would always* happen, it is better to see this as a product of various components of an ecology coming together at a singular point, performing through sign-exchange and choice. No matter how much *we* test, due to choice, malware will always exceed our capability to conceive its politics. I argue that the performance of Stuxnet, in the environments of computation, through a reading of signs, lead to this – malware’s reading of signs, their choices, are made in unique iterations of ecologies that reveal a politics of malware.

Zero Days

Discussions on cyber weapons, security and war became widespread after Stuxnet. In the aftermath of the attacks against Estonia in 2007, this brought further evidence of ‘destructive’ code and the potential for attacks to *transcend* demarcations between *cyberspace* and the ‘real’ world (Schmitt, 2013). Here we see evidence of the second kind of malware politics, its capacity and power to disrupt. This had been captured in multiple newspaper entries, popular books, and the docufilm, *Zero Days*. In *Zero Days*, vistas of code sweep the screen, where mainly men speak about Stuxnet, one is brought into a dramatised tale about the broader political project of the Olympic Games conducted by the USA and Israel against Iran. This includes shadowy figures from intelligence agencies, classified information, and the murder of Iranian scientists. The aim of the film is to let the “computer code speak for itself” (*Zero Days*, 2016, 5:40 - 05:42), however it still focuses almost exclusively on its authors, and those who ‘discovered’ it.




Figure 12: Theatrical one-sheet for *ZERO DAYS*, a Magnolia Pictures release. Photo courtesy of Magnolia Pictures.

Yet *Zero day's* focus on the humans involved neglects an attempt to understand Stuxnet as a form in itself; and plays into broader pathologies in how publics comprehend malware. Hence, we see the translations of a pathology into popular (re)imaginings of Stuxnet, that in attempting to move away by letting it “speak for itself,” due to the stranglehold on linear intention, becomes re-enrolled into a certain narrative. Through using words such as *infection*, *highly sophisticated hackers*, *bombs*, and other terms that exclude the agency of malware, we are left with malware as some form of exceptional *human* product – rather than a co-constitutive, more-than-human political encounter. This is perhaps due to a need to develop a *plot* that makes sense – pathologies draw neat, linear explanations. This is also present in the discussions from newspapers which quickly turned to discussing the relationship between Iran, Israel and the USA. This meant that the malware, from early reports from MALs, were translated into how it was a ‘cyber weapon’ saturated in the discourses of war – where the intentions of the USA and Israel flow through cyberspace as a non-stratified, Cartesian-like expanse. This is clearly depicted in the poster to the film in the announcement that Stuxnet is part of “World War 3.0” (Figure 12). Yet some interesting alternative, minor narratives emerge in *Zero Days*. One individual, Ralph Lagner, a crucial interlocutor of the translation of how we understand Stuxnet, says that it became too noisy. What is meant by too noisy? This, as part of the choices made by Stuxnet in co-performance with its ecologies produces an *excess* of not only its choices, but also its affective presence on systems, that meant it exceeded the intent of its author, forming through ecologies, new political formations that

allow for new concepts, translations, and ideas around the possibility of malware to be a ‘cyberweapon’ that can only come in a more-than-human collaboration.

Case 2 | The Dukes

In this case, in contrast to my discussion of Stuxnet, I am interested in how malware agencies are stripped, rolled back, and brought under pathologies in order to render them attributable to authors. The Dukes are a malware family that have been connected to an ‘advanced persistent threat’ (APT) group; a classification used to describe malware and human groups that attempt to gain access to targeted individuals, systems, organisations, and networks persistently over time. Thus, how do MALs *tie* together different malware forms and families to a group of human authors: how do they attribute them? I wish to expose how the Dukes had a disruptive power – through the case of the hacking of the US Democratic National Convention (DNC) that was reliant on a more-than-human collaboration with malware choice. This allows for a perspective on the Dukes (otherwise known as CozyBear, CozyDuke, or APT 29) to be seen in a new light in a development of a more-than-human attributional politics. So, how do various analysts, analyses, detections, reports, intelligence agencies and others come to conclusions about ‘who’ wrote the Dukes? How do humans come to assign certain malware to (human) political projects? How and where does the click of a link and a certain internet connection generate a more-than-human political project? What traces, knowledge, and practices lead to discussions of international significance in geopolitics?

One of the core reports on the Dukes written by Artturi Lehtiö, from the Finnish cybersecurity company F-Secure, says;

“the Dukes are a well-resourced, highly dedicated and organized cyberespionage group that we believe has been working for the Russian Federation since at least 2008 to collect intelligence in support of foreign and security policy decision-making (2015, p.3)... [but] are unable to conclusively prove responsibility for any specific country for the Dukes. All of the available evidence however does in our opinion suggest that the group operates on behalf of the Russian Federation” (2015, p.26).

With this, what is the evidence that supports these levels of attribution, and more importantly what connects various malware forms given names such as MiniDuke, CosmicDuke, OnionDuke, CozyDuke, HammerDuke, and CloudDuke?

Kaspersky Labs first coined the name MiniDuke, after the malware was announced by FireEye (Lin, Bennett and Haq, 2013). They suggested that MiniDuke used an exploit similar to that used by similar ItaDuke, and the name reminded them of another threat, Duqu (thought to be connected to Stuxnet in the *Olympic Games*), due to comments in its shellcode from Dante's *Divine Comedy* (Raiu et al., 2013, p.1). Yet, according to Lehtiö, "it is important to note that there is no reason to believe that Duke toolsets themselves are in any way related to the ItaDuke malware, or to Duqu for that matter" (2015, p.4). MiniDuke had 59 unique victims in 23 countries (Raiu et al., 2013; CrySyS Malware Intelligence Team, 2013) and is thus different to the other cases, in that it did not propagate as much as Stuxnet and WannaCry/(Not)Petya. There are few instances of the Dukes – arising from their differing materialities. The Dukes tightly target individuals through spear-phishing⁷⁸ emails and do not use propagation mechanisms, apart from certain OnionDuke variants that used alternative, but equally restrictive distribution methods (Lehtiö, 2015, p.25).

The authors did not want the Dukes to propagate, and therefore limited this, as their intention was *likely* to be stealthy. For example, with CozyDuke, a modular⁷⁹ malware form (F-Secure: a 'toolset'), a victim received a tailored email, containing a link to a compromised file, that when downloaded and opened, would install on the computer. A decoy file also opened, often a PDF, but in one case a video of monkeys (F-Secure, 2015, p.2). MiniDuke meanwhile used Twitter to connect to the author's command and control (C2) server, in order to appear like 'normal' traffic (Tivadar, Balázs and Istrate, 2013). The Dukes demonstrate carefully selected logical outputs and frameworks: demonstrated in thought on how to connect with the malware and the limited amount of samples that have been found. This is how (one could say 'good') malware authors develop a more-than-human collaboration – by sensing and appreciating the limits of computation and malware as well as their capacity to control it.

⁷⁸ Highly specialised emails targeted to individuals. In this case, likely to refer to organisation in which an individual works and seek to make it as 'realistic' as possible.

⁷⁹ This refers to malware that is made of interchangeable components that can be updated and modified by the authors.

Attributing the Dukes

Yet, in attribution, how malware analysts stitch multiple malware forms back to the authors is far more complex. Code sharing is frequently the greatest determinant in the allocation of attribution. As discussed in chapter six, the identification of such code, that would have been presented to me as a ‘gene,’ can point towards shared authorship. However other contextual factors such as language, time zones, and who it targets (potential intent) can also be used as evidence. For example, PinchDuke had the error message, Ошибка названия модуля! Название секции данных должно быть 4 байта [Module Name Error! The name of the data section must be 4 bytes], suggesting potential Russian-language authors (Lehtiö, 2015b, p.26). Code sharing can also imply similar development environments (where malware are written, the ecologies of their production). For example, OnionDuke and CozyDuke share a particular component, *show.dll*, that has the same export tables, internally called “UserCache.dll”, tying these malware forms to potentially one place of production.

Consequently, the environments and ecologies from which malware emerge are not some innocent bystander, a ‘clean’ or neutral container, but are teeming with conditions of emergence and become present in traces that are often picked up in attribution. However, making these connections entail making certain assumptions. For example, that these malware forms i) can only exist together in certain spaces, ii) that they cannot be deliberately mimicked or manufactured to ‘look like’ they are produced by another human actor, and important for my argument iii) that malware could not produce something *unexpected*. The performativity of software, where code becomes in different environments and ecologies, often absorbing components of a particular computing environment it passes through in order to operate, is often forgotten, making it difficult to simply trace malware back to its human author.

Delving further into these attributional politics, contextual factors, some of which I detailed in the MAL are used by governments and ‘threat research’ businesses to align different malware targets (only those where they are ‘found,’ remember) and whether they fit into a ‘strategic’ cycle. This is used in attribution of the Dukes – such as the targets, language used, time zones where they are compiled, and so on. Though this may get us close, as malware follow the logics of computation, they are not its extent. Other evidence, such as watching

human authors can be far more compelling, as happened with the Joint Sigint Cyber Unit of the Dutch Intelligence Agency (AIVD) and the Dutch Military Intelligence and Security Service (MIVD). As two newspapers reported in January 2018, *de Volkskrant* (Modderkolk, 2018) and *Nieuwsuur* (Bosch van Rosenthal, 2018), the Dutch intelligence services had access to the Russians who wrote the Dukes (CozyBear in the articles) who “are in a space in a university building near the Red Square. The group’s composition varies, usually about ten people are active. The entrance is in a curved hallway. A security camera records who enters and who exits the room” (Modderkolk, 2018). This, if true, is a fundamentally different form of attribution, as it is an observation. Rather than relying on more-than-human collaborations, that are infused with choices and performativities of ecologies, observation bypasses the tenuous links to human authorship. This is why tying code together through code structures, whether timestamps or timings of attacks conform to a particular time zone, or similar environments, will never *reveal* with confidence the malware authors. Yet it does not make it impossible either, *as computation is logical*, we can observe constructs for outputs and frameworks, meaning some tracing can be made, but not necessarily about the connection between intent and impact. Malware are more-than-human collaborations that do not lead to easily attributable authors due to the choices made by malware (and human obfuscation) – and ones that authors can work with to generate this uncertainty.

Democratic Politics

The US DNC hack, as it became known, was reportedly conducted by two Russian groups, the Dukes (CozyBear) and FancyBear that targeted the US Democratic Party ahead of the 2016 Presidential election. A DNC lawsuit shows that the DNC were unaware that many of their systems had been compromised until at least July 27 2015 (Democratic National Committee v. Russian Federation, et al., p.23), with user credentials exfiltrated, permitting entry to its systems. The DNC invited CrowdStrike, the cybersecurity business, to assist with preventing the attack after it was discovered on 28 April 2016 (Democratic National Committee v. Russian Federation, et al., p.22), several months before the election. This led to the DNC decommissioning 140 servers and rebuilding 11 servers and 180 computers. However, when thinking ecologically, through more-than-human collaborations, the damage did not come *directly* from the malware. It was the information extracted and its subsequent publication by the whistleblowing website *Wikileaks*, that purportedly damaged the

Democratic nominee, Hillary Clinton, by releasing politically-embarrassing information (Nakashima and Harris, 2018).

The particular form used, SeaDuke (Crowdstrike: SeaDaddy) (Alperovitch, 2016), was described by Symantec as the “latest weapon in the Duke armoury,” a “low-profile information-stealing Trojan [that] is used only against high-value targets” (2015). They continue to say that:

While the Duke group began to distribute Cozyduke in an increasingly aggressive manner, Seaduke installations were reserved only for select targets. Seaduke victims are generally first infected with Cozyduke and, if the computer appears to be a target of interest, the operators will install Seaduke.”

(Symantec, 2015)

The high configurability that was found, and cross-functionality across Windows, Mac and Linux, show the intricate work of its authors. I see a dual operation at work; between more-than-human collaborations clearly occurring between the materialities of CozyDuke being a broader ‘tool’ in which to determine (with intent of its authors) whether to install SeaDuke for greater stealth; and a second in an attributional negotiation that occurs between MALs, threat intelligence business, and intelligence agencies to trace back who has written malware. Even with a tightly logically-bounded malware an attributional politics can still be unclear. There is *never* going to be a perfect alignment, and even divergence, between intent, logics, and performance of malware. Malware may be highly-targeted and specialised by their authors, but more-than-human collaborations ensure that there is no easy track *from the malware alone* to their human authors for attribution in cybersecurity. An ecological reading of the Dukes then sees them as more-than-human collaborations, which make attribution difficult, and that the affective, disruptive power of malware’s choices can disturb and affect ‘human’ politics itself. That is, without malware choice, we would have unlikely have seen the 2016 US Presidential election play out in the same way. By this, the authors leant on the more-than-human choices of malware in order to disrupt and see human politics in a new light, which the Dukes made possible.

Case 3 | WannaCry/(Not)Petya

With WannaCry and (Not)Petya, I move beyond attribution to highlight *how* malware became translated in the short-term, high-impact events of 2017 through my experience at Sophos to expose how these events generated confusion, mistakes, and revelations that show that there are logical outputs and frameworks, as well as disruptive power. The spatio-temporality of analysis, comprehension, knowledge, detection, and politics of the MAL worked to eradicate choice, to render malware pathological in order to detect. I connect them together to explore how the ‘wild’ interacted with my experience of the MAL, as well as the frameworks that came from the release of exploit ‘tools’ by the *ShadowBrokers* (thought to be hacking tools of the USA’s National Security Agency, which were subsequently used for *effective* self-propagation⁸⁰). In my first days in the MAL, the ShadowBrokers made an appearance through an earlier dump of ‘tools,’ with Terry announcing over Jabber that “we’ve gotta catch ‘em all” in a clear reference to the then popular smartphone game, *Pokémon Go* (Research Diary, 13 January 2017). However, it was not until the fifth leak from ShadowBrokers on 14 April 2017, where the EternalBlue and EternalRomance exploits were released⁸¹, that their connection with WannaCry and (Not)Petya became important. Yet, even with the release of such frameworks, Daniel said that detecting them was not at the “core of what we [SophosLabs] do” (Research Diary, 19 April); writing detections for these tools was a research project rather than one based on primary task of detecting malware.

WannaCry

Ransomware has become popular for cyber criminals (AV-Test, 2018, p.3) due to the relative ease of generating value by encrypting files and then offering to decrypt these files for a fee. An (apparent) early hash⁸² of WannaCry came from a sample submitted to the malware repository, Virus Total, on 10 February 2017, with a timestamp of 9 February 2017 (Balaban, 2017; Shevchenko and Nish, 2017), which had a *.wry* extension. These extensions provide the name added to files after they have been encrypted by a ransomware. So, in this case a file

⁸⁰ Incidentally this was not a ‘zero day’ vulnerability and had been patched by Microsoft in Windows systems before their release in the ‘wild’ (BetaFred, 2017) on 14 March 2017, ahead of WannaCry’s main impact on 12 May 2017.

⁸¹ These both exploit a vulnerability in Windows Server Message Block 1.0 (SMBv1) that is a network file sharing protocol.

⁸² SHA256: 3e6de9e2baacf930949647c399818e7a2caea2626df6a468407854aaa515eed9; MD5: 9c7c7149387a1c79679a87dd1ba755bc



Figure 13: Left: 10 February 2017 sample (MD5: 9c7c7149387a1c79679a87dd1ba755bc). Screenshot taken from Twitter (S!Ri, 2017). Right: The updated decryptor (ransom note) screen and the change from ‘WCry’ to ‘WannaCry’ from 27 March 2017. Screenshot from Twitter (Malware Hunter Team, 2017).

‘andrew_dwyer.docx’ would then become ‘andrew_dwyer.wry’ and I would not be able to open that file without the decryption key. These early versions predate the release of the ShadowBrokers tools and already used network shares to propagate, with a unique section very similar to ‘Contopee’ backdoor linked to the Lazarus group – a North Korean state hacking group (Shevchenko and Nish, 2017). The North Korea connection was highlighted by a Google researcher in a tweet after the main WannaCry release (Mehta, 2017).

By 27 March 2017, a new decryptor user interface emerges similar to one that affected computers around the world, where the name it uses for itself, changing from ‘WCry’ to ‘WannaCry’ (See Figure 13). However, as was said in an interview;

“for example, if you check some of the claims of some of the AV companies making concerning attribution, you would see they are referring to a February WannaCry sample, but again that’s just a naming convention. This February WannaCry variant, at least from my perspective, had very few elements in common in terms of infection mechanism with the ones we saw in May.”

(ENISA, 18 October 2017)

This demonstrates the contested histories of malware; however, I think there is something to be traced here in the development of the ransomware *prior* to the ShadowBrokers dump, while noting the significance of this dump in shaping the potential for WannaCry to cause disruption *perhaps* beyond the intention of its authors.

Incidentally, Friday 12 May 2017 was the only day I ‘took-off’ when WannaCry ‘hit,’ and I was on a Eurostar train from London to Brussels for a long weekend. In an interview at the UK National Cyber Security Centre (19 October 2017) they told me how they had sought

‘patient zero,’ in its pathological resonance, that had been traced to *Telefónica*, the telecommunications business, in Argentina through DNS⁸³ connections. From Argentina, WannaCry quickly propagated to organisations such as Deutsche Bahn and the UK’s National Health Service (NHS). As I wrote in my research diary (15 May 2017), there were thought to be around 200,000 *infections* in around 100 countries. An independent audit written by William Smart (2018, p.10) for the UK Department for Health and Social Care (DHSC) said, that eighty out of 236 NHS trusts in England (thirty-four directly and forty-six reporting disruption), including 595 GP practices were affected (Morse, 2017, p.6). This led to a total cost for the NHS of £92 million (Department for Health and Social Care, 2018). Disruption at several hospitals lasted several days, and even weeks after.

When I returned to the MAL, there was a condensation of knowledge on the internal ‘wiki’ pages. There was concern that the exploit could be used again (Research Diary, 16 May 2017). However, by 17 May there were at least three variants of the ransom note, with some looking similar to the (in)famous CryptoLocker ransomware, with various extensions including ‘.wry’ and ‘.wncry’ making the classification of each malware form difficult to assess. By 19 May, “there was a notification on the hub for all Sophos employees to look at, which was an overview of the WannaCry ransomware.” This was a condensation of all the varied initial suggestions by analysts and unexplored options, into a neat, well-structured and ‘clean’ PowerPoint presentation that excluded these unanswered questions; such as that ‘patient zero’ or the initial infection vector (how it had started), as this was yet unknown.

WannaCry can be seen as both an intended consequence of the authors (giving logical frameworks to the malware), allowing it to perform, extended away from the human, whilst displaying the disruptive power of this malware in its ecology, making choices, in ways unknown and incomprehensible. *The authors may not have known how disruptive their malware was going to be, and this is the surprise – as in it followed the logics of the framework of its author – of ecologies. They may have known it was likely to propagate, but to what extent would be unknown and dependent on the performance of different ecological elements.* For example, even though the authors set up a framework for WannaCry, would they have known the NHS would have been

⁸³ Domain Name System – as WannaCry attempted to connect to a URL (that would later be identified as a ‘kill switch’) it is possible to trace the first connections to this URL. More information on these URLs and tracking them by the UK’s Top-Level DNS is available (Nominet, 2019).

severely impacted? Very unlikely. This was not the first malware to use the EternalBlue exploit, with at least the Adyluzz cryptominer⁸⁴ using this and closing the vulnerability, that likely led to fewer computers *being* vulnerable to WannaCry (Barak, 2017; Dwyer, 2018). As Whalen (2017) notes in a blog post, there were only three anti-virus engines that detected WannaCry through heuristic or generic detection before 12 May 2017 (i.e. no endpoint detection engine had a signature detection). There had been previous identification of WannaCry on systems in February, but due to its low volume, no particular detection had been written as it was not commercially viable for different MALs to do so – these early forms of WannaCry did not have the (in)famous worming component. This links to the discussion in the previous chapter around curation, on what work to pursue next, where to focus attention in the commercial logics of MALs. Yet, after 12 May 2017, most engines detected WannaCry as ransomware, with some incorrectly doing so as another ransomware, *Locky*.

The Varying Geographies of WannaCry

Over time, it emerged that WannaCry had interesting geographies based on the EternalBlue exploit. I asked Nathan what he had been doing in the days after WannaCry. He appeared frustrated, and increasingly resolute in the face of not getting WannaCry to ‘work’ on Windows XP (Research Diary, 31 May 2017). XP kept crashing, due to an implementation error in the SMBv1 protocol it used to propagate, meaning that WannaCry, in many instances, could not affect computers. As Nathan said in a later interview (21 July 2017), “I think it was just how the SMB[v1] was assembled, it was wrong, that is all I can think of.” This observation strikes similar to Greenhough’s (2012) common cold scientists and their unruly environments in attempting to grow viruses in laboratory conditions – malware, like biological viruses, *make* choices in broader computational ecologies that make them ‘unruly’ to us. It was logical, but seemed to be doing things that escaped the explanation of Nathan.

As Costin Raiu (2017), from Kaspersky Labs posted on Twitter on 19 May 2017 (see Figure 14 for a perspective on affected systems), the vast majority of infections did not come from

⁸⁴ A cryptominer is a piece of software that uses computing processing power to ‘mine’ for cryptocurrencies in order to generate more currency and therefore capital.

Windows XP, but the newer operating system, Windows 7. As research from Kryptos Logic⁸⁵ (2017) showed, on Windows XP, the EternalBlue exploit was not able to execute properly, meaning it could not spread through this mechanism. However, another *tool* in the ShadowBrokers dump, DoublePulsar, a backdoor that runs in the computer’s kernel mode (which grants high-level permission to take actions on a computer), was used by WannaCry as a propagation method in addition to the EternalBlue SMB exploit, so it did not mean XP was completely unaffected, but not by the propagating framework of the EternalBlue exploit. This DoublePulsar backdoor was reported to have several tens of thousands of infections even prior to WannaCry on 12 May (Goodin, 2017). As Kryptos Logic (2017, emphasis added) say, “As ETERNALBLUE exploits a heap overflow, *successful exploitation is nondeterministic and requires multiple attempts*. WannaCry gives up on an IP after 5 attempts.” Whilst it is unnecessary to consider what a heap overflow⁸⁶ is, even when the author incorporates the exploit in their software – the performative aspects of computation meant it was not deterministic. This challenges much of the discourse at the time that claimed Windows XP was the reason for the NHS’s problems, due to their older IT estate, even if this was not the case (Morse, 2017).




Figure 14: Graph showing WannaCry infection rates across differing operating systems, photo from Twitter on 19 May 2017 (Raiu, 2017).

⁸⁵ Kryptos Logic is an employer of ‘MalwareTech’, Marcus Hutchins who discovered the ‘kill switch’ in WannaCry that slowed the spread of the malware. He is currently under investigation by US prosecutors over the accusation he created a credential harvester, *Kronos*, among other malware forms, for more information see (Newman, 2018).

⁸⁶ This is part of a broader set of buffer overflows, which can be used to overwrite memory in areas used for dynamic allocation (i.e. used when a program is running)

There was much anxiety in the media about how many computers would be affected by WannaCry and whether it could be stopped. One researcher, Marcus Hutchins, found a ‘kill switch’ (that stops the ransomware component encrypting files). Hutchins claims he did not know what this was, but found WannaCry connecting to a website which was unregistered, and bought this domain for \$10.69 (Newman, 2017). This dramatically slowed down WannaCry encryption, as if the domain sent back a response, WannaCry would not encrypt and would shut-down. *However*, this did not mean that the WannaCry worm component had stopped and it is still propagating to ‘unpatched systems’ – the kill switch only prevented the ransomware component encrypting files, meaning it is still propagating, interacting, but is being stopped by a *sinkhole*⁸⁷ domain at the time of writing this thesis. This is what a politics of logical frameworks could be described as, where the author knew what they have logically programmed and this had extended away from them, but also in the sense of a disruptive power, in that the choices of this more-than-human collaboration in computational ecologies caused huge disruption (even though it only earned around £105,000 for its authors in ransom payments (Lee, 2017)). Hence, seeing malware as some form of singular object carrying maliciousness around is critical to challenge – as one may not see WannaCry without the ransomware component as particularly harmful, but the worm component is still active and can still infect those systems that are vulnerable. It was the coming together of different components into the WannaCry form at a particular performative moment, imbricated in certain ecologies, that the disruptive power of malware was enabled in a more-than-human collaboration through logical frameworks.

With this case, I highlight something different in comparison to Stuxnet and the Dukes; that ecologies, computing environments, also ‘speak’ back to malware and limit those choices and political impacts. WannaCry was reliant on choice, but it is also the ecologies themselves that made it a global phenomenon. Yet, this is not how it became known; as WannaCry was quickly enrolled into a human-dominant reading of politics. Though quick attribution was made to the ‘Lazarus Group’ (North Korea), in speaking with the UK’s National Cyber Security Centre, they stated that their attribution standard had to be far greater than private enterprise (that relied on code sharing (GReAT, 2017)) – ‘they have to *see*’ (Interview, 19 October 2017). WannaCry became intertwined in the 2017 UK General Election over the

⁸⁷ A sinkhole domain is a domain that gives out false information, and can be used to monitor malware connections, but can also be used to limit the impacts of DDoS attacks against websites.

vexed issue of NHS funding. As the *Observer* headline – “Cyber-attack sparks bitter political row over NHS spending” (Doward and Townsend, 2017) – demonstrated, funding became synonymous with the ecologies in which WannaCry performed, and one that was based on older XP operating systems. As I have attempted to argue throughout, it is not possible to separate malware from its environment, and what WannaCry demonstrates is how *essential* ecologies are to both (de)establishing propagation and interaction but also how they come into contact with human politics, with *mistranslations* about Windows XP and whether further funding would have even eased the NHS’s exposure.

(Not)Petya

After WannaCry, the MAL started returning to normal (Research Diary, 7 June 2017) with preparations, in a meeting of all members of the MAL, to improve reporting mechanisms in response to large-scale attacks, with wiki pages in the future to be used to collate all information (Research Diary, 21 June 2017). It would only be days later that this new procedure would be required. I was sitting at my computer after lunch, thoroughly looking forward to the end of that Tuesday, dealing with the stupor of a late night with friends. I had a few minor tasks to complete. At around 2.15pm on 27 June 2017, the lab suddenly changed however:

“Adam ran into the room, speaking very fast, out of breath to Daniel. Daniel at first just asked him to slow down. Out of utter boredom I just listened in. He was saying there was some new ransomware variant, but that it was affecting a major company very bad and was spreading like WannaCry, through its network. Daniel tried to brush him off saying did we have a SHA⁸⁸ [a unique cryptographic identifier]? However, Adam explained a bit more what was going on which evidently caught Daniel’s attention. Adam left and then Daniel said, ‘who’s on Frontline⁸⁹?’ Charlie responded saying he was... The tempo and rhythm of the lab markedly changed, I could feel my body quicken, coming out of its slumber, eager for information about what this may be.”

I searched Twitter for a sample SHA using ‘#ransomware,’ but Kai, another malware analyst, had got there before me and had forwarded it to Daniel. When he received this, and with increasing chatter among several of us on Twitter, Daniel asked how many of us were free to look into the new ‘ransomware variant.’

⁸⁸ A ‘SHA’ is short for Secure Hash Algorithm, typically a SHA-256 (32 bytes) that uniquely identified software and are used as a common identifier across MALs.

⁸⁹ Frontline was the analyst dedicated to responded to emerging incidents and issues.

“The entire dynamic of the lab was now feverish, gone were the quiet moments, the sipping tea; now everyone was fixed on the screen. I looked at the sample on my screen which I was unpacking, and it came out with two DLLs, one which was a 32-bit and another 64-bit. I looked into this further, submitting the samples to [Sophos’ automated sandbox, Sandstorm]. At this point, an email came in from Daniel saying that we were to put anything we find onto a Labs wiki page (from the lessons learnt from WannaCry)... Daniel shouted [across] to Nathan to put in a generic⁹⁰ [signature detection] for the file that had come in to ensure there was coverage. Daniel was running around coordinating a response as news flooded in... he ran to those in system development and quality assurance (also in the MAL) to expediate the testing and get out a generic as a special release as quickly as possible.”

In this fast-paced environment, pathology took over with even greater vigour; dissecting, observing, containing, detecting. I found several *pipes* linked to credential extraction and a GUI (Graphic User Interface), which I quickly assumed must be some form of graphic for a ransom note (see Figure 15). However, the calls to *shutdown* were something that struck me as odd. What was this human-directed logical output doing in the malware? However, at that point, this was a distraction as I was searching for things to put into a detection as *we already knew it was malicious*.

Widespread reporting suggested the form was some variant of the Petya ransomware, Daniel called out across the lab and asked if anyone was against calling this malware Petya; there was a grumble of agreement with its comparison to Petya due to its *visual* similarity in the ransom note. Several (Not)Petya signature detections were released that day (Sophos Knowledge Base, 2017). It was only later that we would realise that the ‘odd’ things indicated the malware was doing something very different. This shows how the expectation of what ecologies *should* be, as delineated by the pathological logics of the MAL (including myself), led to the analysts ignoring the choices made by malware; which routines were executed, what it did in ‘real’ environments. Because we believed it was ransomware as it *behaved and looked like what we expected ransomware to be*.

⁹⁰ A type of (static) signature detection particular to a malware form – i.e. tailored to (Not)Petya.

```

Oops, your important files are encrypted.

If you see this text, then your files are no longer accessible, because they
have been encrypted. Perhaps you are busy looking for a way to recover your
files, but don't waste your time. Nobody can recover your files without our
decryption service.

We guarantee that you can recover all your files safely and easily. All you
need to do is submit the payment and purchase the decryption key.

Please follow the instructions:

1. Send $300 worth of Bitcoin to following address:

1Mz7153HMuxXTuR2R1t78MGsdzaAtNbBWx

2. Send your Bitcoin wallet ID and personal installation key to e-mail
w0wsm1th123456@posteo.net. Your personal installation key:

zRNagE-CDBMfc-pD5Ai4-vFd5d2-14mhs5-d7UCzb-RYjq3E-ANg8rK-49XFX2-Ed2R5A

If you already purchased your key, please enter it below.
Key: _

```

Figure 15: The (Not)Petya ransom note. Taken from Sood and Hurley (2017).

Ecologies of Interaction

Kaspersky Labs, the day after the attacks, said that rather than being a conventional ransomware form it was in fact a ‘wiper’ (Ivanov and Mamedov, 2017). This was due to its modification of a piece of crucial code that told the computer where the operating system was located so that it could ‘boot’ – the Master Boot Record (MBR). Wiping could occur be either through corrupting or overwriting this (Hung and Joven, 2017). With all ‘normal’ ransomware, there is a recording of the unique encryption key in order for its authors to be able to decrypt the files for that computer. However, due to how its authors had written and implemented the cryptosystem, (Not)Petya produced a randomly generated string so that there was no ‘key’ for decryption – meaning it had no logical framework or output to decrypt its own encryption. Outputs and frameworks allow the identification of a likelihood that authors *intend* to do something. The lack of these to decrypt files – a deliberate omission – is one tightly bound to its author.

Its method of propagation borrowed heavily from the ShadowBrokers dump, in the same way as WannaCry. It used the EternalBlue SMB exploit and Microsoft had ‘patched’ the EternalBlue vulnerability in the Windows operating systems through its MS17-010 security

bulletin back on 14 March 2017. However in the light of WannaCry, many organisations and individuals had been advised to install the patch to be ‘inoculated’ from the vulnerability (Windows Defender Research, 2017). Yet (Not)Petya had other propagation techniques at its disposal including the collection of credentials (users and passwords) in addition to all the IP addresses on a network. It used this information to propagate using network shares, meaning that even if systems were patched with MS17-010, (Not)Petya could still potentially propagate depending on the ecologies it encountered.

What I wish to draw out is how (Not)Petya’s materialities were constructed differently, changing the way it could perform. Unlike WannaCry, (Not)Petya’s EternalBlue exploit exhibited a crucial distinction – that it was not designed to propagate via the internet, but only through local shares – meaning (Not)Petya was less likely to spread outside of local networks, limiting its speed of external propagation compared to WannaCry. Hence, we encounter a crucial distinction in how the same exploit can be written by authors differently; leading to materially alternative malware geographies. Though both WannaCry and (Not)Petya were developed in a more-than-human collaboration; the possibilities for rendering *choice* were limited in (Not)Petya in ways commensurate to logical frameworks.

(Not)Petya was initially spread through the accounting software, MEDoc, in what is called a ‘supply-chain’ attack (Warren and Hutchinson, 2000). The software, widely used in Ukraine, was first reported by the Ukrainian ‘cyberpolice’ on Facebook and Twitter (Cyberpolice Ukraine, 2017) where they state this was the most likely source of (Not)Petya. This demonstrates a particular geography and ecology of interaction, where different elements came together to introduce malware onto other networks and places; where the logics of computation are utilised for their *both* their logics (and likely choices) and disruptive power. One business, AP Møller-Mærsk, became seriously affected as captured in a fascinating piece of journalism by Andy Greenberg (2018). (Not)Petya affected all parts of the business, costing it in the region of US\$250-300million. As its Chairman, Jim Hagemann Snabe, said at the 2018 World Economic Forum (colloquially known as Davos), “we had to reinstall an entire infrastructure, we had to install 4,000 new servers, 45,000 new PCs, 2,500 applications and that was done in a heroic effort over 10 days” (Securing a Common Future in Cyberspace, 2018, 4:05 – 4:21). Though the choices of (Not)Petya were limited by its authors, the ecologies

of Mærsk led to it spreading rapidly due to limited network segregation; leading to it nearly losing its entire IT infrastructure.

However (Not)Petya had a relatively *limited* impact compared to WannaCry, so why did an individual from the UK NCSC conclude that (Not)Petya was much more “dangerous” (Interview, 19 October 2017)? I suggest this is due to the differential more-than-human collaboration with ecologies, in which MEDoc software was part. MEDoc is accounting software used widely in Ukraine (around 80% of all businesses (Stubbs and Polityuk, 2017)), but was the target of an attack where credentials of an administrator had been stolen. This granted them access to servers in order to introduce (Not)Petya through an update of the software (see Maynor *et al.* (2017) for more on how this happened). If its author intended to target Ukrainian businesses, this would make sense. Using the logical frameworks, along with malware choice, it *was unlikely to propagate too far but embrace ecological alignments within businesses to achieve their goal – with the impact unknown in its entirety*. The new spatialities of computation, greater interconnectedness, IoT devices, automation, and so on, pose new questions around how malware propagate and what choices are available to authors – and whether intent can be strongly tied to these logical outputs and frameworks themselves. With (Not)Petya the authors bounded their malware. They intended it to spread – but not via the internet – which meant the number of computational devices that could be affected at speed was reduced. Yet, the authors also did not *know* where exactly (Not)Petya would end up – as it still had to make choices, to which networks it was to spread internally, and would not know how companies affected would be interconnected. Potentially this was the aim, to have a severe impact, but not one as great as WannaCry but more targeted to those who did business in Ukraine.

Malware Materialities

A malware politics is different to pathological contextualisation; it treats malware as coming from somewhere, making choices, and is constructed within and limited by its emergent ecologies. This emphasis on ecologies allows for a generative use of new materialisms, recent work on political ecology, and the work of N. Katherine Hayles on the role of cognition in the formation of political actors. As Stiegler (2017, p.139) argues in a discussion of Erich Hörl’s *General Ecology*, that “it must as such articulate the questions of *selection* and of

decision.” Hence, how we select and decide what software is malicious, how malware makes choices, and what it includes and excludes, are essential questions. Does a software form that, in one ecology appears malicious and, in another ecology as ‘clean,’ mean that a piece of software is *always* malicious? Or is it that the selections and decisions (and choices) of a multitude of more-than-humans in ecologies that lead to particular performances of maliciousness really based on the intent of their authors? This takes us to a cognitive and thus political ecology, where choices are mediated through various more-than-human signs; computational and human. Identifying where human and malware choices begin and end is difficult to distinguish and are specific and often exceptionally nuanced, making it all the more important to understand – indeed where political agency begins and ends. Assessing malware in an ecology is then both a scientific and political question, and these are indivisible. However, this exists on two levels; one based on selection and decisions of humans, as no doubt we play an important part in our own politics, but also the *choices* of malware and their disruptive capacity – which intertwine to inform a more-than-human, malware, politics.

As George Osborne, the former UK Chancellor of the Exchequer said at the UK Government Communications Headquarters (GCHQ) on 16 November 2015:

“We cannot create a hermetic seal around the country... but with the right systems and tools, our private internet services providers could kick out a high proportion of the malware in the UK internet.”

(George Osborne, 2015)

This explicit recognition of there being no possible “hermetic seal” shows how malware exceeds the control of the state, and not just because of humans. In this way I think it useful to put Grosz’s chaos into a discussion with cybersecurity and cyberspace where “chaos here may be understood not as absolute disorder but rather as a plethora of orders, forms, wills – forces that cannot be distinguished or differentiated from each other, both matter and its conditions for being otherwise, both the actual and the virtual indistinguishably.” Chaos is not pure disorder but ecological; which may seem to exhibit no structures, but which is full of logics that give it a semblance of such through complexity. A balance is thus required between understanding *how* ecology conditions malware, and how malware’s choices, as part of broader computation, condition the ecologies in which it performs those choices.

A turn to thinking of malware as material is crucial, in what its possible logical avenues are. I have demonstrated through the three cases in this chapter – how they are constructed by humans is crucial, but this is not the entire story. The materiality of the code enable certain potentialities for malware agency and capability. However, materiality has come to mean something different for many – one which causes ‘physical’ effects. Whereas I am less inclined to see a distinction in the binary between the ‘non-physical’ and ‘physical’ – it becomes of great importance when understanding whether ‘harm’ has been caused in international law and how this is re-enrolled in inter-state actions (Schmitt, 2013). I do not wish to engage in the debates on what ‘harm’ means in international law or its ethics (see Taddeo, 2016 for this discussion) – but to instead offer to the debate *what* the ‘physical’ and materiality means for a malware politics, and the choices and ecologies in which this emerges.

Unlike Stuxnet, and more recent attacks against the Ukrainian power grid (E-ISAC, 2016), there have been few malware that have had perceived direct, physical effects. Though this distinction may be superficial, it is not to discount the work of those who explain the *effects* of malware; that can have serious and dangerous effects that could cause certain forms of harm. I do not see malware as distinct in whether it controls a motor, or the movement of other forms of *hardware*; there are particular materialities no doubt, but it is one of degree rather than binary separation. However, it does affect how people relate to malware and whether they morph into a particular politics. For example, with (Not)Petya, does wiping a computer not constitute a ‘physical’ destruction of a computer? In thinking of the interpretation of signs through the layers of computation, the electronic depiction of bits, are all physical – otherwise cosmic rays would not be able to flip bits as an expression of the material power of ecologies. Rather than thinking of malware along its variants of physical/non-physical, I propose considering malware in its *cyberspatial* becoming. Malware performs electronic movements but makes choices not on this distinction but on how it reads signs and chooses based on this through the ecologies in which it is situated and crafts.

Movements of electrons, to the manipulation of values on an Iranian centrifuge is one of scale – not of fundamental difference. These scales are important, but we should discard notions of ‘cyber-physical’ action (as this category is often called, see Cardenas, Amin and Sastry, 2008; Wolf, 2009; Lee and Sokolsky, 2010) as if this binary already exists, and question to what scale the effects of malware such as Stuxnet are. The ecologies in which




Figure 16: Liam O'Murchu demonstrates how a Programmable Logic Controller (PLC) can be manipulated to create 'unexpected' action through the popping of a balloon. Screenshot from YouTube video (Perez, 2014).

malware exist are far more important to understand the affectual capacities of malware, and how choices made, through more-than-human collaborations, that always have effects that vary in intensity across a variety of scales. I think the distinction is retained to preserve intentionality as a linear extension of the human author; if we accept a radical new understanding of malware as exhibiting choice – what happens if a human is killed or other serious consequences that are not the clear intent of its author(s)? If, in detailing the potential politics of malicious presence, if there is a uranium enrichment plant controlled by computers, or if an NHS MRI machine is unpatched from the SMB exploit used by WannaCry and (Not)Petya, then this will have important effects on what we humans *immediately* recognise. But recognising malware agency is about a reconsideration of our entire relation to intention, space, and security.

Liam O'Murchu a Symantec researcher, in Figure 16, changes code on a Programmable Logic Controller (PLC) like that targeted by Stuxnet, where the settings are set so the balloon expands and pops. These atmospheric demonstrations, in similar ways to McCormack's (2018) description of how balloons generate a sense of affects, allow for the visibility of code to somehow state that Stuxnet is both quantitatively and qualitatively different from previous malware. This *performative* gesture of a 'physical' attack which pops a balloon may be intended to show the (potential) destructive capacity of malware; but for me is a demonstration of how ecologies of malware have to come together to achieve such an act. Yes, the *intent* of the author in developing logical bounds is required, but also on the PLC,

on the elasticity of the balloon's material, and ensuring (on the author's behalf) that the malware *interprets signs*, makes *appropriate* choices, and imbricated with its ecologies, the balloon pops.

I therefore challenge the central role given to Stuxnet and other 'physical' malware forms; that there has to be an ecology of interaction, and a reading and performance. As Balzacq and Calvelty (2016, p.195) say with regard to Stuxnet, "authorship for the malware has never been confirmed by any political actor... it is the basis for a widely accepted version of 'the truth' to emerge, which has changed the cyber-security discourse fundamentally." The ecology of Stuxnet's interactions then mixed with our collective pathological notions of malware, in order to establish this piece of malware as uniquely 'physical.' This, in turn, comes with a growing awareness of the changes in computational capacity and spatiality for malware to express greater choices and therefore exhibit greater affectual power. This is not to say that anyone involved saw it this way; but without this attentiveness we are left with an underdeveloped view on malware that instead prioritises 'physical' forms of malware over others. The different forms of political agencies – whether outputs, frameworks, or surprises – and affectual capacities I have outlined in this chapter are frequently hard to disentangle and I have attempted to provide some insights at these different gradients of agency for malware. Yet, as I turn to in the concluding chapter, the eco-logical is already with us – and I have opened up questions on who and what produces (cyber)security when malware performs in a more-than-human collaboration with its authors.

Chapter Eight || Ecological Conclusions

Malware have become the nouveau threat of contemporary societies. They are infused with pathologies that are both part of their lineages and how they become practiced and reiterated in, through, and beyond the spaces of the MAL. The diffusion of computation in the home, workplaces, ‘critical national infrastructures,’ and militaries has been generative of a collective vulnerability (Stevens, 2015; Barnard-Wills and Ashenden, 2012), with the internet becoming central to this unease in its radical decentralisation (Graham, 2013). In this context, I believe an eco-logic is required more than ever to offer a greater appreciation of more-than-human computational agency that challenges and (re)constitutes ‘human’ politics. In this concluding chapter I do not set out a singular practice of ecology, as I hope to have demonstrated, this would be an impossible, even irresponsible, act. However, it is to offer directions, new sketches, and mappings in the way that Gerlach and Jellis (2015) argue – each iteration, each engagement with malware, requires new alliances, new sensitivities, and new readings. I consider how humans have already encountered malware in a tantalising glimpse towards what I term a proto-ecological politics, through exploring the ‘Conficker Working Group’ and their attempts to limit Conficker. I then turn to the implications of ecology in practice – and how this radically decentres the human in discussions ranging from the ‘cyber domain’ to ‘cyber weapons.’ I argue that a lack of appreciation of the vibrancy – in terms of how Jane Bennett uses the term – and as I have argued throughout this thesis on choice, of environments, computation, and malware, mean that we miss an *actor* in our politics, one that means *how we do politics and cybersecurity must change*.

When there is still too little focus on understanding how computation works with, for, and through different marginalised groups, such as refugees (Coles-Kemp and Hansen, 2017), it may seem to be a luxury to emphasise a more-than-human politics. However, it is precisely this focus on malware’s politics that allows us to consider the dominance of a pathological approach to malware and cybersecurity that focuses on, and often presumes, a human that is all-too-often a white, western, middle-aged man. This abstracts computation from its environments, treats malware as an object moving with deterministic intent from its author, and suggests that they are somehow *a priori* always malicious; a universal ‘bad’ from which *we* need to defend ourselves and our spaces. In challenging these patho-logics, a malware

ecology is not only a decomposition of malware, but of how we approach computation and security too. My claim here is that this is a radical undoing of cybersecurity to challenge states and human control as our central concern to move to a plurality of agencies and actors. Drawing on Hayles' cybersemiosis and computational cognition, my argument has been that the interactions between computation (and associated infrastructures) and humans produce ecologies that are always partially unexpected and unknown. An ecological approach insists that different environments, places, and configurations produce outcomes shaped by, and shaping, the *choices* made by malware, humans, and other computational technologies make. But, for computation, there are logical bounds to this more-than-human politics. These choices become imbricated in a complex more-than-human politics – that often makes them difficult to distinguish from a 'humanly' politics.

Therefore, in this concluding chapter I seek to address the claims that I outlined in the introduction of this thesis: how we come to understand software as malicious, how malware are a political actor and influence various forms of political agency, and how this thesis challenges conventional (pathological) notions of cybersecurity to rethink who and what negotiates security. To do so, I first look at a proto-ecological practice that takes contextuality to understand and intervene in ecologies through one final case, the Conficker Working Group. I then turn to how the cyber domain and cyber weapons have been understood and how an ecological re-reading of these concepts opens up and fragments who and what participates in cybersecurity, before I third, conclude the thesis by offering future directions for research.

The Conficker Working Group

“In an unprecedented act of coordination and collaboration, the cybersecurity community, including Microsoft, ICANN, domain registry operators, anti-virus vendors, and academic researchers organized to block the infected computers from reaching the domains – an informal group that was eventually dubbed the Conficker Working Group (CWG).”

(The Rendon Group, 2011, p.ii)

I am turning to Conficker in this concluding chapter to consider a proto-ecological response that is not just a contextualisation but an explicit intervention in ecologies to limit malware agency. Conficker was first 'discovered' by Symantec on 21 November 2008 (Symantec, 2013) when it was quickly propagating across unpatched Windows operating systems using remote

code execution, bringing it to the attention of the computer security industry. Over five distinct versions, the worm was detected during the course of late 2008 and early 2009 (see Table 4 for an overview of each variant). Researchers at the time claimed this to be one of, if not the, most important malware incident to date (The Rendon Group, 2011, p.12), with suggestions that the total number of computers affected range between five to thirteen million (ibid, 2011, p.10). Conficker formed a botnet, allowing for updates that made it a *potentially* serious threat. By this, the performance of Conficker itself was redefining what *was* the security threat. Initially, Conficker exhibited little ability to make *choices* but to remain aligned to the tight, logical, outputs written by its author(s). This morphed later, with Conficker given greater ability to choose through logical frameworks where to propagate as the author attempted to evade actions by the Conficker Working Group (CWG) to limit the propagation of the worm. As Mueller *et al.* (2013) claim, “no Internet security incident has raised the question of the scalability of the networked approach more assertively than the so-called Conficker botnet” (2013, p.95). Not until WannaCry over eight years later would there be such public awareness of propagating malware.

Conficker is interesting in that it exposes a response to malware that is not wholly contextual, but in parts, is a new political formation developing responses in response to the emergent politics of the ecologies in which Conficker formulated and worked within. Tracing the actions of the CWG allows for an observation of how humans came to analyse, detect, and disseminate information about Conficker in ways which suggest an appreciation of its more-than-human agency and politics, rather than solely against its authors. This can provide a glimpse to proto-ecological practice already with us, as a stepping stone to thinking of malware anew. That is, how did humans – as individuals, corporations, non-profits and states – come together to respond to a malware politics, its choices, and its formulations? I am not claiming this is the only example of this activity, but it is one of the first that attempts to intervene in ecologies themselves as a way to limit propagation. Similar forms of proto-ecological response – i.e. contextual responses that respond to ecologies themselves – can be seen in the UK’s ‘Protective DNS’ that limits internet connectivity for known malicious activity, as part of what is known as ‘Active Cyber Defence’ (NCSC, 2017). The Conficker Working Group (CWG), initially known as the ‘Conficker Cabal,’ was however an *ad hoc* group formed to tackle Conficker. In particular, the group focused on preventing the propagation of Conficker by registering internet domains – in order to stop its authors from

attaining them and being able to update Conficker through this method. As Kamluk details, and with much foresight, “while Kido [Conficker] can be pigeon-holed as an ‘old school’ network worm, its success has drawn much media attention and it is a distinct possibility that we may be about to witness a resurgence of such malware attacks” (Kamluk, 2009, p.7). In this section, I wish to disavow the CWG’s response being seen as solely against Conficker’s author(s) – but instead a response to a more-than-human collaboration, where author actions, malware materialities and choices, generated new forms of (in)security, that can be read as proto-ecological.

Table 4: Overview of the five Conficker variants - following Microsoft’s classification scheme. Developed from the Conficker Working Group’s Lessons Learned document (The Rendon Group, 2011).

| Conficker Version | ‘Release’ Date | Detail |
|--------------------------|-----------------------|--|
| <i>Conficker.A</i> | 21 November 2008 | Generation of 250 domains over 5 TLDs ⁹¹ . |
| <i>Conficker.B</i> | 29 December 2008 | Generation of 250 domains over 8 TLDs. |
| <i>Conficker.C</i> | 20 February 2009 | Slight modifications to the code to avoid connection to an internet rendezvous to loads executables to a Conficker host. For information see Porras <i>et al.</i> (2009). |
| <i>Conficker.D</i> | 4 March 2009 | Generation of 50,000 domains, contacting only 500 of these over 110 TLDs. In addition, it introduces a ‘peer-to-peer’ P2P component to update without connecting domains, with additional measures including disabling safe mode and deletes restore points. |
| <i>Conficker.E</i> | 8 April 2009 | Variation on Conficker.D. Installed another malware form, <i>Waladec</i> (scareware – that attempts to get users to download fake anti-virus software), which was programmed to revert back to Conficker.D on 3 May 2009. |

For the CWG, of crucial importance was being able to reverse-engineer the domains Conficker was generating in order to connect to servers controlled by Conficker’s author(s). Conficker, in order to know which domains to connect to, included a Domain Generation Algorithm (DGA). This was pseudorandom (i.e. some variables that are random but with deterministic output). By being pseudorandom, it enabled the author to know which

⁹¹ TLDs are Top Level Domains that refer to highest level of the ‘Domain Name System’ or DNS. These include such examples as .com, .gov., .uk, .eu and so on.

domains to register in order to communicate with the various Conficker forms on millions of computers in advance. As I argued in the previous chapter, this shows how there can be a logical output that is written by the author – such as the pseudorandom DGA – that in some ways reflects the bounds of computational possibility (some values are pre-set) but which also allows the possibility of (randomly generated) malware actions. In the early variants of Conficker, its materiality was tightly bound in order for domain names to be registered so that the author(s) had the opportunity to update their malware. Yet, as this strategy was challenged by the actions of the CWG, this was extended, stretched, and eventually side-stepped, passing greater choices to Conficker, in order for it to evade the actions of the group.

It could be argued that Conficker’s ability to make choices were extremely bounded – and I would agree – at least in versions A, B, and C. Once one moves to version D however, this becomes far more complex. As shown in Table 4, Conficker.D adds another level of ‘randomness’ in which it chooses 500 domains out of 50,000 to connect to. With the introduction of ‘peer-to-peer’ (P2P) networking in D, propagation could occur between computers *without direct human control* – this is where choice became exceptionally important. Conficker could now replicate without direct instruction – meaning it must read environments, and make choices on where to go, choices which are shaped by logical frameworks, but which are none-the-less non-deterministic. It must read the signs of computing environments in order to assess if it is appropriate to propagate. This is not to say that I think Conficker somehow (reflexively) *chose* what to do – but that in imbrications with human choice and decision to create its logics and frameworks, performance was conditioned in this more-than-human register, that enabled an extension away from its author(s) intent, where signs were read in ecologies beyond human control alone.

The CWG’s review document, commissioned by the US Department of Homeland Security’s Science and Technology Directorate, which includes 15 in-depth interviews, provides a unique snapshot of the more-than-human response to Conficker. Though speedy malware analysis was crucial to produce the target domain names (The Rendon Group, 2011, p.20) – it also required a whole group of individuals and organisations, such as the high-level

internet registry, ICANN⁹², to effectively register domains. This required a coming together – in response to a more-than-human politics – between author(s), the logical bounds of computing, the environments of computing (that had *already* been patched by Microsoft but not applied on many computers), and a reading of signs by malware. Previous attempts at restricting the Srizbi botnet, that was responsible for a lot of ‘spam’ at the time, had been (unsuccessfully) attempted earlier that year, serving as a blueprint for the action – with the majority of the individuals involved in this effort being involved with the CWG (ibid, 2011, p.28). This required shared computational resources, such as sinkholes that collected data on the Conficker botnet by redirecting connections from Conficker instances on computing devices that were connecting to domains controlled by the CWG, which were eventually hosted at Georgia Tech in the US. In this, we see how an *ad hoc* more-than-human response was required, with registered domains, humans, sinkholes, servers at Georgia Tech, and a system of passing domains for registration – which apart from two instances of non-registration of domains – were deemed to be broadly successful.

Conficker had a particular geography of movement according to ‘pirated’ (unofficial) versions of the Microsoft operating system, Windows XP. This was due to a October 2008 patch provided by Microsoft in MS08-067 (2008) which secured the vulnerability that Conficker was exploiting, prior to its emergence. Where there were more pirated versions of XP – such as in China, Brazil, and Russia – they were more likely to be affected as they could not update their computers with the patch (Kamluk, 2009, p.6), and there were worries in early January that only 50% of vulnerable computers had been patched (Lidzborski, 2009). There were different disruptive capacities within varying ecologies. As Nathan, from the Sophos MAL said, reflecting on his experience of the early Conficker versions:

Nathan: “If you had patched for Conficker it wouldn't have got in because it was a noob⁹³ malware, you know [laughter]?”

Andrew: “Yeah.”

Nathan: “So one machine gets infected, it can't infect other machines, you've massively reduced the risk... Conficker like there were big worms

⁹² The Internet Corporation for Assigned Names and Numbers (ICANN) is a non-profit organisation that facilitates agreements on rules of governance and maintains the top-level structures of the internet.

⁹³ ‘noob’ is a slang term that refers to a new person or process. Sometimes this can refer to a stupid, ignorant or silly mistake.

that hit. Conficker was probably one of the worst and then you know, at least it wasn't too damaging it didn't ransomware your network.”

(Interview, 21 July 2017)

Though the patch was an essential part of providing protection against Conficker, the particular ecologies of computing at that time, which meant it was still able to propagate, and this revealed how pathological renderings are required in order to analyse, categorise, and detect malware. An ecological malware politics realises that analysis and detection is required – but incorporates it to say that this is not the end of the story.

The *ad hoc* response by the CWG, according to the materialities of Conficker, demonstrated a formation of more-than-human proto-ecological politics of malware – distinct from previous contextual engagements. Though there are many theories of whether Conficker was to experiment new functionality such as propagation frameworks (Touchette, 2016), I am less interested in the debates around attribution. Conficker became a big media story, in part due to its particular materialities as a worm – and how the CWG generated interest through its response. One particular date, 1 April 2009, became significant. Conficker.D was programmed to update to Conficker.E from April Fool's Day 2009. Conficker hit the UK's House of Commons (reported in the Guardian by Arthur (2009)), the UK Ministry of Defence, as well as the German Bundeswehr (the German armed forces), heightened its threat to conventional state politics. This led the German newspaper, *Der Spiegel*, to ask “Alles nur April, April?” [Is it all an April fools?] (Patalong, 2009). Between the joke and the materiality of Conficker, particular articulations in the media allowed it to be both *recognised* as a threat, yet at the same time, render it potentially as harmless – almost acting as a ‘joker’ of computing. The CWG's report on Conficker laments how the media over-hyped the event, but through reading headlines from the period, there appears to be much more nuance across media; providing advice, and a general increase in cybersecurity awareness.

No doubt, much of the work around Conficker, its reporting in the media, and how people approached the malware were infused with a pathological lens. However, what I think Conficker shows is how there was a moment of proto-ecological response. In this, how the CWG responded to malware and the ecologies of its connection, its network, and how information was shared to ensure computers were patched. None of this grouping would have been possible without a particular Conficker materiality, certain logics constructed by

its author. Yes, these logics were changed over time by an author *to permit a greater agency, or choice*, to Conficker as part of a more-than-human collaborative politics. But, Conficker did not simply *execute the intention of the author*, it performed the logics of its code, but how it interacted with different environments, the ecologies of patching, and increasingly given greater flexibility to read signs – such as with the introduction of P2P and more domain names generated in Conficker.D – allowed for new political formations to be generated. This led to analysts having to come together in new ways to respond to its propagation, help refine rules around registration of domains, and even challenge how states response to malware. As former US President Barack Obama said in May 2009;

“[W]hen it comes to cybersecurity, federal agencies have overlapping missions and don't coordinate and communicate nearly as well as they should -- with each other or with the private sector. We saw this in the disorganized response to Conficker, the Internet "worm" that in recent months has infected millions of computers around the world.

This status quo is no longer acceptable -- not when there's so much at stake. We can and we must do better.”

(Obama, 2009)

The CWG as an *ad hoc* collective primarily made up of volunteers was based on personal connections, at least in the first instance. This is still another important aspect of contemporary malware analysis, as I found observing and interacting with analysts at the Virus Bulletin conference. Efforts to coordinate began with the formation of an email listserv by the Shadowserver Foundation on 28th January 2009 (that had swelled to at least 300 members by late March (The Rendon Group, 2011, p.21)). A meeting at the Global DNS Security, Stability, and Resiliency Symposium in Atlanta, USA between 3 – 4 February 2009 was seen by many as the start of the CWG, with an official announcement of its creation on 12 February 2009 (ibid, 2011, p.18). This was accompanied by a US\$250,000 offer for the arrest of Conficker's author by Microsoft. Yet, this response was not led by conventional security actors – but by private individuals and organisations where, “the contribution of traditional security provisioning organizations, such as law enforcement agencies, intelligence agencies, military forces, and even national Computer Emergency Response Teams (CERTs) was negligible” (Mueller, Schmidt and Kuerbis, 2013, p.96) – which is also reported in the CWG's ‘Lessons Learned’ report. This shows how cybersecurity, and malware agency, were part of a reconfiguration of who acts and constructed new alliances

according to its materiality in different versions – that is on privately-owned internet infrastructures, processed through corporate MALs, and other entities. Malware’s politics is not all extensive, but produces new forms of political responses that intertwine with other forms of political activities to produce something new. This sort of activity can now be seen in the promotion of the ‘cyber kill chain’ (Hutchins, Cloppert and Amin, 2011) and its variants, that aim to look at manipulating ecologies to limit the agencies of malware.

Conficker, along with previous events in Estonia in 2007, were formative of a new kind of response; which was not simply about tackling the *human authors* of malware or reverse engineering, but about developing whole new ways of coming together to respond to malware and its ecologies themselves – one that takes malware agency seriously. Unlike endpoint detection products that focus on monitoring, contextualising, and detection, these proto-ecological practices actively intervene in ecologies. It is from this period that there was an increasing awareness from states of a need to become more actively involved as malware agency clearly impinges on conventional state interests. This is not to say the CWG or Conficker ever understood their own work in the terms I am using here. Yet, in forming an *ad hoc* grouping, I re-read its story as not necessarily a response to Conficker’s author(s), but to a broader more-than-human collaborative action and intervening in ecologies themselves. This provides the clearest indication of why an ecological approach is needed: where malware moves beyond its author(s) and means that we must go further than the proto-ecological that is still pathological in order to take more-than-human agency seriously, so that we see its reconfiguration of cybersecurity itself, especially as capacity grows such as through machine learning. It is about realising that software performs, that this has impacts on our politics, and forces us to change our institutions, practices, and ways we think about and in relation to software.

The Cyber domain

More generally perhaps, responses to malware now frame cyberspace as the ‘fifth domain’ of warfare. This is a political response to ecologies, whilst trying to render them as pathological once again – to be human-constructed, controlled, and with space as a neutral variable. As the former US General Michael Hayden claims;

“[C]yber as a domain. It is now enshrined in doctrine: land, sea, air, space, cyber. It trips off the tongue, and frankly I have found the concept liberating when I think about operationalizing this domain. But the other domains are natural, created by God, and this one is the creation of man. Man can actually change this geography, and anything that happens there creates a change in someone’s physical space.”

(Hayden, 2011, p.4)

This sense that cyberspace is ‘the creation of man’ is a seductive one – surely all networks, protocols, devices, and so on have been crafted by ‘us’ – so they should be controllable? Cyberspace is situated as a geography that *can be changed*, unlike the supposedly ‘natural’ geographies of other domains. Though I do not wish to delve too deeply into the disciplinary debates in international relations, strategic studies and related areas, it is important to note how cyberspace has been treated, which has important implications for what is a broadly pathological understanding of how malware performs in these spaces. Though I agree that debates around cyberspace have been repeated “*ad nauseum*” (Smeets, 2018, p.7), it is because they rarely interrogate the premise that cyberspace is something that is full of fissure, friction⁹⁴, and withdrawal away from us through computation. It’s that it has not been explored enough which is the problem – the surprising lack of engagement with literatures on space within some work on cybersecurity and international relations is at major fault here.

Recently, Kello’s (2017) call for a division, in the new field of cyber studies (Smeets, 2018), between cyberspace and the ‘cyber domain’, where the former concerns the technology and the latter the social, political and cultural implications that emerge from this ‘technological’ space. This approach has a seriously underdeveloped sense of spatiality and computation that allows Kello to claim that cyberspace has led to the demise of geography. This is not a tenable distinction to make, as this suggests space is some sort of indistinct, abstracted, non-stratified variable to the more important practices of human action. It places humans on a pedestal above the ecologies in which they work – and suggests we can somehow separate ourselves from the world into a sphere of pure strategy. Early work in the study of computing does not try and distinguish between a technological and ‘human’ construct but through a patho-logic drawing upon informational equivalence, sees these as integrally connected. As Channell (1991, p.120) says, “beginning in the 1930s, several people involved with the

⁹⁴ An article by Justin Kosslyn (2018), “The Internet Needs More Friction” provides an interesting example of how friction could be a concept to improve security, and what it means for understanding cyberspace.

development of the computer began to apply ideas drawn from research in the physiology and philosophy of the mind to the design of computer systems.” During the creation of the (electronic) computer, there was no explicit divide between primarily women’s labour, and technology; they were co-constitutive. Yet, I do not wish to repeat the errors of organicism and cybernetics I explored in chapter two by claiming that through an equivalence we should not divide domains and spaces, but instead explore the various capacities and choices of things. Thus, though *cyber studies* may be a fruitful activity, it does not adequately account for a more-than-human engagement that transcends artificial distinctions between domains, computation, humans, and spaces; they are integrally tied.

Work in the field of international relations does however offer an interesting alternative approach to understanding cybersecurity. Most interestingly this has emerged at the intersection with science and technology studies. In a recent edited collection, *Technology and Agency in International Relations*, Hoijtink and Leese (2019) in their introduction, start to develop how computation and agency are affecting international relations through algorithms and robotics. Likewise Balzacq and Cavelti (2016) attempt to combine actor-network theory, cybersecurity, and international relations to offer some useful insights on malware (albeit limited by an over emphasis on networks, connection, and assemblage). In this theorisation we can then see how malware, agency, and international relations can be brought together – where software must engage with environments, social relations, and a variety of others to be determined as good or bad. But the insistence on actor-network theories or assemblage within work on technology and international relations (Balzacq and Cavelti, 2016; Hoijtink and Leese, 2019; Stevens, 2018; Collier, 2018) leaves little room for the “breaks and gaps, interruptions and intervals, caesuras and tears” (Harrison, 2007, p.592) that I think are necessary for a proper engagement with cyberspace, computation, and more-than-human agencies. Observing what does not connect, what falls short, means that we cannot simply connect, but need to understand disconnection as well. I am asking for something more radical than simply a rethinking of security – but also one that reconfigures understandings of (cyber)space itself in cybersecurity. This does not mean that we have to step away from domains, as borders are essential for establishing knowledge, but that attempting to somehow exclude the broader ecologies, that are saturated with the lineages that I have outlined in this thesis, are a fallacy.

The object-ive cyberweapon?

There is a new actor in our politics. The mis-understandings of malware's spatiality has implications for how we respond to the threats it represents. In the recent attention given to 'cyberwar' and 'cyberweapons' as a means of achieving the former, there have been many pages dedicated to the topic. Unfortunately, much of this work makes a mistake of claiming that one can ignore the materiality of malware – and even that malware are something of no interest to international relations or politics unless they acquire the intent required for malware to be a 'weapon' (Kello, 2017; Landau, Lin and Bellovin, 2017; Rid and McBurney, 2012). This is due to the fundamental misunderstanding of computation – with comparisons to nuclear arms (Nye, 2011), or other weapons as tools missing the point (see Singer and Shachtman, 2011, for this argument). Not one that is just unruly and expressing impressive affective capacity – but one that *cognises* and develops choices as part of a computing ecology. Kello says that “the virtual weapon's payload, which, like all computer code, consists of just 0's and 1's in the form of electrons can travel through cyberspace to destroy nuclear enrichment turbines, paralyze the operations of large firms, incapacitate tens of thousands of workstations at large multinational firms, or disrupt stock-trading platforms” (2017, p.5). This does little to theorise what is distinctive between different software forms. Nor does it comprehend that, in 'traveling' through cyberspace, choices are distinct from liberal notions of *reflective*, conscious choice, but instead through cybersemiosis develop a politics through the layers of computation that interact with choices and decisions of humans to produce new coagulates of more-than-human performativity.

Though malware are frequently seen as an important *thing*, there has been little explicit theorisation of what they are, apart from when dealing with 'cyberweapons' (Langner, 2011; Barzashka, 2013; Stevens, 2017). This reduces malware to a discursive object that are neutral to environments and, if they can be deployed with the correct configuration, malware are simply a tool which execute based on their author's intent, rather than the complex interactions of political agency and affective capacity I outlined in chapter seven. Work on *cyberwar* (Rid, 2013) often does not explicitly address the materiality of malware but instead concerns itself with contributions to theory based on weak assumptions about how computation and software logically operate. Though 'emergence' in the cyber domain is discussed (McGraw, 2013), this focuses only on identifying 'collateral' damage. Collateral

damage exists because we think that the rational malware author has made a mistake – but this denies a role to the performative role of not only malware, but the spaces in which they come to exist. What is this collateral damage if not performative ecologies and the choices of computation and malware?

The NATO CCDOE-commissioned first-edition of the *Tallinn Manual* states, that;

“for the purposes of this manual, cyber weapons are cyber means of warfare that are by design, use, or intended use capable of causing either (i) injury to, or death of, persons; or (ii) damage to, or destruction of objects, that is, causing the consequences required for qualification of a cyber operation as an attack.” (Schmitt, 2013, p.141, Rule 41 (2))

The Manual is seen as the *de facto*, if not *de jure*, word on cyber engagements in the ‘west’; what do, and do not, constitute acts of cyberwar. I am less interested in the particularities of the manual’s detailed explanations, but how it treats malware. In the glossary, there is a rejection of malware for ‘malicious logic,’ where the cyber weapon is an advanced form which must have *physical* effects. Stuxnet is clearly an inspiration, its name littering the manual’s pages. A weapon thus must have intent somehow *imbibed* within ‘malicious logic,’ to follow the manual’s description. Or as one reverse engineer said at the UK’s National Cyber Security Centre, “it’s all about *intent*” (Interview, 19 October 2017). Yet, as I have argued, there is no such thing as a wholly clear *malicious* logic embedded, and it requires certain environments, ecologies in which software (and other components) perform within which to become malicious as such. Thus, I refute that cyber weapons are somehow distinct to other forms of malware (apart from the strategic intent of their authors – as much as that influences its performance!). Though there have been extensive reservations to think of the ‘strategic consequences’ of cyber weapons (Kello, 2017), this is to an exclusion of understanding how cyber weapons are constructed and intertwine with other malicious forms. Malware are an important part of the story – understanding how software do not always operate as expected. This comes in the glitch and the rebooting computer seen with Stuxnet.

As Kim Zetter notes in *Countdown to Zero Day*:

“It has been a long and improbable ride that was made possible only be a series of unfortunate events and flubs that should never have occurred – from the zero-day exploits that launched Stuxnet on its wild odyssey through thousands of machines around the world to the crashing machines in Iran that first gave it away; from the green researchers in Belarus who lacked the skill and experience to tackle a threat like Stuxnet to the Symantec researchers who bumbled their way through the PLC code; from the Wiper tool in Iran that led the Kaspersky team to uncover Flame to

the server in Malaysia that locked the attackers out and preserved a mountain of forensic evidence for researchers to find. There were so many things that had to go wrong for Stuxnet and its arsenal of tools to be discovered and deciphered that it's a wonder any of it occurred.”

(Zetter, 2014, p.306)

However, it is the environments, the ecologies in which malware operate and perform, that meant there was always a potential for Stuxnet and other malware such as Conficker, the Dukes, and WannaCry/(Not)Petya to exert affects, and more importantly, political agency, through an active interpretation of signs that may or may not align with human expectation. Or as David Sanger has written on Stuxnet;

“Why was a computer worm that had been painstakingly designed to release itself only if detected by computer controllers connected to a specific array of centrifuges at Natanz suddenly zipping through the Internet like a newly released videogame? The answer appeared to be one that Microsoft and every software manufacturer has discovered sooner or later: poorly tested new releases of software can generate all kinds of unanticipated results.”

(Sanger, 2012, p.204)

A malware politics cannot be disassociated from its ecologies – that is computation, social affects, and multiple others – in the same way we should not abstract animals, bacteria, and humans from theirs. In the pathologic of saying that malware can be statically analysed, through to the abstractions of contextual analyses and detections, is but a particular flavour of a computational politics, in which ecology helps to condition and allow for the potential for emergence and performance to be generative of certain actions, discourses, and politics. Stuxnet led to the removal of centrifuges and accusations of US and Israeli involvement in Iranian affairs by the then President of Iran, Mahmoud Ahmadinejad saying “they [the USA and Israel] succeeded in creating problems for a limited number of our centrifuges with the software they had installed in electronic parts” (Yong and Worth, 2010). Ecology does not reduce politics or international relations, but orientates it towards a far more complex, fuller realisation of who and what acts.

There can be no *objective* cyberweapon, however. Malware are more-than-human collaborations, making choices, conditioning and being conditioned by ecologies, forming a politics through cybersemiosis and cyberspace, that intersect with a humanly politics in ways complex, folding, and unforeseen. Malware are not a tool – they are cognisers. As malware cannot be abstracted from its environments, the sign-reading and interpretation of malware

do not always work in ways intended. Authors can maintain a tight control over their more-than-human collaboration through a logical output – but this is always prone to a new environment, a reboot, or an unexpected ‘error.’ A mechanical tool follows a deterministic process, which can be accurately modelled (to an extent of the vibrancy of environments), but due to electronic computational choice, this becomes complex, logical, but not determinate. This is contrast to a paper by Thomas Rid and McBurney, ‘Cyber-Weapons,’ where they say:

“We understand a weapon as a *tool* that is used, or designed to be used, with the aim of threatening or causing physical, functional, or mental harm to structures, systems or living things. This general definition is an essential building block for developing a more precise understanding of cyber-weapons, and be extension cyber-conflict.”

(Rid and McBurney, 2012, p.7, emphasis added)

Others have referred to cyberweapons as “instruments” (Zegart and Lin, 2017) and as “tools” (Dewar, 2017). When talking with experts in cybersecurity, Silomon and Roeling (2018) expose issues regarding the discursive overlap between software, malware, and weapon – which shows a great flexibility, and perhaps confusion, between the terms. This is further complicated in that their use may be ‘defensive’ or ‘offensive’ (or as is sometimes, discussed, ‘dual-use’ – i.e. for malicious and non-malicious activities). As Stevens (2017, pp.2–3) contends, due to this dual-use “malware would cease to be a weapon *per se*” if it is used for beneficial reasons, and that reveals the “fuzzy boundaries” between a cyberweapon and software. Then, if we consider in ways Balzacq and Cavelty have, that malware are performative, as I have argued in this thesis, along with computational cognition and ecology, then a cyberweapon cannot exist apart from during its performance. Yet, this performance in itself is full of more-than-human *choice* and alliance, meaning that, though it may be closely tied to the logics of the author, it does not necessary directly map on to the ecologies of malware’s emergence. Thus, can we call such an alliance a weapon? One that is not simply an object that moves through a neutral space, but is constantly reading, interpreting, and choosing, being reformed, and reformulating ecologies? I wager that we are experiencing the growth of more-than-human collaborations – where malware become constructed as ‘cyberweapons’ to fit into discourses of previous articulations of weapon, tool, and war – but they are distinctly different. Thus, in questioning what forms a cyberweapon, I open-up broader questions around computation in cybersecurity to permit a reconsideration of who

acts and performs in its politics, war, and whether we can ever have such a human-centred view of such things are drones, cyber-attacks, autonomous vehicles, and supporting computational processes permeate not only our homes through the IoT but saturate militaries and their state infrastructures.

Drawing a Close

Throughout this thesis, I have attempted to draw a politics of cybersecurity that addresses several components of what I see as a current *patho*-logic to how we understand malware, and what thinking through an *eco*-logic can provide to this perspective. Not to counter it, but to think of ways to extend and compliment current study. The reason I retain the term 'logic' is to signal that there is a boundedness to computation in that it cannot exceed the calculative modes it operates upon. This has required rethinking computation's role in the world – that draws extensively on N. Katherine Hayles' work on cybersemiosis and the role of cognition – along with work in new materialism and geography – to consider how this is a more-than-human construct rather than something that is beyond humans. Through doing so, I hope to have moved differently to the terms 'digital' and 'virtual' to describe and make knowable malware – and offer cyberspace as a critical intervention to talk of a computational agency and a malware politics.

Malware make choices that open-up a limited politics and lead to more-than-human collaborations that intertwine human and computation to produce something new, such as the crashing computers in Iran with Stuxnet. This has implications for security, as when there are more-than-human choices being formed, security is written, defined, articulated beyond us. Hence, in ways often imperceptible, malware constructs cybersecurity itself. The irreducibility of ecologies, where malware choices are political (through spaces of withdrawal) mean that not only are we faced with emergence but something much more unsettling. That is an (in)security that escapes us. We do not have exclusive choices on how to secure but must recognise how malware are a political actor that shape other political formulations – such as around sensation, affective identities, geopolitics, and so on. Where we articulate, for instance, the gradients between choice and decision as a complex recursive form of choice require further resolution, but I think choice is something that is currently broadly unaccounted for. I have explored the Conficker Working Group to understand how

they responded to Conficker by attending to ecologies themselves, but through a contextual lens, as an example of alternative responses to more-than-human threats beyond endpoint protection. The attention between the two logics could be summarised as one that seeks to contain and limit the potential for malware to be *political* whereas the ecological embraces this potential. Attending to ecologies does not necessarily change the way we detect but make us aware that *we* are not in control of our cybersecurity; it is also shaped by the choices of more-than-humans both malware and those computational technologies that seek to secure in the MAL.

I did not intend to side-step the lineages of how we have come to our present juncture; where pathology informs how we respond to computing security – informed through biological, medical, and militaristic narratives. Yet, without this critical attention, we are left with a controllable cyberspace, neutral, one that is ‘man-made’ and informed through long running intersections from organicism and cybernetics that I explored in chapter two. The lineages of understandings between humans and others have led to an equivalence of machines being able to transfer contextless information, yet simultaneously placing them below animals due to their lack of organicity and human construction – leading to them being of a lower status. Thus, we have a dual bind where we can equate malware – through terms such as virus and worm – to their biological equivalents and their emergent potential whilst also being able to say they do not exhibit choice and are bounded by the intent of their human authors. This cannot stand if we are to understand the way malware works and the ecologies within which it generates and forms a politics. Yet, in recognising these lineages, it might become possible to revive and open them up for critique that allows for an ecological conceptualisation. Computation, through its layers, and through malware, make choices and are thus political actors in our world – on a different gradient from humans, but in similar ways to choices made by animals and plants, in an expanded view of who and what can be political. By ‘staying with the trouble’ as Donna Haraway implores, it is essential to recognise that computational, and thus malware, agency is limited and is tied in a more-than-human collaboration with its authors, the environments it performs in, and the social relations that curate and disseminate knowledge within and beyond MALs – where authors set malware’s logical bounds, but these are not the end of the story.

When I entered the MAL, these biopolitical lineages echoed alongside ecological resonances for the *becoming-analyst*, in the t-shirts worn by other analysts and in how the books I learnt from discussed terms such as confidentiality, integrity, and availability. They hung in the air as I could not escape words like infection, the ‘wild,’ and rendering malware as something that can be analysed, detected, and abstracted from its environments or monitored in virtual machines. Yet, in becoming-analyst I was able to sense an alternative reading of and working with malware that takes surprises, bugs, and errors to be expressions of malicious choice, that then ties with collaborations with authors who code malware to make choices. In the MAL, the way we make sense of malware comes itself from a more-than-human collaboration between *static* and *contextual* strategies for analysis and detection. Through the various tools I discussed in chapter five, it is clear how complex this process of delineating between the ‘clean’ and ‘malicious’ becomes. As malware do not exist *a priori*, there is extensive labour that go into the process of using different techniques with various datasets to render knowable something as abnormal; where a signal of abnormality must be distilled from anomalies. As the MAL moves from affective and embodied ways of analysing malware that conventionally coalesced around static strategies, to contextual use of data and behaviour mean that where the abnormal and anomaly lie are becoming increasingly difficult to distinguish.

As I explored in chapter six, this then becomes part of the curation of the MAL. The various tools and analysts must come together to delineate the malicious and clean to avoid the false positive. This has led to a division between the ‘experimental’ process that an analyst goes through to analyse malware along with an industrious output in the detection that can be rigorously tested and then sent to endpoint detection engines. I argue that the process of curation is a in itself a co-curated *political* act by analysts and the range of supporting others including Quality Assurance and computing infrastructures. That is, what are the commercial logics that condition such curations, and how the (ab)normal is defined? No more so when this becomes evident in the ‘Potentially Unwanted Application’ (PUA) which, as Daniel described may be created for legal reasons, but also shows the fallacy of the malicious without the ecologies it operates in. This makes the MAL a translator, and thus curator, of malware even though it frequently has little access to all the ecologies it operates in – and the PUA is part of its realisation. In other ways, the MAL also is a crucial site of the

formation of a politics of cybersecurity, in what it chooses to detect, what to prioritise, and how to curate malware as knowledge for publics beyond the MAL.

I could have focused exclusively on the humans who analyse and detect malware, who translate these knowledges into visualisations and blog posts to be consumed outside the MAL. These practices do influence how we pathologically know of malware and how this affects politics and international relations. Instead, I questioned malware as something that exists prior to its encounter with a computing device or the MAL. This draws on a range of work on more-than-human geographies from exploring spaces of engagement with animals (Dowling, Lloyd and Suchet-Pearson, 2016; Madden, 2014; Probyn, 2014), through to the microbiome and biological viruses (Greenhough, 2012; Lorimer, 2016; Greenhough et al., 2018). In a departure from current study, I worked with malware, to become with it, in spaces of dissection, reformulation and categorisation. Unlike most work on more-than-human geographies thus far, I took the vibrancy of ‘non-vital’ objects to be central, following work by Barry (2010, 2013), Forman (2018), and Glouftsiou (2018). I employed post-phenomenological methodologies in developing an acquaintance with malware, dealing with its resistance, and its politics, in ways advocated by Sarah Whatmore’s (2006) call for a materialist return in a more-than-human world. In this way, I followed Thrift’s claim that “this new world is one of not-quite-life but so close to the conduct of life that it is not-quite-inert either” (2005, p.473). By this, computation cannot simply be seen as inert, but occupies spaces that do not easily fall into our conventional categorisations.

The curations of multiple MALs are commonly the only points of access for publics to understand how malware operates and becomes designated as so. This is not to say publics may understand that in this way, but in what informs policy, newspaper articles, and visualisations that they come into contact with. This is important when considering how malware have a politics – in that these reports often pathologically refer back to the human author (if they do attribution) and abstract performativity from malware. So, in chapter seven, I attempted to provide an alternative reading of the cases of Stuxnet, the Dukes, and WannaCry/(Not)Petya. Apart from the latter case, I was limited to as to how I can express the ecological as I was not there, in those ecologies, and thus can only partially reconstruct some strands to explore a malware agency through its pathological recording. It is within this chapter that I tried to codify different forms of malware agency – from tightly-bounded

logical outputs to surprise – without forgetting the affective capacities of their presence and disruptive power to effect change in practices, formations, and politics.

Future Directions

Thinking ecologically requires new ways of engaging with and participating with malware; as something that expresses a politics that extends away from us. This is not easy – as it requires fundamentally challenging the role of computation in our societies. This is something that I think is already starting to occur in work exploring the role of machine learning and algorithmic practices (Amoore, 2018, 2019, Forthcoming; Amoore and Raley, 2017; Aradau and Blanke, 2016, 2018; Halpern, 2015; MacKenzie, 2015; Noble, 2018; Parisi, 2017; Hayles, 2017). To these debates I contribute to an argument for thinking ecologically, appreciating that software is more than a coded expression and execution of human intent. This thesis is by no means a complete overview of what ecology may and could do for expanding the scope of malware and its agency. Further work may continue, for example to challenge how we talk of cyber weapons, and what that means for future more-than-human collaborations as states increasingly invest in these techniques, which are responsible for acts of great harm for individuals and organisations affected. It must also reckon with how choice and political agency intersect and what difference this means for politics, and how this in particular relates to other forms of affectual relations. For me, however, this is about questioning cybersecurity’s core foundations around attribution, intent, execution, the role of the hacker, and even weapons and war. Yet, this is only the start of such questioning when choices mean it is not the ‘ecology’ or ‘environment’ always speaking back, but something else, meaning can we necessarily still call this ‘collateral damage’?

In the four years since the germination of this doctoral research, cybersecurity has gained increased attention and dominance within (inter)national debates. Discussions over what is a ‘cyberweapon,’ the media attention over the wiping of computers in several businesses by the malware form (Not)Petya, and the impact on the UK’s National Health Service in May 2017 by WannaCry (Dwyer, 2018; Morse, 2017), have made this thesis both lively and difficult to bind. Cybersecurity is constantly reforming at bewildering speed, making many things written quickly fall out of contemporary debate. This makes the area exciting, but leads to issues in developing substantive, detailed engagement with computation which are not based

on the next ‘new’ thing. We have a severe deficit in understanding malicious software beyond its technological and strategic implications. This thesis is intended to draw a tenuous thread between these perspectives and offer something new. This required a blending of technical and social science techniques to achieve a more holistic (and yet still far from complete) approach. More of the research required do not sit in disciplinary silos – it is a difficult but essential task to cross these. This does not mean we cannot talk to, or situate, cybersecurity in our disciplinary homes, but that it is imperative to talk outwards, make mistakes, and perhaps achieve something distinctive.

An ecological approach requires talking to international relations that seeks to develop doctrines and strategies of dealing with malicious software. There are serious implications for geography’s lack of engagement thus far – with Lucas Kello (2017, 2013) being able to argue that ‘geography’ no longer matters in discussions on cybersecurity – and how authors frequently lament that geography’s impact is limited. This is based on an unfortunate conflation between geography and cartesian geometry. In allowing for the discipline of geography to have a conversation in cybersecurity, I believe, there can be fruitful findings. Geography is well placed to tackle ‘big’ problems that are made manifest in particular locales as shown through work on climate change, global development, on the varying strands of geopolitics, through to analysing the ‘War on Terror,’ and the impact of data and surveillance. Cybersecurity has thus far had relatively little attention within geography – apart from some isolated examples (Crampton, 2019; Kaiser, 2015; O’Grady and Dwyer, 2020, Forthcoming; Simon and de Goede, 2015; Warf, 2015). The contribution of geography is its relation to varying scales, places, and times – allows for a nuanced, and detailed consideration of how things become secured and become manifest through various infrastructures, discourses, affects, and, in the case of this thesis, through malware being treated as neither human nor wholly non-human, but as *more-than-human*.

This thesis has called for a new way of perceiving the politics of malware that fully embraces computation as an actor through the use and extension of the work of N. Katherine Hayles in *Unthought* (2017) and subsequent work on cybersemiotics (2018, 2019, Forthcoming). This means challenging who and what makes choices – and that it is choice itself that forms the political. This is a departure from much work from new materialisms that often does not engage with opinion and choice. Though something may have immense affectual capacity,

to be powerful, this does not make it political. Politics is reserved for those that performatively cognise (read, interpret, and act on signs) at particular junctures within ecologies. Malware choice, within logical bounds, can become enrolled through outputs and frameworks into malicious activity, and at the same time extend away, and based on ecologies formulate choices and thus action based on this. Furthermore, though this may be hard for us to recognise, when we infuse *our* political discourses with an ability to (apparently) effect change in a 'rational' way, our understanding is always shaped by what is around us, our ecology. I am claiming no different for malware. This is not to say that malware are somehow consciously aware of these decisions (as a particular form of cognition that allows self-awareness; to humans and some other animals), indeed these signs may have been learnt, may be *programmed*, but what computation does allow is a greater ability for software to *choose*, making malware a political interlocutor, constructing a cybersecurity that extends away from us and that is the challenge: how we even articulate what incorporating truly alien forms of cognition into our politics means and will do for our future.

Bibliography

- AAG, 2017. *The Association of American Geographers 2017 Annual Meeting Program*. [Annual Meeting 2017] Boston. Available at: <https://web.archive.org/web/20170812113811/http://www.aag.org/galleries/conferenc-e-files/AAG_2017_Printed_Program_FULL.pdf> [Accessed 12 Aug. 2018].
- Adey, P., 2014. Security Atmospheres or the Crystallisation of Worlds. *Environment and Planning D: Society and Space*, 32(5), pp.834–51.
- Adey, P. and Anderson, B., 2011. Anticipation, Materiality, Event: the Icelandic Ash Cloud Disruption and the Security of Mobility. *Mobilities*, 6(1), pp.11–20.
- Adey, P. and Anderson, B., 2012. Anticipating Emergencies: Technologies of Preparedness and the Matter of Security. *Security Dialogue*, 43(2), pp.99–117.
- Agamben, G., 2005. *State of Exception*. Chicago: University of Chicago Press.
- Agrawal, S., Goswami, K. and Chatterjee, B., 2010. The Evolution of Offshore Outsourcing in India. *Global Business Review*, 11(2), pp.239–256.
- Alperovitch, D., 2016. *Bears in the Midst: Intrusion into the Democratic National Committee*. [online] Available at: <<https://perma.cc/S6YC-ZHBQ>> [Accessed 10 Oct. 2016].
- Amoore, L., 2009. Algorithmic War: Everyday Geographies of the War on Terror. *Antipode*, 41(1), pp.49–69.
- Amoore, L., 2011. Data Derivatives: On the Emergence of a Security Risk Calculus for Our Times. *Theory, Culture & Society*, 28(6), pp.24–43.
- Amoore, L., 2013. *The Politics of Possibility: Risk and Security Beyond Probability*. Durham, N.C.: Duke University Press.
- Amoore, L., 2014. Security and the incalculable. *Security Dialogue*, 45(5), pp.423–439.
- Amoore, L., 2015. Cloud Geographies: Computing, Calculation, Sovereignty. In: *Progress in Human Geography Lecture*. [online] RGS-IBG Annual International Conference. Exeter, UK. Available at: <<http://conference.rgs.org/AC2015/186>>.
- Amoore, L., 2016. Cloud geographies: Computing, data, sovereignty. *Progress in Human Geography*. [online] Available at: <<http://phg.sagepub.com/content/early/2016/08/10/0309132516662147.abstract>>.
- Amoore, L., 2018. Doubtful Algorithms: Of Machine Learning Truths and Partial Accounts. *Theory, Culture & Society*. [online] Available at: <<http://dro.dur.ac.uk/26913/1/26913.pdf>>.
- Amoore, L., 2019. *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others*. Durham, N.C.: Duke University Press.
- Amoore, L. and de Goede, M., 2008. Transactions after 9/11: the banal face of the preemptive strike. *Transactions of the Institute of British Geographers*, 33(2), pp.173–185.
- Amoore, L. and Hall, A., 2009. Taking People Apart: Digitised Dissection and the Body at the Border. *Environment and Planning D: Society and Space*, 27(3), pp.444–464.
- Amoore, L. and Hall, A., 2013. The clown at the gates of the camp: Sovereignty, resistance and the figure of the fool. *Security Dialogue*, 44(2), pp.93–110.
- Amoore, L. and Piotukh, V., 2019. Interview with N. Katherine Hayles. *Theory, Culture & Society*, 36(2), pp.145–155.
- Amoore, L. and Raley, R., 2017. Securing with algorithms: Knowledge, decision, sovereignty. *Security Dialogue*, 48(1), pp.3–10.
- Anderson, B., 2015. Boredom, excitement and other security affects. *Dialogues in Human Geography*, 5(3), pp.271–274.
- Anderson, B. and Gordon, R., 2016. Government and (non)event: the promise of control. *Social & Cultural Geography*, pp.1–20.

- Anderson, L., 2006. Analytic Autoethnography. *Journal of Contemporary Ethnography*, 35(4), pp.373–395.
- Aradau, C., 2010. Security That Matters: Critical infrastructure and Objects of Protection. *Security Dialogue*, 41(5), pp.491–514.
- Aradau, C. and Blanke, T., 2016. Politics of prediction: Security and the time/space of governmentality in the age of big data. *European Journal of Social Theory*, 20(3), pp.373–391.
- Aradau, C. and Blanke, T., 2018. Governing others: Anomaly and the algorithmic subject of security. *European Journal of International Security*, 3(1), pp.1–21.
- Arntz, P., 2017. *Explained: YARA rules*. [online] Malwarebytes Labs. Available at: <<https://archive.is/VCMTy>> [Accessed 20 Jun. 2019].
- Arthur, C., 2009. House of Commons network hit by Conficker computer worm. *The Guardian*. [online] 27 Mar. Available at: <<https://web.archive.org/web/20190215133502/https://www.theguardian.com/technology/2009/mar/27/conficker-downadup-parliament-virus-april-1>> [Accessed 15 Feb. 2019].
- Ash, J., 2015. *The Interface Envelope: Gaming, Technology, Power*. New York: Bloomsbury Academic.
- Ash, J., Anderson, B., Gordon, R. and Langley, P., 2017. Unit, vibration, tone: a post-phenomenological method for researching digital interfaces. *cultural geographies*, 25(1), pp.165–181.
- Ash, J., Kitchin, R. and Leszczynski, A. eds., 2019. *Digital Geographies*. London: SAGE.
- Ash, J. and Simpson, P., 2016. Geography and post-phenomenology. *Progress in Human Geography*, 40(1), pp.48–66.
- AV-Test, 2017. *Security Report 2016/2017*. [online] Magdeburg, Germany. Available at: <https://web.archive.org/web/20170925201935/https://www.av-test.org/fileadmin/pdf/security_report/AV-TEST_Security_Report_2016-2017.pdf> [Accessed 4 Jan. 2018].
- AV-Test, 2018. *Security Report 2017/18*. [online] Magdeburg, Germany: AV-Test. Available at: <https://web.archive.org/web/20181129124941/https://www.av-test.org/fileadmin/pdf/security_report/AV-TEST_Security_Report_2017-2018.pdf> [Accessed 29 Nov. 2018].
- Aycock, J., 2006. *Computer Viruses and Malware*. New York: Springer Science & Business Media.
- Balaban, D., 2017. February 2017: The Month in Ransomware. *The State of Security*. Available at: <<https://perma.cc/UYK2-Z3TC>> [Accessed 9 Oct. 2018].
- Balzacq, T. and Cavelti, M.D., 2016. A theory of actor-network for cyber-security. *European Journal of International Security*, 1(02), pp.176–198.
- Barad, K.M., 2007. *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*. Durham, N.C.: Duke University Press.
- Barak, G., 2017. *Multiple Groups Have Been Exploiting ETERNALBLUE Weeks Before WannaCry*. [online] Secdo. Available at: <<https://web.archive.org/web/20181009095819/https://blog.secdo.com/multiple-groups-exploiting-eternalblue-weeks-before-wannacry>> [Accessed 9 Oct. 2018].
- Barnard-Wills, D. and Ashenden, D., 2012. Securing Virtual Space: Cyber War, Cyber Terror, and Risk. *Space and culture*, 15(2), pp.110–123.
- Barry, A., 2010. Materialist Politics: Metallurgy. In: B. Braun and S.J. Whatmore, eds. *Political matter: Technoscience, Democracy and Public Life*. Minneapolis, Minn.: University of Minnesota Press. pp.89–117.
- Barry, A., 2013. *Material Politics: Disputes Along the Pipeline*. Oxford: John Wiley & Sons.
- Barthes, R., 1967. *Elements of Semiology*. Translated by A. Lavers. and Translated by C. Smith. New York: Hill and Wang.

- Barua, M., 2014. Volatile Ecologies: Towards a Material Politics of Human—Animal Relations. *Environment and Planning A*, 46(6), pp.1462–1478.
- Barzashka, I., 2013. Are Cyber-Weapons Effective? *The RUSI Journal*, 158(2), pp.48–56.
- Bassaganya-Riera, J., 2015. *Computational Immunology: Models and Tools*. Oxford: Elsevier.
- Bateson, G., 1979. *Mind and Nature: A Necessary Unity*. London: Wildwood House.
- Baz, M., 2017. Dridex v4 - Atombombing and other surprises. *Virus Bulletin*. Madrid, Spain. pp.27–30.
- Beck, U., 1992. *Risk Society: Towards a New Modernity*. London: Sage.
- Ben-Asher, N. and Gonzalez, C., 2015. Effects of cyber security knowledge on attack detection. *Computers in Human Behavior*, 48, pp.51–61.
- Benjamin, W., 1992. *Illuminations*. New ed. ed. London: Fontana.
- Bennett, J., 2010. *Vibrant Matter: A Political Ecology of Things*. Durham, N.C.: Duke University Press.
- Berlant, L., 2007. On the Case. *Critical Inquiry*, 33(4), pp.663–672.
- BetaFred, 2017. *Microsoft Security Bulletin MS17-010 - Critical*. [online] Available at: <<https://perma.cc/6XUK-88WB>> [Accessed 9 Oct. 2018].
- Betz, D.J. and Stevens, T., 2011. *Cyberspace and the State: Towards a Strategy for Cyber-power*. Oxford: Routledge.
- Bingham, N., 1999. Unthinkable complexity? Cyberspace otherwise. In: M. Crang, P. Crang and J. May, eds. *Virtual Geographies: Bodies, Space and Relations*. Oxford: Routledge. pp.244–260.
- Bochner, A. and Ellis, C., 2016. *Evocative Autoethnography: Writing Lives and Telling Stories*. Oxford: Routledge.
- Bogost, I., 2012. *Alien Phenomenology, or What It's Like to Be a Thing*. Minneapolis, Minn.: University of Minnesota Press.
- Boldewin, F., 2010. Reconstructor. *Reconstructor*. Available at: <<https://web.archive.org/web/20110722005423/http://www.reconstructor.org/main.html>> [Accessed 22 Jul. 2010].
- Bosch van Rosenthal, E., 2018. Dutch intelligence first to alert U.S. about Russian hack of Democratic Party. *Nieuwsuur*. [online] 25 Jan. Available at: <<https://web.archive.org/web/20181104194603/https://nos.nl/nieuwsuur/artikel/2213767-dutch-intelligence-first-to-alert-u-s-about-russian-hack-of-democratic-party.html>> [Accessed 4 Nov. 2018].
- Bowker, G.C. and Star, S.L., 1999. *Sorting Things Out: Classification and its Consequences*. Inside Technology. Cambridge, Mass.: MIT Press.
- boyd, danah and Crawford, K., 2012. Critical Questions for Big Data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, Communication & Society*, 15(5), pp.662–679.
- Brain: Searching for the first PC virus in Pakistan*. 2011. F-Secure Available at: <<https://www.youtube.com/watch?v=lnedOWfPKT0>> [Accessed 19 May 2019].
- Bratton, B.H., 2015. *The Stack: On Software and Sovereignty*. Software studies. Cambridge, Mass.: MIT Press.
- Braun, B., 2007. Biopolitics and the molecularization of life. *cultural geographies*, 14(1), pp.6–28.
- Braun, B. and Whatmore, S., 2010. *Political Matter: Technoscience, Democracy, and Public Life*. Minneapolis, Minn.: University of Minnesota Press.
- Brier, S., 1996. Cybersemiotics: A New Interdisciplinary Development Applied to the Problems of Knowledge Organization and Document Retrieval in Information Science. *Journal of Documentation*, 52(3), pp.296–344.
- Brier, S., 2008. *Cybersemiotics: Why Information Is Not Enough!* London: University of Toronto Press.

- Brier, S., 2014. The Transdisciplinary View of Information Theory from a Cybersemiotic Perspective. In: *Theories of Information, Communication and Knowledge*. Springer. pp.23–49.
- Bruni, A., 2005. Shadowing Software and Clinical Records: On the Ethnography of Non-Humans and Heterogeneous Contexts. *Organization*, 12(3), pp.357–378.
- Bryant, L., 2011. *The Democracy of Objects*. Online: Open Humanities Press.
- Buchanan, B., 2016. *The Cybersecurity Dilemma: Hacking, Trust, and Fear Between Nations*. Oxford: Oxford University Press.
- Buller, H., 2014. Animal geographies II: Methods. *Progress in Human Geography*, 39(3), pp.374–384.
- Butz, D. and Besio, K., 2009. Autoethnography. *Geography Compass*, 3(5), pp.1660–1674.
- Canguilhem, G., 1989. *The Normal and the Pathological*. Translated by C. Fawcett R. and Translated by R. Cohen S. New York: Zone Books.
- Cardenas, A.A., Amin, S. and Sastry, S., 2008. Secure Control: Towards Survivable Cyber-Physical Systems. In: *The 28th International Conference on Distributed Computing Systems Workshops*. IEEE. pp.495–500.
- Carr, M., 2016. Public–private partnerships in national cyber-security strategies. *International Affairs*, 92(1), pp.43–62.
- Carr, M. and Tanczer, L.M., 2018. UK cybersecurity industrial policy: an analysis of drivers, market failures and interventions. *Journal of Cyber Policy*, 3(3), pp.430–444.
- Cass, N., Schwanen, T. and Shove, E., 2018. Infrastructures, intersections and societal transformations. *Technological Forecasting and Social Change*, 137, pp.160–167.
- Castree, N., 2005. The epistemology of particulars: Human geography, case studies and ‘context’. *Geoforum*, 36(5), pp.541–544.
- Castree, N., 2015. The Anthropocene: a primer for geographers. *Geography*, 100 (Summer 2015), pp.66–75.
- Cavelty, M.D., 2015. Cyber-security and Private Actors. In: R. Abrahamsen and A. Leander, eds. *Routledge Handbook of Private Security Studies*. Oxford: Routledge.
- CCDCOE, 2019. *About us*. [online] NATO Cooperative Cyber Defence Centre of Excellence. Available at: <<https://web.archive.org/web/20190128132850/https://ccdcoe.org/about-us/>> [Accessed 28 Jan. 2019].
- Channell, D.F., 1991. *The Vital Machine: A Study of Technology and Organic Life*. Oxford: Oxford University Press.
- Chen, H., 2011. *Dark Web: Exploring and Data Mining the Dark Side of the Web*. London: Springer Science & Business Media.
- Chun, W.H.K., 2005. On Software, or the Persistence of Visual Knowledge. *Grey Room*, Winter 2005(18), pp.26–51.
- Clough, P.T., Gregory, K., Haber, B. and Scannell, R.J., 2015. The Datalogical Turn. In: P. Vannini, ed. *Non-Representational Methodologies: Re-Envisioning Research*. London: Taylor & Francis. p.146.
- Cobb, S. and Lee, A., 2014. Malware is called malicious for a reason: The risks of weaponizing code. Cyber Conflict (CyCon 2014), 2014 6th International Conference On. IEEE. pp.71–84.
- Cockayne, D.G. and Richardson, L., 2017. Queering code/space: the co-production of socio-sexual codes and digital technologies. *Gender, Place & Culture*, pp.1–17.
- Cohen, F., 1989. Computational aspects of computer viruses. *Computers & Security*, 8(4), pp.325–344.
- Cohen, F., 1991. *A Case for Benevolent Viruses*. [online] Pittsburgh, PA: ASP Press. Available at: <<https://web.archive.org/web/20190603082120/https://pdfs.semanticscholar.org/24f3/b65261c7ac1795ef7a3cb0096d6269a12753.pdf>> [Accessed 3 Jun. 2018].

- Cohen, F.B., 1994. *It's Alive! The New Breed of Living Computer Programs*. New York: John Wiley & Sons.
- Coles-Kemp, L. and Hansen, R.R., 2017. Walking the Line: The Everyday Security Ties that Bind. In: T. Tryfonas, ed. *Human Aspects of Information Security, Privacy and Trust: 5th International Conference, HAS 2017, Held as Part of HCI International 2017, Vancouver, BC, Canada, July 9-14, 2017, Proceedings*. [online] Cham: Springer International Publishing. pp.464-480. Available at: <https://doi.org/10.1007/978-3-319-58460-7_32>.
- Collier, J., 2018. Cyber Security Assemblages: A Framework for Understanding the Dynamic and Contested Nature of Security Provision. *Politics and Governance*, 6(2), pp.13-21.
- Connolly, W.E., 2017. *Facing the Planetary: Entangled Humanism and the Politics of Swarming*. Durham, N.C.: Duke University Press.
- Cooper, G., Hine, C., Rachel, J. and Woolgar, S., 1995. Ethnography and Human-Computer Interaction. In: P.J. Thomas, ed. *The Social and Interactional Dimensions of Human-Computer Interfaces*. pp.11-36.
- Crampton, J.W., 2015. Collect it all: national security, Big Data and governance. *GeoJournal*, 80(4), pp.519-531.
- Crampton, J.W., 2016. Assemblage of the vertical: commercial drones and algorithmic life. *Geogr. Helv.*, 71(2), pp.137-146.
- Crampton, J.W., 2019. Geopolitics. In: J. Ash, R. Kitchin and A. Leszczynski, eds. *Digital Geographies*. London: SAGE. pp.281-290.
- Crandall, J.R., Ensafi, R., Forrest, S., Ladau, J. and Shebaro, B., 2009. The Ecology of Malware. Proceedings of the 2008 Workshop on New Security Paradigms. ACM. pp.99-106.
- CrySyS Malware Intelligence Team, 2013. *Miniduke: Indicators*. [online] Budapest, Hungary: Laboratory of Cryptography and System Security. Available at: <https://web.archive.org/web/20170703172125/https://www.crysys.hu/miniduke/miniduke_indicators_public.pdf> [Accessed 3 Jul. 2017].
- Cyberpolice Ukraine, 2017. *Кіберполіцією попередньо встановлено, що перші вірусні атаки на українські компанії могли виникнути через вразливості ПІЗ М.Е.doc*. [online] Twitter. Available at: <<https://web.archive.org/web/20170627194707/https://twitter.com/CyberpoliceUA/status/879772963658235904>> [Accessed 27 Jun. 2017].
- Dalakov, G., 2018. *First computer virus of Bob Thomas*. [online] History Computer. Available at: <<http://archive.is/92zCg>> [Accessed 19 Jun. 2018].
- Dang, B., Gazet, A. and Bachaalany, E., 2014. *Practical Reverse Engineering: X86, X64, ARM, Windows Kernel, Reversing Tools, and Obfuscation*. Indianapolis, Ind.: Wiley.
- Deleuze, G., 1992. Postscript on the Societies of Control. *October*, 59(Winter 1992), pp.3-7.
- Deleuze, G. and Guattari, F., 2013. *A Thousand Plateaus: Capitalism and Schizophrenia*. Bloomsbury revelations. Translated by B. Massumi. London: Bloomsbury Academic.
- Democratic National Committee v. Russian Federation, et al.* (U.S. District Court, Southern District of New York) Available at: <<https://perma.cc/UE7Z-RYLF>> [Accessed 4 Nov. 2018].
- Denning, T., 2014. *Human-Centric Security and Privacy for Emerging Technologies*. [Doctoral Thesis] University of Washington. Available at: <https://web.archive.org/web/20190603084600/https://digital.lib.washington.edu/researchworks/bitstream/handle/1773/26958/Denning_washington_0250E_13497.pdf> [Accessed 3 Jun. 2019].
- Department for Health and Social Care, 2018. *Securing cyber resilience in health and care: Progress update October 2018*. [Implementation Update] London. Available at: <https://web.archive.org/web/20181015105106/https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/747464/securing-cyber-resilience-in-health-and-care-september-2018-update.pdf> [Accessed 15 Oct. 2018].

- Derrida, J., 1981. Plato's Pharmacy. Translated by B. Johnson. In: *Disseminations*. Chicago: University of Chicago Press. pp.61–172.
- Des Chene, D., 2001. *Spirits and Clocks: Machine and Organism in Descartes*. Ithaca, N.Y.: Cornell University Press.
- Descartes, R., 1972. *Treatise on Man*. Translated by T. Hall. Cambridge, Mass.: Harvard University Press.
- Develennes, C. and Dillet, B., 2018. Questioning New Materialisms: An Introduction. *Theory, Culture & Society*, 35(7–8), pp.5–20.
- Dewar, R., 2017. *Cyberweapons: Capability, Intent and Context in Cyberdefense*. [online] Center for Security Studies (CSS), ETH Zurich. Available at: <<https://doi.org/10.3929/ethz-b-000210449>>.
- Dewdney, A.K., 1984. Computer Recreations: In the game called Core War hostile programs engage in a battle of bits. *Scientific American*, 250(5), pp.14–22.
- Dewdney, A.K., 1985. Computer Recreations: A Core War bestiary of viruses, worms and other threats to computer memories. *Scientific American*, 252(3), p.14.
- Digital Evolution Laboratory, 2018. *Devolab | Digital Evolution Laboratory*. [online] Devolab | Digital Evolution Laboratory. Available at: <<https://web.archive.org/web/20181119174258/http://devolab.msu.edu/>> [Accessed 19 Nov. 2018].
- Digital Life Laboratory, 2009. *The Digital Life Lab at Caltech*. [online] Available at: <<https://web.archive.org/web/20090503111727/http://dllab.caltech.edu:80/>> [Accessed 19 Nov. 2018].
- Dillon, M. and Lobo-Guerrero, L., 2008. Biopolitics of security in the 21st century: an introduction. *Review of International Studies*, 34(02), pp.265–292.
- Dix, A., 2009. Human-computer interaction. In: L. Liu and M.T. Özsu, eds. *Encyclopedia of Database Systems*. New York: Springer. pp.1327–1331.
- Dodge, M., 2001. Cybergeography. *Environment and Planning B: Planning and Design*, 28(1), pp.1–2.
- Doward, J. and Townsend, M., 2017. Cyber-attack sparks bitter political row over NHS spending. *The Observer*. [online] 14 May. Available at: <<https://web.archive.org/web/20181005051718/https://www.theguardian.com/technology/2017/may/13/cyber-attack-on-nhs-sparks-bitter-election-battle>> [Accessed 5 Oct. 2018].
- Dowling, R., Lloyd, K. and Suchet-Pearson, S., 2016. Qualitative methods II: 'More-than-human' methodologies and/in praxis. *Progress in Human Geography*, 41(6), pp.823–831.
- Ducklin, P., 2017. New Petya ransomware: everything you wanted to know (but were afraid to ask). [Sophos] *Naked Security*. Available at: <<https://web.archive.org/web/20190118091618/https://nakedsecurity.sophos.com/2017/06/28/new-petya-ransomware-all-you-wanted-to-know-but-were-afraid-to-ask/>> [Accessed 18 Jan. 2019].
- Duggan, M., 2017. Questioning “digital ethnography” in an era of ubiquitous computing. *Geography Compass*, 11(5), p.e12313.
- Dunn Cavelt, M., 2013. From Cyber-Bombs to Political Fallout: Threat Representations with an Impact in the Cyber-Security Discourse. *International Studies Review*, 15(1), pp.105–122.
- Dwyer, A., 2018. The NHS cyber-attack: A look at the complex environmental conditions of WannaCry. *RAD Magazine*, 44, pp.25–26.
- Egloff, F., J. and Wenger, A., 2019. *Public Attribution of Cyber Incidents*. [CSS Analyses in Security Policy] Zürich: CSS Zürich. Available at: <<https://web.archive.org/web/20190511085131/http://www.css.ethz.ch/content/dam/ethz/special-interest/gess/cis/center-for-securities-studies/pdfs/CSSAnalyse244-EN.pdf>>.

- E-ISAC, 2016. Analysis of the cyber attack on the Ukrainian power grid: Defence Use Case. *Electricity Information Sharing and Analysis Center (E-ISAC)*. [online] Available at: <https://web.archive.org/web/20190121100932/https://ics.sans.org/media/E-ISAC_SANS_Ukraine_DUC_5.pdf> [Accessed 21 Jan. 2019].
- Elden, S., 2016. *Foucault's Last Decade*. Cambridge, UK: Polity Press.
- Elwood, S. and Leszczynski, A., 2018. Feminist digital geographies. *Gender, Place & Culture*, pp.1-16.
- Esposito, E., 2017. An Ecology of Differences: Communication, the web, and the question of borders. In: E. Hörl, ed. *General Ecology: The New Ecological Paradigm*. London: Bloomsbury Academic. pp.285-301.
- Falliere, N., O'Murchu, L. and Chien, E., 2010. *W32.Stuxnet Dossier*. [online] Symantec. Available at: <https://web.archive.org/web/20101003183230/http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf> [Accessed 3 Oct. 2010].
- Falliere, N., O'Murchu, L. and Chien, E., 2011. *W32.Stuxnet Dossier*. [online] Symantec. Available at: <https://web.archive.org/web/20181003105114/http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf> [Accessed 3 Oct. 2018].
- Fazi, M.B., 2018. *Contingent Computation: Abstraction, Experience, and Indeterminacy in Computational Aesthetics*. London: Rowman & Littlefield International.
- Ferbrache, D., 1992. *A Pathology of Computer Viruses*. London: Springer-Verlag.
- Flemming, J., 2018. *Speech at CyberUK18*. Manchester. Available at: <<https://web.archive.org/web/20180807145923/https://www.gchq.gov.uk/sites/default/files/Director%20GCHQ%20speech%20CyberUK%202018.pdf>> [Accessed 7 Aug. 2018].
- Folkers, A., 2017. Existential provisions: The technopolitics of public infrastructure. *Environment and Planning D: Society and Space*, 35(5), pp.855-874.
- Forbes, N., 2004. *Imitation of Life: How Biology is Inspiring Computing*. Cambridge, Mass.: MIT Press.
- Forman, P.J., 2018. Circulations beyond nodes: (in)securities along the pipeline. *Mobilities*, 13(2), pp.231-245.
- Foucault, M., 1991. *Discipline and Punish: The Birth of the Prison*. Translated by A. Sheridan. Harmondsworth: Penguin Books.
- Foucault, M., 2003a. *Society Must Be Defended: Lectures at the Collège de France, 1975-1976*. Translated by D. Macey. London: Penguin Books.
- Foucault, M., 2003b. *The Birth of the Clinic: An Archaeology of Medical Perception*. 3rd Edition ed. Translated by A. Sheridan. London: Routledge.
- Foucault, M., 2008. *The Birth of Biopolitics: Lectures at the Collège de France, 1978-1979*. Basingstoke: Palgrave Macmillan.
- Foucault, M., 2016. *Abnormal: Lectures at the Collège de France, 1974-1975*. Translated by G. Burchell. London: Verso.
- Fransen, F., Smulders, A. and Kerkdijk, R., 2015. Cyber security information exchange to gain insight into the effects of cyber threats and incidents. *e & i Elektrotechnik und Informationstechnik*, 132(2), pp.106-112.
- F-Secure, 2015. *CozyDuke*. [online] F-Secure. Available at: <<https://perma.cc/HZ6Y-P7KJ>>.
- Galloway, A.R., 2004. *Protocol: How Control Exists after Decentralization*. Leonardo (Series). Cambridge, Mass.: MIT Press.
- Geers, K., 2009. The Cyber Threat to National Critical Infrastructures: Beyond Theory. *Information Security Journal: A Global Perspective*, 18(1), pp.1-7.

- Gerlach, J. and Jellis, T., 2015. Guattari: Impractical philosophy. *Dialogues in Human Geography*, 5(2), pp.131–148.
- Gibson, W., 1984. *Neuromancer*. London: Grafton.
- Gibson, W., 1988. *Burning Chrome*. London: Grafton.
- Gibson-Graham, J.K., 2011. A feminist project of belonging for the Anthropocene. *Gender, Place and Culture*, 18(01), pp.1–21.
- Glouftsiou, G., 2018. Governing circulation through technology within EU border security practice-networks. *Mobilities*, 13(2), pp.185–199.
- de Goede, M., 2014. The Politics of Privacy in the Age of Preemptive Security. *International Political Sociology*, 8(1), pp.100–104.
- de Goede, M., 2018. Counter-Terrorism Financing Assemblages After 9/11. In: C. King, C. Walker and J. Gurulé, eds. *The Palgrave Handbook of Criminal and Terrorism Financing Law*. [online] Cham, Switzerland: Springer International Publishing, pp.755–779. Available at: <https://doi.org/10.1007/978-3-319-64498-1_31>.
- de Goede, M., Leander, A. and Sullivan, G., 2016. Introduction: The politics of the list. *Environment and Planning D: Society and Space*, 34(1), pp.3–13.
- de Goede, M. and Randalls, S., 2009. Precaution, Preemption: Arts and Technologies of the Actionable Future. *Environment and Planning D: Society and Space*, 27(5), pp.859–878.
- Goodbun, J., 2010. Gregory Bateson's Ecological Aesthetics—An Addendum to Urban Political Ecology. *Field Journal*, 4(1), pp.35–46.
- Goodin, D., 2017. >10,000 Windows computers may be infected by advanced NSA backdoor. [online] Arstechnica. Available at: <<https://web.archive.org/web/20181009061752/https://arstechnica.com/information-technology/2017/04/10000-windows-computers-may-be-infected-by-advanced-nsa-backdoor/>> [Accessed 9 Oct. 2018].
- Gordon, R., 2012. *Ordering Networks: Motorways and the Work of Managing Disruption*. [Doctoral Thesis] Durham University. Available at: <http://etheses.dur.ac.uk/6347/1/Ordering_Networks_Motorways_and_the_Work_of_Managing_Disruption_-_Rachel_Gordon.pdf?DDD14+>.
- Goulden, M., Greiffenhagen, C., Crowcroft, J., McAuley, D., Mortier, R., Radenkovic, M. and Sathiaselan, A., 2017. Wild interdisciplinarity: ethnography and computer science. *International Journal of Social Research Methodology*, 20(2), pp.137–150.
- Graham, M., 2013. Geography/internet: ethereal alternate dimensions of cyberspace or grounded augmented realities? *The Geographical Journal*, 179(2), pp.177–182.
- Graham, M. and Shelton, T., 2013. Geography and the future of big data, big data and the future of geography. *Dialogues in Human Geography*, 3(3), pp.255–261.
- GReAT, 2017. *Wannacry and Lazarus Group – the missing link?* [online] Securelist. Available at: <<https://web.archive.org/web/20190115161606/https://securelist.com/wannacry-and-lazarus-group-the-missing-link/78431/>> [Accessed 15 Jan. 2019].
- Greenberg, A., 2018. *The Untold Story of NotPetya, The Most Devastating Cyberattack in History*. [online] WIRED. Available at: <<https://web.archive.org/web/20180822121630/https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/>> [Accessed 22 Aug. 2018].
- Greenhough, B., 2010. Vitalist geographies: life and the more-than-human. In: B. Anderson and P. Harrison, eds. *Taking Place: Non-Representational Theories and Geography*. pp.37–54.
- Greenhough, B., 2012. Where species meet and mingle: endemic human-virus relations, embodied communication and more-than-human agency at the Common Cold Unit 1946–90. *Cultural Geographies*, 19(3), pp.281–301.
- Greenhough, B., 2014. More-than-human Geographies. In: A. Passi, N. Castree, R. Lee, S. Radcliffe, R. Kitchin, V. Lawson and C. Withers, eds. *The Sage Handbook of Progress in Human Geography*. London: Sage Publications, pp.94–119.

- Greenhough, B., Dwyer, A., Grenyer, R., Hodgetts, T., McLeod, C. and Lorimer, J., 2018. Unsettling antibiosis: how might interdisciplinary researchers generate a feeling for the microbiome and to what effect? *Palgrave Communications*, [online] 4(149). Available at: <<https://doi.org/10.1057/s41599-018-0196-3>>.
- Gregory, D., 2011. From a View to a Kill: Drones and Late Modern War. *Theory, Culture & Society*, 28(7-8), pp.188-215.
- Grosz, E., 2008. *Chaos, Territory, Art: Deleuze and the Framing of the Earth*. New York: Columbia University Press.
- Guattari, F., 2014. *The Three Ecologies*. Bloomsbury revelations. London: Bloomsbury.
- Guilfanov, I., 2014. *The story of IDA Pro*. [online] Available at: <<https://web.archive.org/web/20181228121031/https://www.youtube.com/watch?v=hLBlck1lTU8>> [Accessed 28 Dec. 2018].
- Guinier, D., 1991. Computer “virus” identification by neural networks: An artificial intelligence connectionist implementation naturally made to work with fuzzy information. *ACM SIGSAC Review*, 9(4), pp.49-59.
- Halpern, O., 2015. *Beautiful Data: A History of Vision and Reason since 1945*. Durham, N.C.: Duke University Press.
- Haraway, D., 1988. Situated knowledges: The science question in feminism and the privilege of partial perspective. *Feminist studies*, 14(3), pp.575-599.
- Haraway, D., 1991. *Simians, Cyborgs, and Women: The Reinvention of Nature*. New York: Routledge.
- Haraway, D.J., 2016. *Staying with the Trouble: Making Kin in the Chthulucene*. Experimental Futures. Durham, N.C.: Duke University Press.
- Harman, G., 2002. *Tool-Being: Heidegger and the Metaphysics of Objects*. Chicago: Open Court.
- Harrison, P., 2007. “How Shall I Say it ... ?” Relating the Nonrelational. *Environment and Planning A: Economy and Space*, 39(3), pp.590-608.
- Harrison, P., 2015. After Affirmation, or, Being a Loser: On Vitalism, Sacrifice, and Cinders. *GeoHumanities*, 1(2), pp.285-306.
- Hatfield, G., 2018. René Descartes. In: E.N. Zalta, ed. *The Stanford Encyclopedia of Philosophy*, Summer 2018. [online] Metaphysics Research Lab, Stanford University. Available at: <<https://plato.stanford.edu/archives/sum2018/entries/descartes/>>.
- Hayden, M.V., 2011. The Future of Things “Cyber”. *Strategic Studies Quarterly*, 5(1), pp.3-7.
- Hayles, K., 1990. *Chaos Bound: Orderly Disorder in Contemporary Literature and Science*. Ithaca, N.Y.: Cornell University Press.
- Hayles, N.K., 1993. Virtual Bodies and Flickering Signifiers. *October*, 66, pp.69-91.
- Hayles, N.K., 1999. *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature, and Informatics*. Chicago: University of Chicago Press.
- Hayles, N.K., 2005. *My Mother Was a Computer: Digital Subjects and Literary Texts*. Chicago: University of Chicago Press.
- Hayles, N.K., 2008. *Electronic Literature: New Horizons for the Literary*. Notre Dame, Indiana: University of Notre Dame Press.
- Hayles, N.K., 2014. Speculative Aesthetics and Object-Oriented Inquiry (OOI). *Speculations: A Journal of Speculative Realism*, 5, pp.158-79.
- Hayles, N.K., 2017. *Unthought: The Power of the Cognitive Nonconscious*. Chicago: University of Chicago Press.
- Hayles, N.K., 2018. *Cybersemiosis: Meaning-Making in Humans, Nonhumans and Computational Media*. [online] Available at: <<https://perma.cc/K7ZE-73Y3>>.
- Hayles, N.K., 2019. Can Computers Create Meanings? A Cyber/Bio/Semiotic Perspective. *Critical Inquiry*.

- Heidegger, M., 1977. *The Question Concerning Technology, and Other Essays*. First ed. Harper colophon books. New York: Harper & Row.
- Helmreich, S., 2000. Flexible Infections: Computer Viruses, Human Bodies, Nation-States, Evolutionary Capitalism. *Science, Technology & Human Values*, 25(4), pp.472–491.
- Helmreich, S., 2004. The Word for World is Computer: Simulating second natures in artificial life. In: M.N. Wise, ed. *Growing Explanations: Historical Perspectives on the Sciences of Complexity*. Durham, N.C.: Duke University Press. pp.275–300.
- Herbert, S., 2000. For ethnography. *Progress in Human Geography*, 24(4), pp.550–568.
- Highland, H., 1987. Computer viruses and sudden death! *Computers & Security*, 6(1), pp.8–10.
- Hinchliffe, S., Allen, J., Lavau, S., Bingham, N. and Carter, S., 2013. Biosecurity and the topologies of infected life: from borderlines to borderlands. *Transactions of the Institute of British Geographers*, 38(4), pp.531–543.
- Hinchliffe, S., Butcher, A. and Rahman, M.M., 2018. The AMR problem: demanding economies, biological margins, and co-producing alternative strategies. *Palgrave Communications*, [online] 4(142). Available at: <<https://doi.org/10.1057/s41599-018-0195-4>>.
- Hoijsink, M. and Leese, M. eds., 2019. *Technology and Agency in International Relations*. Emerging Technologies, Ethics and International Affairs. Abingdon, UK: Routledge.
- Holt, T.J., Strumsky, D., Smirnova, O. and Kilger, M., 2012. Examining the social networks of malware writers and hackers. *International Journal of Cyber Criminology*, 6(1), pp.891–903.
- Hopkins, N., 2011. Stuxnet attack forced Britain to rethink the cyber war. *The Guardian*. [online] 30 May. Available at: <<https://web.archive.org/web/20181031121845/https://www.theguardian.com/politics/2011/may/30/stuxnet-attack-cyber-war-iran>> [Accessed 31 Oct. 2018].
- Hörl, E., 2017. *General Ecology: The New Ecological Paradigm*. London: Bloomsbury Academic.
- Hruska, J., 1990. *Computer Viruses and Anti-Virus Warfare*. London: Ellis Horwood.
- Huawei Cyber Security Evaluation Centre (HCSEC) Oversight Board, 2019. *Huawei Cyber Security Evaluation Centre (HCSEC) Oversight Board Annual Report 2019*. [online] Huawei Cyber Security Evaluation Centre. Available at: <https://web.archive.org/web/20190401042734/https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/790270/HCSEC_OversightBoardReport-2019.pdf> [Accessed 1 Apr. 2019].
- Hughes, L.A. and DeLone, G.J., 2007. Viruses, Worms, and Trojan Horses: Serious Crimes, Nuisance, or Both? *Social Science Computer Review*, 25(1), pp.78–98.
- Hui, Y., 2019. *Recursivity and Contingency*. Media philosophy. London: Rowman & Littlefield International, Ltd.
- Hung, G. and Joven, M., 2017. *Petya's Master Boot Record Infection*. [Fortinet Threat Research] Available at: <<https://web.archive.org/web/20181103141010/https://www.fortinet.com/blog/threat-research/petya-s-master-boot-record-infection.html>> [Accessed 3 Nov. 2018].
- Hutchins, E.M., Cloppert, M.J. and Amin, R.M., 2011. Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. *Leading Issues in Information Warfare & Security Research*, 1(1), p.80.
- Invincea, 2017. *Next-Generation Antivirus for Government Organizations*. [online] Invincea. Available at: <<https://www.invincea.com/solutions/next-generation-antivirus-for-government-organizations/>> [Accessed 7 Mar. 2018].
- Ivanov, A. and Mamedov, O., 2017. *ExPetri/Petya/NotPetya is a Wiper, Not Ransomware*. [online] Secure List. Available at: <<https://web.archive.org/web/20170629085447/https://securelist.com/expetripetyano-tpetya-is-a-wiper-not-ransomware/78902/>> [Accessed 29 Jun. 2017].

- Jacobsen, A., 2015. *The Pentagon's Brain: An Uncensored History of DARPA, America's Top-secret Military Research Agency*. New York: Little, Brown, and Company.
- Jellis, T., Gerlach, J. and Dewsbury, J.-D. eds., 2019. *Why Guattari? A Liberation of Cartographies, Ecologies and Politics*. Routledge Series in Human Geography. Abingdon, UK: Routledge.
- Johnston, J., 2008. *The Allure of Machinic Life: Cybernetics, Artificial Life, and the new AI*. Cambridge, Mass.: MIT Press.
- Johnston, J., 2009a. Mutant and Viral: Artificial Evolution and Software Ecology. In: J. Parikka and T.D. Sampson, eds. *The Spam Book: On Viruses, Porn, and other Anomalies from the Dark Side of Digital Culture*. Cresskill, NJ: Hampton Press. pp.23–38.
- Johnston, J., 2009b. *Technological Turf Wars: A Case Study of the Computer Antivirus Industry*. Philadelphia, Pa.: Temple University Press.
- Jordan, T. and Taylor, P., 2008. A sociology of hackers. *The Sociological Review*, 46(4), pp.757–780.
- Kaiser, R., 2012. Reassembling the Event: Estonia's 'Bronze Night'. *Environment and Planning D: Society and Space*, 30(6), pp.1046–1063.
- Kaiser, R., 2015. The birth of cyberwar. *Political Geography*, 46, pp.11–20.
- Kamuk, V., 2009. Here's looking at you Kido. *Network Security*, 2009(3), pp.6–8.
- Kello, L., 2013. The Meaning of the Cyber Revolution: Perils to Theory and Statecraft. *International Security*, 38(2), pp.7–40.
- Kello, L., 2017. *The Virtual Weapon and International Order*. New Haven: Yale University Press.
- Kinsley, S., 2014. The matter of 'virtual' geographies. *Progress in Human Geography*, 38(3), pp.364–384.
- Kinsley, S., 2015. Memory programmes: the industrial retention of collective life. *cultural geographies*, 22(1), pp.155–175.
- Kitchin, R., 2014. Big Data, new epistemologies and paradigm shifts. *Big Data & Society*. [online] Available at: <<https://doi.org/10.1177/2053951714528481>> [Accessed 18 Jan. 2019].
- Kitchin, R., 2017. Thinking critically about and researching algorithms. *Information, Communication & Society*, 20(1), pp.14–29.
- Kitchin, R. and Dodge, M., 2011. *Code/Space: Software and Everyday Life*. Cambridge, Mass.: MIT Press.
- Kline, S., Dyer-Witheford, N. and De Peuter, G., 2003. *Digital Play: The Interaction of Technology, Culture, and Marketing*. Montreal, Canada: McGill-Queen's Press.
- Klinke, I., 2019. Vitalist temptations: Life, earth and the nature of war. *Political Geography*, 72, pp.1–9.
- Kneale, J., 1999. The virtual realities of technology and fiction: reading William Gibson's cyberspace. In: M. Crang, P. Crang and J. May, eds. *Virtual Geographies: Bodies, Space and Relations*. Oxford: Routledge.
- Kosslyn, J., 2018. *The Internet Needs More Friction*. [online] Motherboard. Available at: <https://web.archive.org/web/20181118163109/https://motherboard.vice.com/en_us/article/3k9q33/the-internet-needs-more-friction> [Accessed 25 Nov. 2018].
- Krebs, B., 2010. Experts Warn of New Windows Shortcut Flaw. *Krebs on Security*. Available at: <<https://web.archive.org/web/20100718101510/http://krebsonsecurity.com/2010/07/experts-warn-of-new-windows-shortcut-flaw>> [Accessed 18 Jul. 2010].
- Krebs, B., 2016. Who Makes the IoT Things Under Attack? — Krebs on Security. *Krebs on Security*. Available at: <<https://krebsonsecurity.com/2016/10/who-makes-the-iot-things-under-attack/>> [Accessed 23 Dec. 2016].
- Kryptos Logic, 2017. *WannaCry: Two Weeks and 16 Million Averted Ransoms Later*. [online] Available at:

- <<https://web.archive.org/web/20181101162723/https://blog.kryptoslogic.com/malware/2017/05/29/two-weeks-later.html>> [Accessed 1 Nov. 2018].
- Kuus, M., 2013. Foreign Policy and Ethnography: A Sceptical Intervention. *Geopolitics*, 18(1), pp.115–131.
- Landau, S., Lin, H.S. and Bellovin, S.M., 2017. Limiting the undesired impact of cyber weapons: technical requirements and policy implications. *Journal of Cybersecurity*, 3(1), pp.59–68.
- Landler, M. and Markoff, J., 2007. In Estonia, what may be the first war in cyberspace. *The New York Times*. [online] 28 May. Available at: <<https://web.archive.org/web/20190401104706/https://www.nytimes.com/2007/05/28/business/worldbusiness/28iht-cyberwar.4.5901141.html>> [Accessed 1 Apr. 2019].
- Langner, R., 2011. Stuxnet: Dissecting a Cyberwarfare Weapon. *Security & Privacy, IEEE*, 9(3), pp.49–51.
- Latour, B., 1987. *Science in Action: How to Follow Scientists and Engineers through Society*. Cambridge, Mass.: Harvard University Press.
- Latour, B., 1996. On actor-network theory: a few clarifications. *Soziale Welt*, 47(4), pp.369–381.
- Latour, B., 2013. *An Inquiry Into Modes of Existence: An Anthropology of the Moderns*. Translated by C. Porter. London: Harvard University Press.
- Latour, B., 2017. *Facing Gaia: Eight Lectures on the New Climatic Regime*. Cambridge, UK: John Wiley & Sons.
- Latour, B. and Woolgar, S., 2013. *Laboratory Life: The Construction of Scientific Facts*. Princeton University Press.
- Law, J., 1994. *Organising Modernity: Social Ordering and Social Theory*. Oxford: Blackwell.
- Law, J., 1999. After ANT: complexity, naming and topology. *The Sociological Review*, 47(S1), pp.1–14.
- Law, J., 2004. *After Method: Mess in Social Science Research*. Abingdon: Routledge.
- Lee, D., 2017. WannaCry ransomware bitcoins move from online wallets. *BBC News*. [online] 3 Aug. Available at: <<http://web.archive.org/web/20170808062123/http://www.bbc.co.uk/news/technology-40811972>> [Accessed 6 Sep. 2017].
- Lee, I. and Sokolsky, O., 2010. Medical cyber physical systems. In: *Proceedings of the 47th Design Automation Conference*. ACM. pp.743–748.
- Lehtiö, A., 2015. *The Dukes: 7 years of Russian cyberespionage*. [online] F-Secure. Available at: <https://web.archive.org/web/20181012022458/https://www.f-secure.com/documents/996508/1030745/dukes_whitepaper.pdf> [Accessed 12 Oct. 2018].
- Leibniz, G.W., 2001. *The Monadology*. Translated by R. Latta. Blacksburg, VA: Virginia Tech.
- Leiner, B.M., Cerf, V.G., Clark, D.D., Kahn, R.E., Kleinrock, L., Lynch, D.C., Postel, J., Roberts, L.G. and Wolff, S., 2009. A brief history of the Internet. *ACM SIGCOMM Computer Communication Review*, 39(5), pp.22–31.
- Leszczynski, A., 2017. Digital methods I: Wicked tensions. *Progress in Human Geography*, 42(3), pp.473–481.
- Leyden, J., 2014. *CONFIRMED: Sophos shifting threat response work to India*. [online] The Register. Available at: <https://web.archive.org/web/20181230093743/https://www.theregister.co.uk/2014/06/04/sophos_moves_threat_response_ops/> [Accessed 30 Dec. 2018].
- Lieberman, T., 2016. *AtomBombing: A Code Injection that Bypasses Current Security Solutions*. [online] Available at: <<https://blog.ensilo.com/atombombing-a-code-injection-that-bypasses-current-security-solutions>> [Accessed 9 Oct. 2017].

- Lidzborski, N., 2009. *Conficker Worm: Patching is Not Fast Enough*. [online] Network Security Blog | Qualys, Inc. Available at: <<https://web.archive.org/web/20171104103643/https://blog.qualys.com/laws-of-vulnerabilities/2009/01/14/conficker-worm-patching-is-not-fast-enough>> [Accessed 4 Nov. 2017].
- Lin, Y., Bennett, J. and Haq, T., 2013. *In Turn, It's PDF Time*. [online] FireEye. Available at: <<https://web.archive.org/web/20181104150316/https://www.fireeye.com/blog/threat-research/2013/02/in-turn-its-pdf-time.html>> [Accessed 4 Nov. 2018].
- Lindsay, J.R., 2013. Stuxnet and the Limits of Cyber Warfare. *Security Studies*, 22(3), pp.365–404.
- Lindsay, J.R., 2015. Tipping the scales: the attribution problem and the feasibility of deterrence against cyberattack. *Journal of Cybersecurity*, 1(1), pp.53–67.
- Lorimer, J., 2012. Multinatural geographies for the Anthropocene. *Progress in Human Geography*, 36(5), pp.593–612.
- Lorimer, J., 2016. Gut Buddies: Multispecies Studies and the Microbiome. *Environmental Humanities*, 8(1), pp.57–76.
- Lorimer, J. and Driessen, C., 2014. Wild experiments at the Oostvaardersplassen: rethinking environmentalism in the Anthropocene. *Transactions of the Institute of British Geographers*, 39(2), pp.169–181.
- Lorimer, J., Hodgetts, T. and Barua, M., 2017. Animals' atmospheres. *Progress in Human Geography*, 43(1), pp.26–45.
- Lorimer, J., Hodgetts, T., Grenyer, R., Greenhough, B., McLeod, C. and Dwyer, A., 2019. Making the microbiome public: Participatory experiments with DNA sequencing in domestic kitchens. *Transactions of the Institute of British Geographers*, (44), pp.524–541.
- Lupovici, A., 2016. The “Attribution Problem” and the Social Construction of “Violence”: Taking Cyber Deterrence Literature a Step Forward. *International Studies Perspectives*, 17(3), pp.322–342.
- Lupton, D., 2018. How do data come to matter? Living and becoming with personal data. *Big Data & Society*. [online] Available at: <<https://doi.org/10.1177/2053951718786314>> [Accessed 2 Oct. 2018].
- Lusthaus, J., 2013. How organised is organised cybercrime? *Global Crime*, 14(1), pp.52–60.
- Lusthaus, J. and Varese, F., 2017. Offline and Local: The Hidden Face of Cybercrime. *Policing: A Journal of Policy and Practice*. [online] Available at: <<http://dx.doi.org/10.1093/polic/pax042>>.
- Machrone, B., 2009. *A Malware Glossary*. [online] Trend Micro on Wall Street Journal. Available at: <<https://web.archive.org/web/20181126112310/http://online.wsj.com/ad/article/protecting-malware>> [Accessed 26 Nov. 2018].
- MacKenzie, A., 2015. The production of prediction: What does machine learning want? *European Journal of Cultural Studies*, 18(4–5), pp.429–445.
- Madden, R., 2014. Animals and the Limits of Ethnography. *Anthrozoös*, 27(2), pp.279–293.
- Malanov, A., 2016. Five myths about machine learning in cybersecurity. *Securelist - Kaspersky Lab's cyberthreat research and reports*. Available at: <<https://web.archive.org/web/20180821143655/https://securelist.com/five-myths-about-machine-learning-in-cybersecurity/76351/>> [Accessed 21 Aug. 2018].
- MalwareBytes, 2018. *So You Want To Be A Malware Analyst - Malwarebytes Labs*. [online] Available at: <<https://web.archive.org/web/20180416003015/https://blog.malwarebytes.com/security-world/2012/09/so-you-want-to-be-a-malware-analyst/>> [Accessed 4 Jun. 2018].
- Massey, D.B., 2005. *For Space*. London: SAGE Publications.

- Massumi, B., 2010. The Future Birth of the Affective Fact: The Political Ontology of Threat. In: M. Gregg and G.J. Seigworth, eds. *The Affect Theory Reader*. Durham, N.C.: Duke University Press. pp.52–70.
- Maturana, H.R. and Varela, F.J., 1980. *Autopoiesis and Cognition: The Realization of the Living*. Boston: D. Reidel Pub. Co.
- Maurer, T., 2018. *Cyber Mercenaries*. Cambridge, UK: Cambridge University Press.
- Maynor, D., Nikolic, A., Olney, M. and Younan, Y., 2017. *The MeDoc Connection*. [online] Talos Intelligence. Available at: <[https://web.archive.org/web/20190108044330/https://www.blogger.com/comment-iframe.g?blogID=1029833275466591797&postID=6684040042737345889&blogspotRpcToken=121962#%07B%022color%022:%022rgb\(255,%020255,%020255\)%022,%022background-color%022:%022rgb\(38,%02040,%02042\)%022,%022unvisitedLinkColor%022:%022rgb\(255,%020255,%020255\)%022,%022fontFamily%022:%022Roboto,%020sans-serif%022%07D](https://web.archive.org/web/20190108044330/https://www.blogger.com/comment-iframe.g?blogID=1029833275466591797&postID=6684040042737345889&blogspotRpcToken=121962#%07B%022color%022:%022rgb(255,%020255,%020255)%022,%022background-color%022:%022rgb(38,%02040,%02042)%022,%022unvisitedLinkColor%022:%022rgb(255,%020255,%020255)%022,%022fontFamily%022:%022Roboto,%020sans-serif%022%07D)> [Accessed 8 Jan. 2019].
- McCormack, D.P., 2013. *Refrains for Moving Bodies: Experience and Experiment in Affective Spaces*. Durham, N.C.: Duke University Press.
- McCormack, D.P., 2016. The circumstances of post-phenomenological life worlds. *Transactions of the Institute of British Geographers*, 42, pp.2–13.
- McCormack, D.P., 2018. *Atmospheric Things: On the Allure of Elemental Envelopment*. Elements. Durham, N.C.: Duke University Press.
- McGrath, A., 1993. 'Beneath the skin': Australian citizenship, rights and aboriginal women. *Journal of Australian Studies*, 17(37), pp.99–114.
- McGraw, G., 2013. Cyber war is inevitable (unless we build security in). *Journal of Strategic Studies*, 36(1), pp.109–119.
- Meehan, K., Shaw, I.G.R. and Marston, S.A., 2013. Political geographies of the object. *Political Geography*, 33, pp.1–10.
- Meehan, K.M., Shaw, I.G. and Marston, S.A., 2014. The state of objects. *Political Geography*, 39(6), p.60–62.
- Mehta, N., 2017. @neelmehta. [online] Twitter. Available at: <<https://web.archive.org/web/20180618085900/https://twitter.com/neelmehta/status/864164081116225536>> [Accessed 18 Jun. 2018].
- Mermin, N.D., 2007. *Quantum Computer Science: An Introduction*. Cambridge: Cambridge University Press.
- Microsoft, 2008. *Microsoft Security Bulletin MS08-067 - Critical*. [online] Available at: <<https://web.archive.org/web/20190215074533/https://docs.microsoft.com/en-us/security-updates/securitybulletins/2008/ms08-067>> [Accessed 15 Feb. 2019].
- Miles, T., 1988. *Neuromancer*. Interplay Productions.
- Mistreanu, S., 2018. Life Inside China's Social Credit Laboratory. *Foreign Policy*. [online] 3 Apr. Available at: <<https://web.archive.org/web/20181128025127/https://foreignpolicy.com/2018/04/03/life-inside-chinas-social-credit-laboratory/>> [Accessed 28 Nov. 2018].
- Modderkolk, H., 2018. Dutch agencies provide crucial intel about Russia's interference in US-elections | De Volkskrant. Translated by L. Negrijn. *de Volkskrant*. [online] 25 Jan. Available at: <<https://web.archive.org/web/20181104194018/https://www.volkskrant.nl/wetenschap/dutch-agencies-provide-crucial-intel-about-russia-s-interference-in-us-elections-b4f8111b/>> [Accessed 4 Nov. 2018].
- Mol, A., 1999. Ontological politics. A word and some questions. *The Sociological Review*, 47(S1), pp.74–89.
- Mol, A., 2000. Pathology and the clinic: an ethnographic presentation of two atheroscleroses. In: M.M. Lock, A. Young and A. Cambrosio, eds. *Living and working with the new medical technologies: intersections of inquiry*, Cambridge studies in medical anthropology. Cambridge: Cambridge University Press. pp.82–102.

- Montfort, N., Baudoin, P., Bell, J., Bogost, I., Douglass, J., Marino, M.C., Mateas, M., Reas, C., Sample, M. and Vawter, N., 2012. *10 PRINT CHR \$(205.5+RND (1));: GOTO 10*. London: MIT Press.
- Morris, J., 2018. *Securing Finance, Mobilizing Risk: Money Cultures at the Bank of England*. Abingdon, UK: Routledge.
- Morse, A., 2017. *Investigation: WannaCry cyber attack and the NHS*. [online] London: National Audit Office. Available at: <<https://web.archive.org/web/20171027101548/https://www.nao.org.uk/wp-content/uploads/2017/10/Investigation-WannaCry-cyber-attack-and-the-NHS.pdf>> [Accessed 15 Oct. 2018].
- Morton, T., 2010. *The Ecological Thought*. Cambridge, Mass.: Harvard University Press.
- Mueller, M., Schmidt, A. and Kuerbis, B., 2013. Internet Security and Networked Governance in International Relations. *International Studies Review*, 15(1), pp.86–104.
- Mundie, D. and McIntire, D.M., 2013. An Ontology for Malware Analysis. 2013 Eighth International Conference on Availability, Reliability and Security (ARES). Regensburg, Germany: IEEE. pp.556–558.
- Myers, N., 2015. *Rendering Life Molecular: Models, Modelers, and Excitable Matter*. Experimental futures. Durham, N.C.: Duke University Press.
- Nakashima, E. and Harris, S., 2018. How the Russians hacked the DNC and passed its emails to WikiLeaks. *Washington Post*. [online] 13 Jul. Available at: <https://www.washingtonpost.com/world/national-security/how-the-russians-hacked-the-dnc-and-passed-its-emails-to-wikileaks/2018/07/13/af19a828-86c3-11e8-8553-a3ce89036c78_story.html> [Accessed 23 Jan. 2019].
- NCSC, 2017. *Protective DNS (PDNS)*. [online] National Cyber Security Centre. Available at: <https://www.ncsc.gov.uk/information/pdns#section_2> [Accessed 10 Jun. 2019].
- Newman, L.H., 2017. How an Accidental ‘Kill Switch’ Slowed Friday’s Massive Ransomware Attack. *WIRED*. [online] 13 May. Available at: <<https://web.archive.org/web/20181029125707/https://www.wired.com/2017/05/accidental-kill-switch-slowed-fridays-massive-ransomware-attack/>> [Accessed 29 Oct. 2018].
- Newman, L.H., 2018. WannaCry Hero’s New Legal Woes Spell Trouble for White Hat Hackers. *WIRED*. [online] 8 Jun. Available at: <<https://web.archive.org/web/20180608200133/https://www.wired.com/story/wannacry-hero-marcus-hutchins-new-legal-woes-white-hat-hackers/>> [Accessed 8 Jun. 2018].
- Noble, S.U., 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: NYU Press.
- Nominet, 2019. *Technical Whitepaper: Tracking the WannaCry Ransomware*. [online] Available at: <<https://web.archive.org/web/20190123142500/https://s3-eu-west-1.amazonaws.com/nominet-prod/wp-content/uploads/2017/11/21123044/WannaCry-Whitepaper1.pdf>> [Accessed 23 Jan. 2019].
- Nye, J., 2011. Nuclear Lessons for Cyber Security? *Strategic Studies Quarterly*, 5(4), pp.18–38.
- Obama, B., 2009. *Remarks by the President on Securing Our Nation’s Cyber Infrastructure*. [online] whitehouse.gov. Available at: <<https://web.archive.org/web/20190214113728/https://obamawhitehouse.archives.gov/the-press-office/remarks-president-securing-our-nations-cyber-infrastructure>> [Accessed 14 Feb. 2019].
- Obama, B., 2016. Barack Obama, Neural Nets, Self-Driving Cars, and the Future of the World. *WIRED*. [online] 24 Aug. Available at: <<https://www.wired.com/2016/10/president-obama-mit-joi-ito-interview/>> [Accessed 24 Oct. 2018].
- OED Online, 2018. ‘datum, n.’. [online] Oxford University Press. Available at: <<http://www.oed.com/view/Entry/47434>>.

- O'Grady, N., 2015. Data, interface, security: Assembling technologies that govern the future. *Geoforum*, 64, pp.130–137.
- O'Grady, N., 2016. Protocol and the post-human performativity of security techniques. *cultural geographies*, 23(3), pp.495–510.
- O'Grady, N., 2017. Mobility, circulation, and homeomorphism: data becoming risk information. In: S. Wittendorp and M. Leese, eds. *Security/Mobility: Politics of Movement*. Manchester: Manchester University Press. pp.74–93.
- O'Grady, N. and Dwyer, A., 2020. Cyber Security. In: *International Encyclopedia of Human Geography*, Second Edition. London: Elsevier.
- Oleg, K. and Ulasen, S., 2010. *Trojan-Spy.0485 And Malware-Cryptor.Win32.Inject.gen.2 Review*. [online] VirusBlokAda.p.7. Available at: <https://archive.org/details/new_rootkit_en> [Accessed 30 Oct. 2018].
- Ollmann, G., 2008. The evolution of commercial malware development kits and colour-by-numbers custom malware. *Computer Fraud & Security*, 2008(9), pp.4–7.
- O'Murchu, L., 2010. Last-minute paper: An indepth look into Stuxnet. [online] Virus Bulletin. Vancouver. Available at: <https://web.archive.org/web/20181030151502/https://www.virusbulletin.com/uploads/pdf/conference_slides/2010/OMurchu-VB2010.pdf> [Accessed 30 Oct. 2018].
- Osborne, G., 2015. *Chancellor's speech to GCHQ on cyber security - Speeches - GOV.UK*. [online] Available at: <<https://web.archive.org/web/20181021162216/https://www.gov.uk/government/speeches/chancellors-speech-to-gchq-on-cyber-security>> [Accessed 21 Oct. 2018].
- Oxford English Dictionary, 2016. 'cybernetics, n.'. [online] Available at: <<http://www.oed.com/view/Entry/46486>>.
- Parikka, J., 2007. *Digital Contagions: A Media Archaeology of Computer Viruses*. New York: Peter Lang.
- Parikka, J., 2009. Archives of Software - Malicious Code and the Aesthesis of Media Accidents. In: J. Parikka and T.D. Sampson, eds. *The Spam Book: On Viruses, Porn, and Other Anomalies from the Dark Side of Digital Culture*. Cresskill, NJ: Hampton Press.
- Parikka, J., 2010. Ethologies of Software Art: What Can a Digital Body of Code Do? In: S. Zepke and S. O'Sullivan, eds. *Deleuze and Contemporary Art*. Edinburgh: Edinburgh University Press. pp.116–132.
- Parikka, J., 2014. *The Anthrobscene*. Forerunners: Ideas First. Minneapolis, Minn.: University of Minnesota Press.
- Parikka, J., 2016. *Digital Contagions: A Media Archaeology of Computer Viruses*. Second Edition ed. New York: Peter Lang.
- Parikka, J. and Sampson, T.D., 2009. *The Spam Book: On Viruses, Porn, and Other Anomalies from the Dark Side of Digital Culture*. Cresskill, NJ: Hampton Press.
- Parisi, L., 2013. *Contagious Architecture: Computation, Aesthetics, and Space*. Technologies of Lived Abstraction. Cambridge, Mass.: The MIT Press.
- Parisi, L., 2017. Computational logic and ecological rationality. In: E. Hörl, ed. *General Ecology: The New Ecological Paradigm*. London: Bloomsbury Academic. pp.75–99.
- Parisi, L., 2019. Critical Computation: Digital Automata and General Artificial Thinking. *Theory, Culture & Society*, 36(2), pp.89–121.
- Parry, B. and Greenhough, B., 2018. *Bioinformation*. Cambridge: Polity Press.
- Patalong, F., 2009. Conficker-Countdown: Alles nur April, April? *Spiegel Online*. [online] 27 Mar. Available at: <<https://web.archive.org/web/20190222121833/http://www.spiegel.de/netzwelt/web/conficker-countdown-alles-nur-april-april-a-615002.html>> [Accessed 22 Feb. 2019].
- Paulson, W.R., 1988. *The Noise of Culture: Literary Texts in a World of Information*. Ithaca, N.Y.: Cornell University Press.

- Pietrek, M., 1994. *Peering Inside the PE: A Tour of the Win32 Portable Executable File Format*. [online] Available at: <<https://web.archive.org/web/20180103104757/https://msdn.microsoft.com/en-gb/library/ms809762.aspx>> [Accessed 3 Jan. 2018].
- Porras, P., Saïdi, H. and Yegneswaran, V., 2009. A Foray into Conficker's Logic and Rendezvous Points. In: *Proceedings of the 2Nd USENIX Conference on Large-scale Exploits and Emergent Threats: Botnets, Spyware, Worms, and More*, LEET'09. [online] Berkeley, CA, USA: USENIX Association.p.7. Available at: <<http://dl.acm.org/citation.cfm?id=1855676.1855683>>.
- Pötzsch, H., 2015. The emergence of iBorder: bordering bodies, networks, and machines. *Environment and Planning D: Society and Space*, 33(1), pp.101–118.
- Povinelli, E.A., 2016. *Geontologies: A Requiem to Late Liberalism*. Durham, N.C.: Duke University Press.
- Probyn, E., 2014. Women following fish in a more-than-human world. *Gender, Place & Culture*, 21(5), pp.589–603.
- Raiu, C., 2017. #WannaCry infection distribution by the Windows version. Worst bit - Windows 7 x64. The Windows XP count is insignificant. [online] Twitter. Available at: <<https://perma.cc/PW8B-RUV8>> [Accessed 11 Jan. 2018].
- Raiu, C., Soumenkov, I., Baumgartner, K. and Kamluk, V., 2013. *The MiniDuke Mystery: PDF 0-day Government Spy Assembler ox29A Micro Backdoor*. [online] Kaspersky Labs. Available at: <<https://perma.cc/6TK5-4CT3>> [Accessed 4 Nov. 2018].
- Ram, A., 2017. *Sophos gets boost from cyber-attacks*. [online] Financial Times. Available at: <<https://www.ft.com/content/700c9430-8cb6-3c63-b9de-ec9bdb1ddf9c>> [Accessed 6 Jun. 2018].
- Rid, T., 2013. *Cyber War Will Not Take Place*. London: C. Hurst & Co.
- Rid, T. and McBurney, P., 2012. Cyber-Weapons. *The RUSI Journal*, 157(1), pp.6–13.
- Rieffel, E. and Polak, W., 2011. *Quantum Computing: A Gentle Introduction*. Scientific and engineering computation. Cambridge: MIT Press.
- Roberts, S.L. and Elbe, S., 2016. Catching the flu: Syndromic surveillance, algorithmic governmentality and global health security. *Security Dialogue*, 48(1), pp.46–62.
- Roberts, T., 2014. From Things to Events: Whitehead and the Materiality of Process. *Environment and Planning D: Society and Space*, 32(6), pp.968–983.
- Robinson, N. and Martin, K., 2017. Distributed denial of government: the Estonian Data Embassy Initiative. *Network Security*, 2017(9), pp.13–16.
- Roe, E. and Greenhough, B., 2014. Experimental partnering: Interpreting improvisatory habits in the research field. *International Journal of Social Research Methodology*, 17(1), pp.45–57.
- Rose, G., 2015. Rethinking the geographies of cultural 'objects' through digital technologies: Interface, network and friction. *Progress in Human Geography*, 40(3), pp.334–351.
- Rose, G., 2016. *Visual Methodologies: An Introduction to Researching with Visual Materials*. London: SAGE.
- Rose, G., 2017. Posthuman Agency in the Digitally Mediated City: Exteriorization, Individuation, Reinvention. *Annals of the American Association of Geographers*, 107(4), pp.779–793.
- Rosenberg, M. and Nixon, R., 2017. Kaspersky Lab Antivirus Software Is Ordered Off U.S. Government Computers. *The New York Times*. [online] 13 Sep. Available at: <<https://web.archive.org/web/20190603100417/https://www.nytimes.com/2017/09/13/us/politics/kaspersky-lab-antivirus-federal-government.html>> [Accessed 3 Jun. 2019].
- Ross, A., 1990. Hacking away at the counterculture. *Postmodern Culture*, 1(1).
- Roth, A., 2019. Russia's great firewall: is it meant to keep information in – or out? *The Observer*. [online] 28 Apr. Available at:

- <<https://web.archive.org/web/20190512161036/https://www.theguardian.com/technology/2019/apr/28/russia-great-firewall-sovereign-internet-bill-keeping-information-in-or-out>> [Accessed 12 May 2019].
- Saldanha, A., 2006. Reontologising race: the machinic geography of phenotype. *Environment and Planning D: Society and Space*, 24(1), pp.9–24.
- Sampson, T.D., 2012. *Virality: Contagion Theory in the Age of Networks*. Minneapolis, Minn.: University of Minnesota Press.
- Sanger, D.E., 2012. *Confront and Conceal*. New York: Random House, Inc.
- Saussure, F. de, 2013. *Course in General Linguistics*. Translated by R. Harris. New York: Bloomsbury.
- Saxe, J. and Berlin, K., 2015. Deep neural network based malware detection using two dimensional binary program features. 10th International Conference on Malicious and Unwanted Software (MALWARE). pp.11–20.
- Saxe, J. and Berlin, K., 2017. eXpose: A Character-Level Convolutional Neural Network with Embeddings For Detecting Malicious URLs, File Paths and Registry Keys. *arXiv preprint arXiv:1702.08568*.
- Saxe, J. and Sanders, H., 2018. *Malware Data Science: Attack Detection and Attribution*. San Francisco: No Starch Press.
- Schmitt, M.N., 2013. *Tallinn Manual on the International Law Applicable to Cyber Warfare*. Cambridge: Cambridge University Press.
- Scientific American, 1997. When and how did the metaphor of the computer ‘virus’ arise? *Scientific American*. [online] 2 Sep. Available at: <<https://web.archive.org/web/20181127181005/https://www.scientificamerican.com/article/when-and-how-did-the-meta/>> [Accessed 27 Nov. 2018].
- Seaman, M.J., 2007. Becoming more (than) human: Affective posthumanisms, past and future. *Journal of Narrative Theory*, 37(2), pp.246–275.
- Securing a Common Future in Cyberspace*, 2018. [Panel] *World Economic Forum Annual Meeting*. Davos-Klosters, Switzerland. 24 Jan. Available at: <<https://www.weforum.org/events/world-economic-forum-annual-meeting-2018/sessions/securing-a-common-future-in-cyberspace>>.
- Shapiro, F.R., 1987. Etymology of the Computer Bug: History and Folklore. *American Speech*, 62(4), pp.376–378.
- Shaw, I.G., 2016a. *Predator Empire: Drone Warfare and Full Spectrum Dominance*. U of Minnesota Press.
- Shaw, I.G., 2016b. Scorched Atmospheres: The Violent Geographies of the Vietnam War and the Rise of Drone Warfare. *Annals of the American Association of Geographers*, 106(3), pp.688–704.
- Shearer, J., 2010. *W32.Stuxnet*. [online] Symantec. Available at: <<https://web.archive.org/web/20180808104436/https://www.symantec.com/security-center/writeup/2010-071400-3123-99>> [Accessed 8 Aug. 2018].
- Shevchenko, S. and Nish, A., 2017. *WanaCryptor Ransomworm*. [online] BAE Systems Threat Research Blog. Available at: <<https://web.archive.org/web/20180323192006/http://baesystemsai.blogspot.com:80/2017/05/wanacryptor-ransomworm.html>> [Accessed 31 Oct. 2018].
- Shoch, J.F. and Hupp, J.A., 1982. The ‘worm’ programs—early experience with a distributed computation. *Communications of the ACM*, 25(3), pp.172–180.
- Sikorski, M. and Honig, A., 2012. *Practical Malware Analysis: The Hands-On Guide to Dissecting Malicious Software*. San Francisco, Calif.: No Starch Press.
- Silomon, J.A.M. and Roeling, M.P., 2018. Assessing Opinions on Software as a Weapon in the Context of (Inter)national Security. In: M.L. Gavrilova, C.J.K. Tan and A. Sourin, eds. *Transactions on Computational Science XXXII: Special Issue on Cybersecurity and Biometrics*. [online] Berlin: Springer. pp.43–56. Available at: <https://doi.org/10.1007/978-3-662-56672-5_4>.

- Simon, S. and de Goede, M., 2015. Cybersecurity, Bureaucratic Vitalism and European Emergency. *Theory, Culture & Society*, 32(2), pp.79–106.
- Simondon, G., 2017. *On the Mode of Existence of Technical Objects*. Translated by C. Malaspina. and Translated by J. Rogove. Minneapolis, Minn.: Univocal.
- Singer, P.W. and Shachtman, N., 2011. The Wrong War: The Insistence on Applying Cold War Metaphors to Cybersecurity Is Misplaced and Counterproductive. *Brookings*. [online] 15 Aug. Available at: <<https://web.archive.org/web/20190222123924/https://www.brookings.edu/articles/the-wrong-war-the-insistence-on-applying-cold-war-metaphors-to-cybersecurity-is-misplaced-and-counterproductive/>> [Accessed 22 Feb. 2019].
- Skibell, R., 2002. The Myth of the Computer Hacker. *Information, Communication & Society*, 5(3), pp.336–356.
- Smart, W., 2018. *Lessons learned review of the WannaCry Ransomware Cyber Attack*. [Independent Report] London: Department for Health and Social Care. Available at: <<https://web.archive.org/web/20181015111225/https://www.england.nhs.uk/wp-content/uploads/2018/02/lessons-learned-review-wannacry-ransomware-cyber-attack-cio-review.pdf>> [Accessed 15 Oct. 2018].
- Smeets, M., 2018. A matter of time: On the transitory nature of cyberweapons. *Journal of Strategic Studies*, 41(1–2), pp.6–32.
- Smith, S. and Marchesini, J., 2007. *The Craft of System Security*. Boston: Pearson Education.
- Software: Practice and Experience, 1972. Computer recreations. *Software: Practice and Experience*, 2(1), pp.93–96.
- Sophos, 2016. *Download Reputation*. [online] Sophos Community. Available at: <<https://web.archive.org/web/20190603101415/https://community.sophos.com/kb/en-us/121319>> [Accessed 3 Jun. 2019].
- Sophos, 2017. *Information on Sophos Extensible List - Sophos Community*. [online] Available at: <<https://web.archive.org/web/20180105103258/https://community.sophos.com/kb/en-us/117936>> [Accessed 5 Jan. 2018].
- Sophos, 2018. *Sandbox Firewall | Sophos Next-Gen Firewall for Advanced Threat Defense and Protection | Sophos*. [online] Available at: <<https://web.archive.org/web/20181218170935/https://www.sophos.com/en-us/lp/sandstorm.aspx>> [Accessed 18 Dec. 2018].
- Sophos Knowledge Base, 2017. *New variant of Petya ransomware (also known as Petrwrap/PetyaWrap)*. [online] Sophos Community. Available at: <<https://web.archive.org/web/20170704200749/https://community.sophos.com/kb/en-us/127027>> [Accessed 4 Jul. 2017].
- Srinivasan, K., 2019. Remaking more-than-human society: Thought experiments on street dogs as “nature”. *Transactions of the Institute of British Geographers*, [online] 0(0). Available at: <<https://doi.org/10.1111/tran.12291>> [Accessed 21 Feb. 2019].
- Stahl, R., 2014. Life is War: The Rhetoric of Biomimesis and the Future Military. *Democratic Communiqué*, 26(2), pp.122 – 137.
- Stallings, W. and Brown, L., 2015. *Computer Security: Principles and Practice*. Harlow, UK: Pearson Education.
- Steffen, W., Broadgate, W., Deutsch, L., Gaffney, O. and Ludwig, C., 2015. The trajectory of the Anthropocene: The Great Acceleration. *The Anthropocene Review*, 2(1), pp.81–98.
- Stengers, I., 2010. Including Nonhumans in Political Theory: Opening the Pandora’s Box? In: B. Braun and S.J. Whatmore, eds. *Political Matter: Technoscience, Democracy, and Public Life*. pp.3–33.
- Stengers, I., 2015. *In Catastrophic Times: Resisting the Coming Barbarism*. Translated by A. Goffey. Open Humanities Press.
- Stevens, T., 2015. *Cyber Security and the Politics of Time*. Cambridge, UK: Cambridge University Press.

- Stevens, T., 2017. Cyberweapons: an emerging global governance architecture. *Palgrave Communications*, [online] 3(16102). Available at: <<https://doi.org/10.1057/palcomms.2016.102>>.
- Stevens, T., 2018. Global Cybersecurity: New Directions in Theory and Methods. *Politics and Governance*, 6(2), pp.1–4.
- Stiegler, B., 1998. *Technics and Time 1: The Fault of Epimetheus*. Stanford, Calif: Stanford University Press.
- Stiegler, B., 2011. Pharmacology of desire: Drive-based capitalism and libidinal dis-economy. *New Formations*, 72(72), pp.150–161.
- Stiegler, B., 2017. General ecology, economy, and organology. Translated by D. Ross. In: E. Hörl, ed. *General Ecology: The New Ecological Paradigm*. London: Bloomsbury Academic. pp.129–150.
- Stubbs, J. and Polityuk, P., 2017. Family firm in Ukraine says it was not responsible for cyber attack. *Reuters*. [online] 3 Jul. Available at: <<https://web.archive.org/web/20181123192207/https://www.reuters.com/article/us-cyber-attack-ukraine-software-idUSKBN19O2DK>> [Accessed 23 Nov. 2018].
- Suchman, L., 2007. *Human-Machine Reconfigurations: Plans and Situated Actions*. Cambridge, UK: Cambridge University Press.
- Sundaramurthy, S.C., McHugh, J. and Ou, X., 2014. An anthropological approach to studying CSIRTs. *IEEE Security & Privacy*, (12), pp.52–60.
- Symantec, 2013. *W32.Downadup*. [online] Symantec. Available at: <<https://web.archive.org/web/20190205101039/https://www.symantec.com/security-center/writeup/2008-112203-2408-99>> [Accessed 5 Feb. 2019].
- Symantec, 2015. “*Forkmeiamfamous*”: *Seaduke, latest weapon in the Duke armory*. [online] Symantec Security Response. Available at: <<https://perma.cc/58ST-CD5L?type=image>> [Accessed 8 Aug. 2016].
- Ször, P., 2005. *The Art of Computer Virus Research and Defense*. Upper Saddle River, NJ: Pearson Education.
- Taddeo, M., 2016. Just Information Warfare. *Topoi*, 35(1), pp.213–224.
- Tanczer, L.M., Brass, I. and Carr, M., 2018. CSIRTs and Global Cybersecurity: How Technical Experts Support Science Diplomacy. *Global Policy*, 9(3), pp.60–66.
- Tau*. 2018. Directed by F. D’Alessandro. Netflix.
- The H, 2010. Trojan spreads via new Windows hole. *The H Security*. [online] 15 Jul. Available at: <<https://web.archive.org/web/20181030120915/http://www.h-online.com/security/news/item/Trojan-spreads-via-new-Windows-hole-1038992.html>> [Accessed 30 Oct. 2018].
- The Rendon Group, 2011. *Conficker Working Group: Lessons Learned*. [online] Available at: <<https://perma.cc/3SBC-VXQ8>> [Accessed 9 Dec. 2016].
- Thornton, P., 2015. The meaning of light: seeing and being on the battlefield. *Cultural Geographies*, 22(4), pp.567–583.
- Thornton, P., 2018. A Critique of Linguistic Capitalism: Provocation/Intervention. *GeoHumanities*, 4(2), pp.417–437.
- Thornton, P., 2019. *Language in the Age of Algorithmic Reproduction: A Critique of Linguistic Capitalism*. [Doctoral Thesis] Royal Holloway, University of London. Available at: <https://web.archive.org/web/20190517152606/https://pure.royalholloway.ac.uk/portals/files/33473592/THORNTON_THESIS_FINALFINAL.pdf> [Accessed 17 May 2019].
- Thrift, N., 2005. From born to made: technology, biology and space. *Transactions of the Institute of British Geographers*, 30(4), pp.463–476.
- Thrift, N. and French, S., 2002. The automatic production of space. *Transactions of the Institute of British Geographers*, 27(3), pp.309–335.
- Thrift, N.J., 2008. *Non-Representational Theory: Space, Politics, Affect*. London: Routledge.

- Timmers, P., 2018. The European Union's cybersecurity industrial policy. *Journal of Cyber Policy*, 3(3), pp.363–384.
- Tivadar, M., Balázs, B. and Istrate, C., 2013. *A Closer Look at MiniDuke*. [online] Bitdefender. Available at: <https://web.archive.org/web/20190603102120/https://labs.bitdefender.com/wp-content/uploads/downloads/2013/04/MiniDuke_Paper_Final.pdf> [Accessed 3 Jun. 2019].
- Tolia-Kelly, D.P., 2013. The geographies of cultural geography III: Material geographies, vibrant matters and risking surface geographies. *Progress in Human Geography*, 37(1), pp.153–160.
- Tolia-Kelly, D.P., 2016. *Landscape, Race and Memory: Material Ecologies of Citizenship*. London: Routledge.
- Touchette, F., 2016. The evolution of malware. *Network Security*, 2016(1), pp.11–14.
- Turkle, S., 1984. *The Second Self: Computers and the Human Spirit*. London: Granada.
- Ulasen, S., 2010. Rootkit.TmpHider. *Wilders Security Forums*. Available at: <<https://web.archive.org/web/20181030112313/https://www.wilderssecurity.com/thread/rootkit-tmphider.276994/>> [Accessed 30 Oct. 2018].
- Vertesi, J., 2015. *Seeing like a Rover: How Robots, Teams, and Images Craft Knowledge of Mars*. Chicago: University of Chicago Press.
- Vira, B. and James, A., 2011. Researching hybrid 'economic'/'development' geographies in practice: Methodological reflections from a collaborative project on India's new service economy. *Progress in Human Geography*, 35(5), pp.627–651.
- VirusBlokAda, 2010. *Rootkit.TmpHider*. [online] VirusBlokAda. Available at: <<https://web.archive.org/web/20100722095105/http://anti-virus.by:80/en/tempo.shtml>> [Accessed 22 Jul. 2010].
- Von Neumann, J., 1961. The General and Logical Theory of Automata. In: A.H. Taub, ed. *John von Neumann: Collected Works*. Oxford: Pergamon Press. pp.288–326.
- Von Neumann, J., 2012. *The Computer and the Brain*. New Haven, Conn.: Yale University Press.
- Von Neumann, J. and Burks, A.W., 1966. *Theory of Self-Reproducing Automata*. Urbana, I.L.: University of Illinois Press.
- Vukov, T. and Sheller, M., 2013. Border work: surveillant assemblages, virtual fences, and tactical counter-media. *Social Semiotics*, 23(2), pp.225–241.
- Warf, B., 2015. Cyberwar: A new frontier for political geography. *Political Geography*, 46, pp.89–90.
- Warren, M. and Hutchinson, W., 2000. Cyber attacks against supply chain management systems: a short note. *International Journal of Physical Distribution & Logistics Management*, 30(7/8), pp.710–716.
- Webster, A.E., 1989. University of Delaware and the Pakistani computer virus. *Computers & Security*, 8(2), pp.103–105.
- Weimann, G., 2016. Going Dark: Terrorism on the Dark Web. *Studies in Conflict & Terrorism*, 39(3), pp.195–206.
- Wesseling, M., de Goede, M. and Amoore, L., 2012. Data Wars Beyond Surveillance: Opening the Black Box of SWIFT. *Journal of Cultural Economy*, 5(1), pp.49–66.
- Whalen, S., 2017. WannaCry ransomware analysis: Samples date back to at least early February 2017. Available at: <<https://web.archive.org/web/20181009091356/https://seanthegeek.net/2017/wannacry-ransomware-analysis/>> [Accessed 9 Oct. 2018].
- Whatmore, S., 2006. Materialist returns: practising cultural geography in and for a more-than-human world. *cultural geographies*, 13(4), pp.600–609.
- Whitson, R., 2017. Review: Jussi Parikka, *Digital Contagions: A Media Archaeology of Computer Viruses*. *Theory, Culture & Society*, 34(7–8), pp.293–298.

- Wiener, N., 1948. *Cybernetics: Or Control and Communication in the Animal and the Machine*. New York: J. Wiley.
- Wilson, M.W., 2018. On being technopositional in digital geographies. *cultural geographies*, 25(1), pp.7–21.
- Windows Defender Research, 2017. *Analysis of the Shadow Brokers release and mitigation with Windows 10 virtualization-based security*. [online] Microsoft Secure. Available at: <<https://web.archive.org/web/20180202073148/https://cloudblogs.microsoft.com/microsoftsecure/2017/06/16/analysis-of-the-shadow-brokers-release-and-mitigation-with-windows-10-virtualization-based-security/?source=mmmpc>> [Accessed 2 Feb. 2018].
- Winkler, I. and Gomes, A.T., 2017. What Is Threat Intelligence? In: I. Winkler and A.T. Gomes, eds. *Advanced Persistent Security: A Cyberwarfare Approach to Implementing Adaptive Enterprise Protection, Detection, and Reaction Strategies*. [online] Cambridge, Mass.: Syngress, pp.143–150. Available at: <<http://www.sciencedirect.com/science/article/pii/B9780128093160000129>>.
- Wolf, W.H., 2009. Cyber-physical systems. *IEEE Computer*, 42(3), pp.88–89.
- Woodward, K., Jones, J.P., Vigdor, L., Marston, S.A., Hawkins, H. and Dixon, D.P., 2015. One Sinister Hurricane: Simondon and Collaborative Visualization. *Annals of the Association of American Geographers*, 105(3), pp.496–511.
- Yong, W. and Worth, R., F., 2010. Bombings Hit Atomic Experts in Iran Streets. *New York Times*. [online] 29 Nov. Available at: <https://web.archive.org/web/20181031164943/https://www.nytimes.com/2010/11/30/world/middleeast/30tehran.html?_r=1> [Accessed 31 Oct. 2018].
- Yusoff, K., 2013. Geologic life: prehistory, climate, futures in the Anthropocene. *Environment and Planning D: Society and Space*, 31(5), pp.779–795.
- Zandi, A., 2014. *USC - Viterbi School of Engineering - Computer Virus: An Origin Story*. [online] archive.is. Available at: <<http://archive.is/iYfs8>> [Accessed 19 Jun. 2018].
- Zapf, H., 2016. Cultural Ecology of Literature – Literature as Cultural Ecology. In: H. Zapf, ed. *Handbook of Ecocriticism and Cultural Ecology*. Berlin: De Gruyter, pp.135–153.
- Zegart, A. and Lin, H., 2017. Introduction to the special issue on strategic dimensions of offensive cyber operations. *Journal of Cybersecurity*, 3(1), pp.1–5.
- Zero Days*. 2016. Directed by A. Gibney. Available at: <<http://www.zerodaysfilm.com/>> [Accessed 25 Oct. 2018].
- Zetter, K., 2014. *Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon*. First Edition. ed. New York : Crown Publishers.
- Zuboff, S., 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. London: Profile Books.

Appendix

Consent Forms

I. Participant Consent Form signed by participants.

School of Geography and the Environment
South Parks Road, Oxford, OX1 3QY.

Dr Beth Greenhough (Supervisor)
Email: beth.greenhough@ouce.ox.ac.uk
Andrew Dwyer (DPhil Student)
Mobile: +44 [REDACTED]
Email: andrew.dwyer@cybersecurity.ox.ac.uk



PARTICIPANT CONSENT FORM

CUREC Approval Reference: SOGE 17 1A-4

Malware Ecologies: Disrupting the Geographies of Cyberspace and Cybersecurity

Purpose of Study: A series of interviews with individuals who have analysed, detected or disseminated information about specific malware (Conficker, Dridex, Stuxnet, The Dukes and/or WannaCry/Petya) to discover how people work with malware.

Please initial each box

- | | | |
|----|--|--------------------------|
| 1 | I confirm that I have read and understand the information sheet version v1.0 dated July 2017 for the above study. I have had the opportunity to consider the information, ask questions and have had these answered satisfactorily. | <input type="checkbox"/> |
| 2 | I understand that my participation is voluntary and that I am free to withdraw at any time, without giving any reason, and without any adverse consequences or academic penalty. | <input type="checkbox"/> |
| 3 | I understand that research data collected during the study may be looked at by designated individuals from the University of Oxford where it is relevant to my taking part in this study. I give permission for these individuals to access my data. | <input type="checkbox"/> |
| 4 | I understand that this project has been reviewed by, and received ethics clearance through, the University of Oxford Central University Research Ethics Committee. | <input type="checkbox"/> |
| 5 | I understand who will have access to personal data provided, how the data will be stored and what will happen to the data at the end of the project. | <input type="checkbox"/> |
| 6 | I understand how this research will be written up and published. | <input type="checkbox"/> |
| 7 | I understand how to raise a concern or make a complaint. | <input type="checkbox"/> |
| 8 | I consent to being audio recorded | <input type="checkbox"/> |
| 11 | I understand how audio recordings will be used in research outputs | <input type="checkbox"/> |
| 12 | I understand that material from my interview may be quoted in research publications but that it should not be possible to identify me from the quoted material | <input type="checkbox"/> |
| 13 | I agree to take part in the above study. | <input type="checkbox"/> |

Name of Participant Date Signature

Name of person taking consent Date Signature

2. Participant Information Sheet read by all participants prior to the signing of the Participant Consent Form.

School of Geography and the Environment
South Parks Road, Oxford, OX1 3QY.

Andrew Dwyer (DPhil Student)
Mobile: +44 [REDACTED]
Email: andrew.dwyer@cybersecurity.ox.ac.uk



Malware Ecologies: Disrupting the Geographies of Cyberspace and Cybersecurity

PARTICIPANT INFORMATION SHEET

Ethics Approval Reference: SOGE 17 1A-4

1. Background and aims of the study

Andrew's doctoral research is dedicated to the analysis of malware ecology, from both a technical and social perspective. As part of this, he is undertaking a series of interviews with individuals who have analysed, detected or disseminated information about specific malware (Conficker, Dridex, Stuxnet, The Dukes, and/or WannaCry/Petya) to generate a greater understanding of how people work with malware. This is to investigate three core themes:

- malware analysis;
- detection, and;
- how this information is disseminated.

This study is funded by the UK Engineering and Physical Sciences Research Council (EPSRC) that is administered by the Centre for Doctoral Training in Cyber Security.

2. Why have I been invited to take part?

You have been invited due to your previous involvement with a specific piece of malware (Conficker, Dridex, Stuxnet, The Dukes, and/or WannaCry/Petya).

3. Do I have to take part?

You are under no obligation to participate. You can ask questions about the study before deciding whether or not to participate. If you do agree to participate, you may withdraw yourself from the study at any time, without giving a reason and without penalty, by advising the researchers of this decision.

4. What will happen in the study?

You will be asked a series of questions about your involvement with a specific piece of malware. These should take no longer than one hour and will include questions emerging from Andrew's previous work researching these malwares and experience with malware analysts. These will take place at a mutually convenient location or online.

5. Are there any potential risks in taking part?

This study poses a low risk to you and there are no foreseeable risks.

6. Are there any benefits in taking part?

There will be no direct benefit to you from taking part in this research.

7. What happens to the data provided?

Your personal information will be kept confidential and anonymous. The only information to be collected on you as an individual will be your role (e.g. malware analyst, journalist). Audio recordings will be securely stored with no personal identifier attached to each audio file. Direct quotes may be used with your authorisation on the consent form, but you shall still not be named. All audio recordings shall be archived

till the end of the doctoral research and destroyed. The anonymised transcripts and audio recordings will be kept in an encrypted folder using Apple Mac's Disk Utility utilising 256-bit AES encryption.

All research data (anonymised transcripts) and records will be stored for a minimum retention period of three years after publication or public release of the work of the research, in accordance with the [University of Oxford's Policy on the Management of Research Data & Records](#).

In addition, data will be kept for a further 7 years after publication according to EPSRC's [requirements](#).

8. Will the research be published?

This research will lead to publication in a variety of different formats, including academic journals. In addition, this research will form the basis for a doctoral thesis that will be published online. Any reference to your interview in these publications will be anonymised.

The University of Oxford is committed to the dissemination of its research for the benefit of society and the economy and, in support of this commitment, has established an online archive of research materials. This archive includes digital copies of student theses successfully submitted as part of a University of Oxford postgraduate degree programme. Holding the archive online gives easy access for researchers to the full text of freely available theses, thereby increasing the likely impact and use of that research.

If you agree to participate in this project, the research will be written up as a thesis. On successful submission of the thesis, it will be deposited both in print and online in the University archives, to facilitate its use in future research. The thesis will be published open access meaning available to every internet user.

9. Who has reviewed this study?

This study has been reviewed by, and received ethics clearance through, the University of Oxford Central University Research Ethics Committee (reference number: SOGE 17 1A-4).

10. Who do I contact if I have a concern about the study or I wish to complain?

If you have a concern about any aspect of this study, please speak to the relevant researcher on +44 7805940616 or their supervisor Dr Beth Greenhough (beth.greenhough@ouce.ox.ac.uk) who will do their best to answer your query. The researcher should acknowledge your concern within 10 working days and give you an indication of how they intend to deal with it. If you remain unhappy or wish to make a formal complaint, please contact the chair of the Research Ethics Committee at the University of Oxford (contact details below) who will seek to resolve the matter in a reasonably expeditious manner:

Chair, Social Sciences & Humanities Inter-Divisional Research Ethics Committee; Email: ethics@sohsc.ox.ac.uk; Address: Research Services, University of Oxford, Wellington Square, Oxford OX1 2JD

11. Contact Details

If you would like to discuss the research with someone beforehand (or if you have questions afterwards), please contact:

Andrew Dwyer
School of Geography and the Environment
South Parks Road, Oxford, OX1 3QY.
Mobile: +44 [REDACTED]
Email: andrew.dwyer@cybersecurity.ox.ac.uk