

Philosophy meets Internet Engineering: Ethics in Networked Systems Research. (GTC workshop outcomes paper)

Authors in order of first name:

- Bendert Zevenbergen, Oxford Internet Institute, University of Oxford.
 - Brent Mittelstadt, Oxford Internet Institute, University of Oxford.
 - Carissa Véliz, Uehiro Center for Practical Ethics, University of Oxford.
 - Chris Detweiler, The Hague University of Applied Sciences.
 - Corinne Cath, Oxford Internet Institute, University of Oxford.
 - Prof. Julian Savulescu, Uehiro Center for Practical Ethics, University of Oxford.
 - Meredith Whittaker, Google Open Source Research.
-

Ethics in Networked Systems Research project website:

<http://ensr.oii.ox.ac.uk/>



Contact OII

Oxford Internet Institute
University of Oxford
1 St Giles
Oxford OX1 3JS
United Kingdom

Telephone: +44 (0) 1865 287210

Fax: +44 (0) 1865 287211

Email: bendert.zevenbergen@oii.ox.ac.uk

Web: <http://ensr.oii.ox.ac.uk/>

Please cite the source of text and data excerpts as: Zevenbergen, B., Mittelstadt, B., Véliz, C., Detweiler, C., Cath, C., Savulescu, J., Whittaker, M. (2015) *Philosophy meets Internet Engineering: Ethics in Networked Systems Research (GTC Workshop outcomes paper)*. Oxford Internet Institute, University of Oxford.

© The University of Oxford for the Oxford Internet Institute 2015. This work may be copied freely for non-commercial research and study. If you wish to undertake any of the other acts restricted by the copyright you should apply in writing to the Director of the Institute at 1 St Giles, Oxford OX1 3JS, United Kingdom.

Special acknowledgements to the participants of the workshop:

- Dr. Peter Boothe, Google, New York.
- Corinne Cath, Oxford Internet Institute, University of Oxford.
- Prof. Jon Crowcroft, Computer Laboratory, University of Cambridge.
- Dr Mathieu Cunche, INSA Lyon, Inria
- Dr Christian Detweiler, The Hague University of Applied Sciences.
- Arturo Filastò, Open Observatory of Network Interference
- Danit Gal, Oxford Internet Institute, University of Oxford.
- Dr. Hannah Maslen, Uehiro Center for Practical Ethics, University of Oxford.
- Dr Brent Mittelstadt, Oxford Internet Institute, University of Oxford.
- Dr Steven Murdoch, Department of Computer Science at University College London.
- Sean Legassick, datamage
- Suzanne Leijten, De Brauw Blackstone Westbroek, Amsterdam
- Jordan McCarthy, Open Technology Institute, Washington DC.
- Dr Anders Sandberg, Future of Humanity Institute, University of Oxford.
- Dr Arjuna Sathiseelan, Computer Laboratory, University of Cambridge.
- Prof. Julian Savulescu, Uehiro Center for Practical Ethics, University of Oxford.
- Dr Mariarosaria Taddeo, Oxford Internet Institute, University of Oxford.
- Dr Thanassis Tiropanis, Web and Internet Science Group, University of Southampton.
- Carissa Veliz, Uehiro Center for Practical Ethics, University of Oxford.
- Helena Webb, Department of Computer Science, University of Oxford.
- Meredith Whittaker, Google Open Source Research, New York.
- James Williams, Oxford Internet Institute, University of Oxford.
- Dr Joss Wright, Oxford Internet Institute, University of Oxford.
- Ben Zevenbergen, Oxford Internet Institute, University of Oxford.

This workshop was been made possible by a kind donation from Google's Open Source Research team and the support from Green Templeton College, University of Oxford, as well as the Open Technology Fund.

Introduction	5
Background	6
Internet measurement	6
Practical ethics	9
Consequentialism.....	10
Deontology.....	11
Virtue Ethics.....	12
Principlism.....	13
Pluralism and casuistry	13
Computer ethics	14
Discussion: Cases.....	15
Case 1) Carna Botnet: Internet census 2012	15
Technologist's view.....	15
Philosopher's view	16
Emerging themes	17
Case 2) FluPhone	18
Technologist's view.....	18
Philosopher's view	19
Emerging themes	20
Case 3) OONI.....	20
Technologist's view.....	21
Philosopher's view	22
Emerging themes	23
Case 4) Analysing Tweets	25
Technologist's view.....	25
Philosopher's view	26
Emerging themes	26
Case 5) Linking wireless devices using Wi-Fi probe requests	27
Technologist's view.....	28
Philosopher's view	28
Emerging themes	29
Analysis.....	30
Internet as socio-technical system	30
Responsibilities resulting from power imbalances.....	31
Meaningful informed consent.....	31
Weighing risks, benefits and values for an ethical analysis	32
Status of easily accessible data.....	32
Condoning potentially unethical research methods	33
Further research	33
Conclusions	33
References	35

Introduction

The Internet is a decentralised network of information networks that uses protocols¹ to transmit an ever-increasing amount of data and information across the globe. Political, technical and business decisions enable an efficient flow of this data, but can also have malevolent motivations such as censorship, malicious attacks, prioritising content for profit and throttling data flows. Detailed metrics are needed to make sense of the state of the Internet network, at local, national or global levels, in order to ensure just governance of the network and to establish a continued enabling environment for its potential for further innovation.

Internet engineering and networked systems research improves our understanding of the underlying technical processes of the Internet. Internet engineers therefore analyse data transfers on the Internet, typically by collecting data from devices of large groups of individuals as well as organisations. The designs of Internet engineering and research projects reflect human decisions and therefore may create new moral systems. This interplay of technology and society creates new practices that can impact the lives of individuals in many ways. These actions can raise new ethical dilemmas, or challenge existing ethics methodologies within the new and complex information environment presented by the Internet.

A discrepancy exists, between human subject research – where there are relatively strict and widely-understood ethical traditions – and Internet engineering where the consideration of these issues is relatively new, and the existing best practices from other fields don't easily translate. Generally, computer scientists and network engineers tend to apply consequentialist reasoning² to project designs – whereby the end justifies the means – to meet research objectives in the most efficient way. This is partly the result of typical engineering training, where inefficiencies are to be dealt with through technical means, but possibly also the pressures of publishing interesting findings in notable journals and conferences.

One of the tasks of ethics is to analyse the impact of technology, and guide decision-making processes to reduce harms and maximise benefits. It remains an open question whether the effect of networked computer systems on society requires new branches of ethics. However, hardware and software are increasingly 'black boxes' for average people, within which the ethics of the engineers will be embedded. This makes it difficult for individuals, ethical boards, policy makers, or affected communities to scrutinize design choices and evaluate the new moral systems created. The resulting lack of factual knowledge about technology can be problematic, because these stakeholders typically reason based on deontology, where the impact of actions and design choices are considered more important than the achieved goals.

¹ Internet protocols allow for the internetworking of information on the Internet through its routing capability. See https://en.wikipedia.org/wiki/Internet_Protocol and Huitema, 2000.

² See the Background section for an explanation of ethics terms.

To further the discussion on Internet research and engineering ethics, the Ethics in Networked System Research (“ESRN”) project³ hosted a workshop at Green Templeton College, University of Oxford, on 13 March 2015. The aim of the workshop was to understand how different disciplines involved in Internet research approach ethical dilemmas and justify their reasoning. To this end, a group of 25 researchers and practitioners from two distinct groups of researchers attended the workshop: (1) Computer scientists, network engineers and other technical researchers who have faced ethical and legal dilemmas in their work, and (2) philosophers, practical ethicists, legal philosophers, and related disciplines who are interested in Internet engineering and the ethical dilemmas posed by the Internet, but may not be aware of the details, subtleties, and dilemmas of the field. Several computer scientists gave short presentations about their projects, which were then discussed in-depth by the workshop participants. The inter-disciplinary discussions led to some interesting confrontations of cross-disciplinary reasoning.

This is a perspectives paper, in which we present several of the cases discussed, as well as the reasoning applied by the different groups. The arguments made during the workshop reveal some underlying assumptions and values, which lead to some emerging themes that in turn uncover particular conceptual gaps between the disciplines. This paper is by no means intended to be a comprehensive overview of computer ethics or Internet research ethics, but merely an exploration of the themes that emerged during the workshop. However, the findings of this paper will feed into the larger ENSR project and the resulting guidelines for networked systems research.

Background

The multidisciplinary nature of this workshop and this resulting paper would necessitate a very broad literature section. In the interest of space, we will merely provide an overview of the necessary background terms and concepts needed to understand the cases and resulting analyses. We first discuss Internet measurements and why these are carried out, and follow with some basic terms in practical ethics that are useful to grasp before reading about the specific types of reasoning applied by the different disciplines during the workshop.

Internet measurement

In this section we discuss what we mean when we talk about Internet measurement, and the nature of its inclusion in a paper focused on ethics. This is not a glossary, nor is it meant to be complete or technically comprehensive. It’s a conceptual introduction, to which much could be added.

For our purposes Internet measurement is, simply, how research attempts to know what the Internet is and how the various decentralised components of the Internet are behaving. The types of behaviour encompassed are vast, and lead to a number of different experimental approaches designed to infer types of behaviour. This can be done in a number of ways, which we’ve simplified for the purposes of this document. The cases discussed in this paper

³ See the project website: <http://ensr.oii.ox.ac.uk>.

address several types of measurement, some extending their aim and methods beyond this description.

Basic approaches to Internet measurement

Active measurement

The “active” in this name refers to “actively initiating.” Active measurement involves actively sending traffic, in a pre-defined way, to (and often from) nodes connected to the Internet. Measurement using this technique will be designed to infer behavioural or topological properties based on the way the synthetic data is treated along the path, the reach of a given stream of traffic (where does a given path lead), and other observed reactions.

This approach has the benefit of consistency – the data transmitted is configured by those doing the measurement – and each time traffic is transmitted it follows the same rules. In the case of performance measurement, it also has the benefit of (potentially) collecting minimal private data from participating users, as it does not require access to “organic” traffic flows, and other more revealing user behaviour. It also presents the benefit of allowing a measurement researcher to configure both sides of a measurement (both nodes), and thus create a more controlled environment. It has the drawbacks of being fairly inflexible, and when used to map topology, it can be used to collect and expose information that can be considered private by network operators and others running Internet infrastructure.

Passive measurement

Passive here refers to “sitting back and observing”, this opposed to actively initiating as above. Passive measurement involves watching “organic” Internet traffic and measuring it. The variables measured can be the same as those measured actively (speed, latency, path, etc.). This type of measurement can be used to determine, for example, how quickly a user can reach Netflix, how quickly, or whether, the Guardian website loads, and potentially what variables contributed to this behaviour (did an Internet user have many applications running while trying to stream their favourite program? Was there a routing error at a critical exchange point that increased latency?). This approach has the benefit of collecting rich and relevant information on actual Internet traffic – something difficult to simulate. However, it is more invasive in terms of data collection when used to measure network performance, which nearly always requires sample population of “real” users and observation of their “actual” Internet use. Its use in mapping topologies is, conversely, considered less invasive, as it collects and interprets signals transmitted by network devices, and doesn’t actively probe or uncover information.

Basic types of measurement (what we measure)

Mapping

Mapping is what it says: attempts to know and understand the topology of the Internet, and the relevant features that make up this topology. It involves measurement – passive or active – designed to know the paths, devices, etc. that make up the Internet. Mapping is often used in conjunction with performance measurement (as when asking the question “how fast did

traffic travel along a given path?”). A traceroute,⁴ which collects information on the path travelled between two nodes, which network addresses resided along that path, is an example of an active mapping measurement. The Carna Botnet, discussed as a case study further along in this paper, is arguably another example of an active mapping measurement – using the viral properties of a botnet to “infect” thousands of internet-addressed devices, and thus produce an intricate map of Internet-connected devices that had not been available previously. Gathering publicly available Border Gateway Protocol (BGP) routing tables⁵ and using these to understand the connectedness of different points on the Internet would be an example of passive mapping measurement.

Performance

Performance measurement attempts to understand how traffic behaves when travelling between Internet-connected nodes. As with mapping, performance measurement can be either passive or active. Examples of performance characteristics often measured include speed (how quickly does a specific type of traffic travel to and from connected nodes), packet loss (how much traffic is lost between nodes (and thus, must be retransmitted)), jitter (what is the variability in performance between separate packets), and a number of other arcane metrics. Note that terms like “speed” do not connote an agreed-upon means of measuring performance. There are many different types of traffic, and many different methods for deducing its behaviour as it traverses the Internet.

Measurement and monitoring have existed since before the Internet. Lawrence Kleinrock's lab at University of Southern California focused on measuring and monitoring the performance of the nascent ARPANET⁶. These measurements were an attempt to understand how the system as-designed was working when-deployed. Measurements produced signals that let designers and implementers know what was broken and what could be improved, and allowed them to monitor the success or failure of theoretical assumptions. All systems, whether technical or human, require some type of feedback and iteration to ensure continued functioning.

Internet measurement present various exemplary case studies of a technical discipline whose ethical implications have received increasing attention in time with the encroachment of the Internet into individuals' private lives and geographies. To put this concretely, when the ARPANet began, it had 4 nodes,⁷ which grew to 113 and then reduced to 68 when the military and scientific networks separated. These nodes consisted of a handful of powerful computers (powerful for the time) located at government facilities and research universities across the US. This was the measurement challenge then. Now, each person with a smartphone represents a “node”; each laptop, each Internet-enabled device, each domain

⁴ A traceroute is a tool to measure the path and latency of an Internet connection. See <https://en.wikipedia.org/wiki/Traceroute> and http://tools.ietf.org/html/rfc1393?&sa=U&ei=an_rUrahFNXZoASGo4KoBg&ved=0CI8CEBYwJw&usq=AFQjCNHorrUNNvVQbTBeegcnKIGCP1ujwA.

⁵ The Border Gateway Protocol selects the best path for data to be transmitted over the Internet. See for example <http://www.enterprisenetworkingplanet.com/netsp/article.php/3615896/Networking-101-Understanding-BGP-Routing.htm>

⁶ The Advanced Research Projects Agency Network (ARPANET) was a network that was the foundation for the Internet. See Abbate, 1994

⁷ Nodes being other places from which one could connect to the ARPANet.

name server -- each connected thing is a node. Nodes abound. Each of these nodes, across all of these humans, plus the nodes (host servers⁸, DNS servers⁹, etc.) these humans connect to, are connected to all others via a nearly infinite and unfixed number of paths, each carrying data adhering to multiple protocols and serving algorithms.

Mapping what this looks like, how this works, how it behaves, where it is broken requires vantage points from as many nodes as possible, across as many paths as possible. Useful measurement data about Internet behaviour contains information on where these nodes are, what Internet Service Provider (ISP) or other network operator they are connecting from, the type of device, the type and quantity of traffic sent to and from a node, and how that traffic behaved along a given path.

The complexity and scale of measuring the Internet completely is so vast that it has not yet been achieved. We do not know what the Internet looks like as a total system, where it starts and stops, how changes made by one party impact another. Given the Internet's ad-hoc and dynamic architectural constitution, it is unlikely we will ever know. However, the lack of even good-enough knowledge has become a problem increasingly recognized by policymakers, researchers, and a number of other constituencies invested in an Internet that will continue to allow access and engagement in an open and free way, without burdening the system as a whole.

It's at this intersection that the current network measurement field exists. Faced on the one hand with a demand for data, for facts, for clearly quantified certainty about what the Internet is and the way it works, how, and for whom. On the other hand, faced with the ethical and practical challenges of collecting and storing data from useful vantage points (human users), collecting this data consistently, and at a scale that makes it useful and representative of The Internet.

Practical ethics

In this section we provide a brief overview of a few of the most common approaches used in practical ethics. For more in-depth overviews of what ethics is and the kinds of problems it deals with, we refer the reader to Blackwell's *A Companion to Applied Ethics* (Frey and Wellman, 2008 and Peter Singer's *Practical Ethics* Singer, 2011).

The task of practical ethics is to identify moral problems in different target situations, clarify what the stakes are in each case, conceptually explore possible courses of actions (considering their most relevant implications), and justify and suggest what the best course of action is likely to be. Practical ethics suggests what is the right thing to do by appealing to moral reasons.

As opposed to legal reasoning, which appeals to what one can and cannot do according to a current system of law, moral reasons appeal not to what is, but to what should be. One way

⁸ Organisations that offer a place from which one can make their website accessible on the worldwide web. See also: https://en.wikipedia.org/wiki/Web_hosting_service.

⁹ Computer servers that act as a directory for Internet addresses. See also: https://en.wikipedia.org/wiki/Name_server.

of thinking about this kind of reasoning is to ask oneself what course of action would make the world the best possible place it can be. Legal and moral actions come apart because staying within the limits of the law might not be sufficient to fulfil one's ethical duties. That is, one may be acting in a perfectly legal way that is nevertheless immoral, as when we betray a friend by disseminating secrets. Conversely, the right thing to do may not always be the legal thing to do, as when people helped Jews in Nazi Germany. That being said, it is worth keeping in mind that, in a democratic society where laws are generally taken to be legitimate, something being legal usually counts in favour of something being moral, because most people would agree that the best possible world we can think of is one in which people act within the confines of the law. In the field of computer ethics, it is a commonly held belief, however, that it is becoming increasingly difficult to apply existing laws to the new informational environment created – and actions enabled – by networked computers. This results in a policy vacuum (Moor, 1985), which may lead to existing laws are increasingly losing legitimacy, or at least being challenged by various stakeholder groups.

In order to determine and justify what they believe is the best course of action, ethicists resort to different ethical frameworks and theories. To make an ethical analysis (albeit presented in a very simplified form here), one needs to take into account the facts of the case at hand, as well as identify the values that are impacted. The relevant risks, harms, and benefits, as well as the probability of occurrence and magnitude of impact should be identified and disclosed. Finally, these elements need to be weighed in light of the context of the situation and the resulting decision must be justified. In what follows we will briefly describe the most common approaches: consequentialism, deontology, virtue ethics, principlism, casuistry, and give a brief description of the field of 'computer ethics'.

Consequentialism

Consequentialism is the view that the moral rightness or wrongness of an action is determined entirely by the states of affairs it causes. In short, the ends justify the means. An arrays of different kinds of *consequentialism* exist.

Utilitarianism – whose classical proponents were Jeremy Bentham, John Stuart Mill, and Henry Sidgwick – holds that the right moral action is the one that maximises utility, where utility can be defined in a variety of ways. For *hedonic utilitarians*, actions are morally right in so far as they promote the greatest pleasure for the greatest number of people – in turn, pleasure can be defined in many different ways. For *preference utilitarians*, morally right actions allow the greatest number of people to live according to their own preference. Other options include *welfare utilitarianism* and *negative utilitarianism*, the latter aiming to prevent the greatest amount of suffering for the greatest number of people.

Another important distinction is between *act consequentialism* and *rule consequentialism*. Act consequentialists believe that, when morally evaluating an action, one must compare the consequences of that act with the consequences of other acts the agent can engage in. For rule consequentialists, the relevant comparison is between the consequences of moral rules of action. An act is therefore morally right if it is allowed by rules justified by their consequences.

It is common practice to analyse problems in research ethics and public policy in consequentialist terms. Such analyses explore the probable consequences of different courses of action (including not acting at all) and try to determine which option is best in terms of the positive consequences it produces.

The most salient problem for these views is that consequentialist theories can justify actions that people typically consider to be morally wrong. Put crudely, in some cases consequentialists will defend the view that – provided that the positive consequences are good enough for a big enough number of people to outrun the bad consequences – the ends justify the means. Therefore, a consequentialist might be willing to violate people's rights if that action will produce enough good consequences. Other things being equal, if an innocent person must be sacrificed for the greater good, from a consequentialist perspective, that is a morally good action.

Other challenges for consequentialism include the uncertainty of the future and the difficulty of predicting consequences (actual consequences may be very different from expected consequences), deciding when to stop assessing consequences to evaluate an action (since consequences are potentially infinite), and balancing different categories of consequences (economic, political, individual, etc.). These challenges are particularly pertinent for evaluating actions in the complex and dynamic Internet environment.

Deontology

Deontological theories stand in opposition to consequentialism. While most deontologists do take consequences into account when evaluating the morality of an action, they do not consider consequences to be the only element that should be taken into account.

Furthermore, deontologists typically think that there are some actions (e.g. torture) that are so morally reprehensible that no amount of good consequences can justify them – that is, some choices are morally forbidden irrespective of the good they can create.

Historically, Immanuel Kant is often identified as the paradigmatic deontologist. The Humanity formulation of his Categorical Imperative – that we should never act in such a way that would make us treat people as a means only, but always as ends in themselves – is frequently appealed to in the practical ethics literature. Very few deontologists, however, are prepared to endorse all of Kant's ethical theory. His Universal Law of Nature formulation of the Categorical Imperative — that we should only “act only in accordance with that maxim through which you can at the same time will that it become a universal law”—committed him to such unpopular stances as arguing that we should never lie, even if it can save an innocent life from murder.

From the point of view of deontology, what makes an action right is its conformity to a moral norm. One way to divide deontology is between *agent-centered* and *patient-centered* theories. According to agent-centered theories, one has permissions and obligations that give one objective reasons for action. Patient-centered theories focus on rights, rather than duties – the scope of morally right actions is delimited by the rights (or claims) people have. According to this view, an action is morally permissible as long as it violates no one's rights.

Another kind of deontological theory is *contractualism*, which seems to cut across the distinction between agent and patient-centered theories. According to contractualism, morally wrong acts are those acts that would be forbidden either (1) by principles that people would accept as part of a just social contract (e.g., Rawls, 2009), or (2) principles that people could not reasonably reject (e.g., Scanlon, 2003).

Like consequentialism, deontological ethics is not without problems. The most important problem is that in some cases deontology seems to commit one to duties that can have very bad consequences (as in the lying case). Deontologists can defend themselves by saying that, along those duties, there are rights (e.g. the right not to be murdered) that create limits to actions and protect people from harm. The challenge then, however, is to decide how to balance rights when they conflict without simply appealing to consequences.

Virtue Ethics

Virtue ethics takes an alternative approach to determining the ethical acceptability of an action. Whereas consequentialism and deontology examine the quality of an action, virtue ethics is concerned with the *character* of the actor — it prescribes “how we should be rather than what we should do” (Darwall, 2003). Virtue ethics can provide guidance for moral action by connecting the acceptability of actions to the character of the actor (Oakley, 2007). Right actions are those that would be chosen by a ‘virtuous’ agent.

Traditionally, virtue ethics was based on the purpose (“telos”) of humanity, meaning that virtues are “habitual dispositions” to act in accordance with the ends of human life which define a ‘good’ life (Pellegrino and Thomasma, 1993). An Aristotelian approach to virtue ethics has been the most influential, with a ‘means’ based approach to defining virtues as the mean between two vices (Pellegrino and Thomasma, 1993; Darwall, 2003). Modern virtue ethicists (e.g. Darwall, 2003; MacIntyre, 1985) move beyond Aristotle’s ‘mean’ based account of the virtues, acknowledging that certain virtues such as ‘justice’ or ‘fidelity to trust’ cannot have a mean (Pellegrino and Thomasma, 1993). Virtues act as a hermeneutic tool (Edgar, 2005) helping individuals interpret, define and pursue the ‘good’ life.

Virtue ethics is therefore an account of “ethically deep aspects of human life” that stands in contrast to morality as practiced as an irreconcilable dichotomy between competing conceptions of ‘good conduct’ (MacIntyre, 1985), focusing instead on the character traits necessary to lead a ‘good life’ (Darwall, 2003). Virtue ethics does not stand in opposition to consequentialism and deontology, but rather points to an additional quality to be considered in ethical analysis. Justification by reference to virtuous character creates a very broad guide for actions which must be refined by defining specific virtues (Oakley, 2007), which can be further specified for individual practices (MacIntyre, 1985), such as medicine or military service. For example, MacIntyre connects virtues to the practical wisdom or norms of good behaviour which are internal to a practice and developed over time through experience. The virtues are therefore traits that dispose the practitioner to fulfilling the ends of the practice, while still being defined within a broader telos or understanding of the ‘good’ life (MacIntyre, 1985).

Principlism

Principlism has evolved into a practical approach for ethical decision-making that focuses on the common ground moral principles of autonomy, beneficence, nonmaleficence, and justice.

The Belmont Report was drafted by the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research in 1978 (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 1978). It was partly prompted by the revelation in 1972 that the Tuskegee Syphilis Experiment had knowingly failed to treat participants in order to study the natural progression of untreated syphilis. The Report proposed three ethical principles that must be followed when carrying out research with human subjects: respect for persons, beneficence, and justice.

Inspired by the Belmont Report, four principles are widely acknowledged and used in the practical ethics literature—particularly within medical ethics: *respect for autonomy* (the obligation to respect autonomous people's choices), *non-maleficence* (the obligation to avoid causing harm), *beneficence* (the obligation to provide benefits and balance risks against benefits), and *justice* (obligations of fairness in the distribution of risks and benefits) (Beauchamp and Childress, 1979). The principle of '*Respect for Law and Public Interest*' was added to this list in the Menlo Report, a guidelines document for Information and Communication Technology Research (Dittrich et al., 2011).

Critics of principlism (e.g. Harris, 2003) have pointed out that, while the principles can be a good checklist for people who are new to ethics or for committees who do not have ethicists on their board, to be tied down by this framework may obscure all of the nuances and stakes of real life ethical dilemmas. (Clouser and Gert, 1990; Clouser and Gert, 1994) have criticised principlism by arguing that it does not give enough guidance to action. They think principlism instructs agents to think about these issues, but it does not give people a method to weigh each principle or handle conflict among principles.

Pluralism and casuistry

Unsatisfied with the problems and limits of applying one abstract ethical theory to real life cases, practical ethicists often take a more pluralistic approach: they appeal to more than one ethical framework, borrowing reasons from consequentialism, deontology, and virtue ethics, they use *principlism*, and/or they appeal to common sense when weighing reasons and balancing risks against benefits. Norman Daniels, for example, has argued that philosophical theories are not sufficiently fine-grained to be applied to concrete cases and must be complemented by moral and political deliberation (Daniels, 1996; Daniels, 2007).

While theoretical philosophers have a top-down, deductive approach, anti-theorists think bottom-up, starting from the factual particularities of the case to be analysed (Arras, 1990). According to Jonsen and Toulmin (1988), we should start by thinking about a paradigmatic case where our intuitions are strong and clear, and reason our way by analogy: we should compare salient features in the paradigm case with the features of the case that we want to analyse. Defenders of *casuistry* believe that it is easier to achieve consensus in pluralistic

societies following this method than arguing about the theoretical underpinnings that motivate our positions (Toulmin, 1982).

Problems with these approaches include the possibility that two or more elements used from different theories or methods will conflict, the difficulty of assessing a case without reference to principles or ethical theories (i.e. our notion of what is morally relevant is likely to be informed or inspired on importance of consequences, rights, etc.), and the difficulty of justifying decisions to the people who will bear the consequences of those decisions without appealing to well-established ethical principles.

Computer ethics

Computer ethics is the study of social and ethical impact of computing technologies, which can inform the formulation and justification of policy and codes specifying ethical uses of computing.

The term 'computer ethics' first appeared in the 1970s (Bynum, 2008); however, discussion of ethical concerns arising from digital computing goes back to the original development of the technology (Bynum, 2008; Wiener, 1954; Wiener, 1964). Norbert Wiener, the father of cybernetics and one of the pivotal figures in the development of digital computers, recognised early that computing had the potential to change many aspects of life and as such was of crucial ethical importance. In transitioning from industrial to information economies, computing technologies have come to pervade most aspects of personal, organizational and social life. This development led to a steadily growing academic discourse on the ethics of computing technologies.

The malleability of computing technology encourages innovative uses and technical artefacts, which necessarily outpace policy and law designed to govern it, creating governance vacuums. Even when relevant laws and policies do exist, they have often been formulated in response to less versatile technologies (Moor, 1985). Computers have increasingly become less recognisable through integration into other technologies and environments, seen for instance in ambient intelligence and ubiquitous computing (Friedewald et al., 2005; Sadri, 2011). The original notion of computers as easily recognisable artefacts has thus become obsolete.

In parallel, computer ethics has expanded its scope of concern to address emerging uses of computing, increasingly blurring the lines between computer ethics and related applied ethics disciplines, e.g. medical ethics, technology ethics or environmental ethics. This diffusion of computers into their surrounding environment raises new ethical questions related to issues such as privacy, surveillance, autonomy or ownership (Lally et al., 2012; Quilici-Gonzalez et al., 2012; Wright et al., 2008). A key job of computer ethics is therefore to identify emerging policy requirements and to construct complementary conceptual frameworks which assist in comprehension of the ethical effects of emerging computing technologies (Moor, 2005). These responsibilities are shared by scholars across a broad range of fields including computing, philosophy, law and social sciences.

Discussion: Cases

During the workshop, technical researchers gave a brief outline of a project in which they had encountered objections on ethical grounds, or had questions about how to proceed. The researchers explained the technical aspects of the project, the goals, and perceived benefits. After the presentation, the philosophers and technical researchers asked questions and debated the ethical implications and justifications, to finally give some advice on how the project could be redesigned.

In this section we describe some of the cases presented during the workshop in Oxford, as well as the resulting perspectives and emerging themes from the discussions between the representatives of different academic disciplines. In the interest of space, the objective and technology implicated in the case are described very briefly, with a reference to additional resources. This is followed by the ethical dilemma posed by a case and a reflection on the perspectives of the different participants.¹⁰

Case 1) Carna Botnet: Internet census 2012

A paper published by an anonymous researcher (or group of researchers) describes how they created a topological map of the Internet, as well as one of the most comprehensive IPv4¹¹ surveys of the Internet currently available (Anonymous, 2012). They infiltrated 420,000 information systems without permission by exploiting default passwords of a forgotten about entry to the Linux operating system. The researchers created a botnet¹², by which code spread to other systems in the networks of the infected system, in order to grow the measurements base.

Technologist's view

In a display of *act consequentialist* reasoning, the authors acknowledge that the method of collecting this data and creating the subsequent map violated laws, but justified the means with the beneficial value of the end – the creation of a map of the Internet. One of the main arguments posed by technologists was the question surrounding the counter-factual question: What are the ethical costs of not having this information presented by the Carna Botnet study? This left the participants divided: Some argued in favour of using the data produced by the research on the grounds that it provided information that was vital to moving forward the policy debates (most notably in the U.S.) on important issues for Internet freedom like net neutrality Commission, 2015. The study was further defended on *virtue ethics* grounds by highlighting the overall respectful attitude of the researchers towards the compromised systems, because they claimed their methods did not cause any further harm. The non-malicious intent and precautions taken to mitigate harms was suggested to be considered as factors that must be weighed when deciding to use (or not use) the data.

¹⁰ Some basic technical as well as ethics knowledge as described in the background is required to understand the descriptions that follow. Footnotes refer to further information about additional technical terms.

¹¹ The Internet Protocol (IP) is responsible for routing Internet traffic. IPv4 is the forth version of this protocol, see also: <https://en.wikipedia.org/wiki/IPv4>.

¹² Botnets are networks of interconnected computers that have been compromised and follow the commands in a hierarchical fashion. See for more information Stone-Gross et al., 2009

Other technical researchers took the deontological approach by arguing that this study should not be recreated nor directly referred to, because it gives legitimacy the unlawful methods. First, the access to the information systems used in the Carna Botnet was unlawful in many jurisdictions. Second, a ‘good worm’ – which allegedly the Carna Botnet was – can also be used for bad by malicious actors, or they could easily reproduce the same botnet with malevolent intentions when the method is published. Third, condoning this method would set a standard and precedent in research, whereby researchers and actors outside of academia could point to this study to justify projects that pose similar or even higher levels of risk and harm to uninformed participants.

The overall conclusion from the technical researchers was that from a consequentialist point-of-view this research is very interesting. However, they acknowledged that once we start considering the kind of consequences it generates and who benefits from these, the analysis gets murkier. Further, some technical researchers questioned the usefulness of a static topological graph of the Internet to better understand the Internet network, since the network itself is dynamic and the increasing uptake of mobile phones renders a spatial analysis meaningless. Instead, they argued, statistical models may prove to be more useful (see for example Haddadi et al., 2008).

Philosopher’s view

The case presenter gave an overview of how he believed various philosophers would evaluate the Carna Botnet. For example, Immanuel Kant would have denounced it, as it does not adhere the Categorical Imperative¹³. *Virtue ethicists* like David Hume, Aristotle and Confucius would consider the research to fall in the realm of a person with superior knowledge making decisions, so perhaps it could be acceptable from their perspective. Jeremy Bentham and John Stuart Mill, both utilitarians and consequentialists, would have appreciated it. John Rawls would likely have argued that this research does not hold up once you ‘start veiling eyes’¹⁴, because its benefits are primarily for the already privileged technologists who are exploiting technologically less literate users and the weaknesses in their systems, without providing them with any direct benefits. This is also a *principalist* argument under the realm of *beneficence*.

The Kantian and Rawlsian perspectives resonated most with the philosophers. They echoed the arguments of the more deontological-leaning technologists against using data from, and replicating, these types of experiments.

Several arguments were presented, many of which seem *rule-consequentialist* in nature (whereby the question is asked what would happen if everyone violated the law in this way). First, the American legal principle of the fruit of the poisonous tree¹⁵ was used as a concept to explain the dangers of using data from this research: the possibility of a ‘good worm’

¹³ See <http://plato.stanford.edu/entries/kant-moral/#CatHypImp>

¹⁴ John Rawls created a thought experiment where people would be placed behind a curtain to negotiate fair social conditions for their society. See <http://plato.stanford.edu/entries/original-position/>

¹⁵ Evidence obtained illegally must be excluded for the court proceedings. See https://www.law.cornell.edu/wex/fruit_of_the_poisonous_tree

containing exploitable vulnerabilities is too big. Further, the issues of the slippery slope and condoning the research methods by reusing the data were re-emphasized. For example, several philosophers agreed that using the data creates a precedent, and might encourage others to replicate the research, possibly with more harmful consequences.

From a technological perspective, another danger of using this research is the fact that if the data is used as an evidence base for setting standards for Internet architecture, it is going to be *de facto* embedded into the technological layer of the Internet, which cannot be easily reverse engineered. There was some practical push-back from technologists against this argument, however, when they claimed that it is impossible to prevent people from reusing or recreating the experiment, as it is not possible (or desirable) to completely remove the research's finding from the Web.

Emerging themes

The Carna Botnet case is an example of conducting unlawful research that does yield beneficial information for some groups. Participants discussed (1) how to weigh facts and values for an ethical analysis, (2) the power imbalance between (active and passive) actors, and (3) whether condoning such research methods is desirable in the long term.

(1) Facts and values

First, the inferred data, the database, and the subsequent map of the Internet were particularly useful for some actors on the Internet, and may be useful for policy makers. It is a fair consequentialist question to ask what the ethical cost of not having this information is for society. The participants were split, even within disciplines, on how to answer this question. So, a bottom-up casuistic analysis followed, where participants attempted to weigh the facts and impacted values. However, a satisfactory answer did not emerge, likely because of the complexity of the case and the lack of time on the day.

(2) Power Imbalances

A second theme that emerged is the issue of power imbalances between highly technical researchers or engineers, and the wider population that do not understand how their devices and their data are being used, remotely or otherwise. *Consequentialist* justifications by engineers tend to ignore this power imbalance, so the question emerged about how to make technical researchers aware of their superior position in this asymmetry. Could consequentialist methods be further justified if the researcher demonstrates they have the traits of a *virtuous person*, in line with the doctrine of *Virtue Ethics*? If so, to what extent would this justify unlawful research methods?

(3) Condoning unethical research methods

Third, it is important to realise that accepted research sets a precedent for future research. Published methods will be replicated. In the case of the Carna Botnet or other botnets, their methods could be used for malicious purposes. Further, if the collected and inferred data is used for formal standard setting of Internet architecture, this unlawfully obtained data will thus be incorporated as a standard to be referred to, and the research may be replicated for further standard setting with regards to the Internet architecture.

Case 2) FluPhone

The FluPhone project studied the spread of epidemics, such as certain types of the flu, by using data gathered from mobile phones (Yoneki, 2011). The project collected data on human encounters to understand how close-contact infections – such as swine flu – spread between different people in low and middle income countries. The project instigators required data subjects to install an application on their mobile phones to record information such as their location and encounters with other mobile phones through Bluetooth connections. Proximity to other humans was measured through a Bluetooth connection, which was used as a proxy to measure physical encounters. Further, the investigators added a manual self-reporting option to the application to record for user symptoms.

Not all people in a given area would install this application (or would have a mobile phone to do so), but initially all Bluetooth connections were collected, even from individuals not participating in the research project. This data and the resulting inferences can be useful to make decisions about how to allocate resources and treat different groups from a public health perspective. One of the features of the project was to provide a dashboard for users, which would show which persons they encountered were infected and which geographical areas in their vicinity were considered to be high risk.

The project included some pre-designed limitations for the data. For example, the researchers located their servers and databases in Cambridge (UK) and encrypted the database as well as the data flows. Only two people had access to this data, which would all be deleted at the end of the project.

The responsible ethics board had three main objections to this research project. First, the researchers were only allowed to collect data on encounters, rather than the precise location of this encounter. Second, the researchers were not allowed to include minors in their population. Third, the dashboard to inform people about current disease risks needed to be excluded from the project, because such fine grained information about infected areas could lead to discrimination or even violence.

Technologist's view

The engineers argued that the project as it was initially designed would be technically feasible and would have a significant social benefit for communities affected by diseases, as well as public health management. Their model of how diseases spread would be based on ground-truths and the dash board would have a direct benefit of showing individuals which persons or which areas to avoid. The focus on the benefits demonstrates consequentialist reasoning.

Their complaints about the ethics board's more deontological decisions was as follows. First, by not collecting the exact location but only the encounters, the ethics board threw out the baby with the bathwater, because the data would be much less useful to understand the actual spread of a disease, as well as alarming people and authorities about specific highly infected areas. Second, minors are very mobile agents in society (e.g. in schools, social gatherings, etc.), so excluding them from the data collection meant that an important factor

in the spread of disease needed to be ignored. The third complaint – the exclusion of the personal dashboard – was accepted by the researchers, because they had not taken into account the social repercussions of providing information about infected people or areas. Not only would this lead to discrimination, but infected houses or neighbourhoods could be burnt down by angry mobs, who would have been given the location data of at risk areas.

The researchers explained how they had run into other ethical issues while using a similar application to measure (Yoneki et al., 2007). For example, whilst conducting measurements at a company, they found that some employees were having a relationship due to the frequent encounters in remote areas of the office building. The application was also used in nightlife areas of a British town, where it appeared to be a convention that people put their gender and sexual orientation in their Bluetooth IDs in order to identify other persons with similar likings. This data was intended to be communicated to a local crowd and not to be collected by researchers whose databases may be published freely, taking this information out of its original context.

Philosopher's view

The philosophers largely agreed with the analysis of the ethical board and raised a number of additional questions. For example, when asked about the incentives offered to data subjects, the researchers stated that the original idea was to offer the personal dashboard in return. When this was excluded from the project, the researchers chose to pay participants directly, partly to cover their expenses.

The collection of all encounters was also criticised, because this would include data from people who did not (or even had chosen not to) participate in the research project. The technical researchers stated – in line with consequentialist reasoning – that this collection of all encounters would have maximised the utility of the data and made inferences much more accurate. Originally, the justification for the data collection about non-participants was that informed consent would not be necessary, since these persons would have their Bluetooth switched on anyway, and that their data would not be further analysed. However, since the means to achieve this utility would be highly invasive for persons who chose not to participate, the researchers agreed to only collect the data from registered participants, and hence only collecting data when research participants encounter each other. They argued, however, that given the technical limitations in the systems design, it would not be possible to trace back the identity of the persons whose Bluetooth connections were not registered.

Finally, the question was raised whether or not the utility, and subsequent benefits of the system, could be further leveraged by measuring additional transmittable medical conditions, such as sexually transmitted diseases. However, due to the social stigmas surrounding more serious illnesses than the flu, the researchers considered that data subjects may not be completely honest or comprehensive in the self-reporting of symptoms. This could affect the usefulness of the collected data.

Emerging themes

(1) Relevant norms and socio-political context

First, the relevant social norms of data subjects may not be fully understood by researchers, whereby collected and disseminated data may lead to potentially disastrous consequences for persons in the area. This was highlighted by the unforeseen consequences of providing a dashboard, informing frightened and vigilante-like mobs where high risk zones were, that may be dealt with through unconventional methods, such as burning down houses or parts of a village. In the end, this boils down to the debate of data utility on the one hand, and ethics, values, norms, and rights on the other hand.

(2) Collecting potentially data sensitive data through mobile phones

Second, the initial idea of collecting data from non-participants phones because they have their Bluetooth switched on, shows that a power and knowledge imbalance between technically-literate researchers and average mobile phone users exists. Even if the utility of databases would be significantly improved with such data collection, researchers need to be aware of these limitations and pro-actively build in limitations to such data collection. Although people's agency to act must not be underestimated when agreeing to participate in a research, they cannot be expected to have in-depth technical knowledge of the possibilities of their mobile phone, and the unforeseen consequences thereof for the people they interact with physically that are not part of the research.

(3) Importance of ethical decision making on per case basis

Third, the limitations set by the ethics board – especially excluding minors – reduced the utility of the dataset for public health management, according to the researchers. Likely, the decision to exclude minors was based on an interpretation of the UK Data Protection Act whereby minors are not considered to be able to give full informed consent. This shows, however, that there is a significant responsibility for all stakeholders in a research projects to understand the consequences of a research project, whereby the limitations should be set through conversation, rather than top-down decision making.

Case 3) OONI

The Open Observatory of Network Interference (OONI), is a global observation network that aims to measure network interference – such as censorship, surveillance, or data discrimination – in countries around the world (Filasto and Appelbaum, 2012). The project uses a software probe¹⁶ installed locally to infer network interference in a given region, for example by making HTTP¹⁷, HTTPS¹⁸ and DNS¹⁹ requests. The project relies heavily on voluntary participation in regions around the world as their only method of deployment, even

¹⁶ A piece of software used to test network conditions.

¹⁷ Hypertext Transfer Protocol (HTTP) is an application protocol that is the foundation for communication of the worldwide web. See also Berners-Lee et al., 1996

¹⁸ HTTPS is the secure version of HTTP, whereby data is encrypted. See also:

<https://tools.ietf.org/html/rfc2818>.

¹⁹ See footnote 8.

though the researchers have also considered remunerating their participants. Further, the OONI project aims to publish their measurement results in an open data²⁰ format. The project's organisers realise the need for meaningful informed consent, especially since the project may include risks for participants, but are unsure about which information to present to participants.

The most accurate data on censorship, or other network inference, would be collected directly within the networks in regions where a government can effectuate its information control policies. However, several examples exist where Internet users who live in regimes with strict information control policies are confronted with a record of their online behaviour when arrested for other reasons. A concrete example was raised during the workshop where a Tunisian protester during the Arab Spring of 2011 was arrested and confronted with his participation in the Herdict system,²¹ a relatively benign project to measure censorship within Tunisia. Actively participating in censorship detection or other network measurement projects may be considered by governments as an act of espionage, or disseminating information to foreign governments. It provides an easy excuse to arrest citizens.

Moreover OONI has been exploring the possibility of remunerating people that run the instrument in places that are generally harder to reach, such as Iran, but this raises the question of whether or not compensation increases the risk of participation for an OONI volunteer.

Technologist's view

Most researchers attending the workshop criticised the approach in several different ways. First, the URLs²² requested may contain unlawful or politically sensitive information in a given jurisdiction. The project is designed to test connections for content that may be censored. This means participants will be put into a situation where they are whacking into government mandated filters that may register who is trying to access the content. A consequence could be that individuals are included on a dissidents list.

The case presenter countered this argument, claiming that people will likely not be arrested merely because they participate in OONI measurements, but mentioned it could be an additional charge that is bought to people arrested for other reasons. However, the presenter admitted they were not, and could not be, aware of harms that had been done to participants because of the OONI project. A workshop participant explained (anecdotally) how a Tunisian activist had been confronted with test that he had carried out for the less-intrusive Herdict project when he was arrested during the Jasmine Revolution of 2011, which emphasises that technical researchers may not be aware of how their participants are treated when political regimes suddenly change or become more oppressive.

²⁰ Open data is defined as “[...] data that can be freely used, re-used and redistributed by anyone – subject only, at most, to the requirement to attribute and sharealike.” See for example:

<http://opendatahandbook.org/guide/en/what-is-open-data/>

²¹ See the project website: <http://www.herdicth.org/>

²² The Uniform Resource Locator (URL) specifies the location of information on computer networks, such as websites on the worldwide web. See for more information:

<http://tools.ietf.org/html/rfc1738?AFRICACIEL=trkim10qgp5sem36pe2gf8tej2>.

The technologists suggested some improvements to ensure that participants running OONI probes are adequately protected. In their view, the means needed to adhere to contextually relevant values for the end to be justifiable. For example, OONI should consider including technical complications into the system, to ensure that only sufficiently technology literate individuals who would likely understand the possible technical risks associated, can run the tests. Further, albeit difficult to transplant directly, the OONI folks were encouraged to look at some similar academic projects and learn technically implement some existing best practices to adhere to specific values.

Philosopher's view

The philosophers echoed many of the concerns raised by the technologists, but from a different point-of-view. The discussion started with the question to what extent the informed consent given by participants is meaningful. It is difficult for the project instigators to assess – and therefore to communicate – what the potential risks are for participants in their given situation. There is a problem not only of identifying these risks, but – given the complex and dynamic nature of the Internet and digital data – defining the risks in the first place. This led to the conclusion that there is a need for research into the negative consequences of Internet research, especially when it necessitates volunteers.

One of the experts suggested that OONI should gather more information about what happens to individuals who have been arrested, and tortured, for participating in similar projects to gain a deeper understanding of the potential risks they are exposed to. Referring to the Herdict project example, if even a relatively simple and innocent tool can be enough reason to arrest and torture an individual, technical researchers engaged in this field need to be more careful about exposing people to risks that are not fully understood.

Risks to participants will vary according to the political as well as socio-economic position of the individual within the society in which they are conducting measurements. Further, risks may not be immediate, but materialise in the long-term. Risks may also not directly affect the individual, but may materialise on a group level, such as discrimination in various forms against minorities. Therefore, researchers should gain relevant knowledge through local lawyers, activists, fellow technical researchers, or others who can provide expertise about the risks of running a potentially sensitive Internet research project.

The researchers should find out which aspects of the projects are considered to present risks, how bad the harm would be, and in which way it would likely materialise. However, it should be taken into account that the dynamic and ever changing environment presented by the Internet makes a risk assessment more unpredictable and thus difficult. For example, it is difficult to predict how a government will use databases of individuals that have tried to access censored content, especially when the political context becomes more oppressive or there is a regime change. Understanding the social impact of a project is important to mitigate potential future issues, which underlies the importance of deontological reasoning.

Overall, consensus of the group was that, in this particular case, the risks for the participants outweighed the potential benefits. It would likely not pass academic ethical scrutiny, partly because of the lack of *principlist* reasoning based on *beneficence*: the project exposes

individuals to a significant and undocumented risk without directly benefitting the people taking part. It was suggested that OONI has a moral obligation to protect its users, implement protections, and to provide a reasonable and accurate overview of the potential risks of participation.

The group provided various practical recommendations to tackle these issues: First, the suggestion was made to cite actual cases when trying to explain possible risks. Second, OONI should under no circumstances pay their volunteers as it makes it easier for the government to prosecute the volunteers for espionage. OONI could benefit from developing protocols that ensure there are adequate pre-emptive protections for volunteers, protecting them from harm and understanding where OONI falls in the legal framework in any given country. Third, the philosophers suggested that OONI could benefit from setting up a panel of experts who could advise them on these matters.

Emerging themes

From this discussion four themes emerged: (1) the trade-off between benefits and harms, (2) the Internet as a socio-technical system, (3) informed consent, (4) remunerating participants.

(1) The trade-off between benefits and harms

Following from the discussion on how to balance facts and values for an ethical analysis in the Carina Botnet case, the philosophers discussed how to think about benefits and harms in sensitive Internet research. From an informed consent point of view, there is a need to disclose potential harms, even if the probability of occurrence is low or their magnitude is not particularly significant. There is thus a clear need to refine existing mechanisms for defining, identifying, classifying, and delimiting the definitions of harm and their appropriate informed consent procedures.

On a philosophical level there are questions about the notion of ‘the harm done’: What constitutes harm? How bad is the harm and in what way? Does it harm individuals or groups? On a more practical level, questions such as the following arose: How can we ensure that risk assessments are accurate in the continuously changing environment of the Internet? Is there a necessity to delimit a precautionary principle for Internet research (see Narayanan et al., 2015)? This question is especially important considering that many risks materialize in the long-term, whereas the related benefits are immediate.

Finally, the question emerged how we measure the consequences of the knowledge we create through “research” and “science”? How should we respond to these, and ensure “no harm”? On a related matter, technical researchers complained that they get mixed signals from different incentive structures. For example, it is supposed to be good for science to publish data in open data format, but doing so can be bad for society due to privacy harms. There is a need to communicate a sense on where (and how) to draw the lines in the consideration of consequences of collected and inferred data.

(2) Internet as sociotechnical system

The Internet has developed into a diverse and globally used sociotechnical system that is influenced by – and affects – societies in ways that are likely not comprehensible by a single

person or actor. Access to the network has been democratised for those who can afford the end terminals, whereby the applicable values, socio-economic positions, and political environments differ drastically among its user base. The Internet has surpassed its original purpose of a scientific and military information network, but has become the backbone and central nervous system for many diverging parts of society, commerce, culture, and government.

The Internet has been (and still is being) designed by a fairly centralised and homogenous group of academically trained, highly technical people, the majority of which is Caucasian and male. Therefore, the predominant value-system that seems to guide research and engineering design decisions is that of a limited group, which is highly unlikely to be representative of the values of internet users worldwide. This is not a critique of the work that has been done by the Internet engineering community, but rather a note to highlight that the community designing the Internet technically and carrying out measurement experiments is not as diverse as the user base. When expanding ethical considerations in the design of Internet engineering projects, it was highlighted that the generality of the Internet should not be compromised in order to encompass local ethical concerns.

(3) Informed consent

Meaningful informed consent is a problematic issue on the Internet. Users often have little intuitive understanding of how personal data is used in the Internet economy, and have a limited technical understanding of their devices. The combination of these two factors renders their 'informed consent' somewhat meaningless, as their agency to make a well-informed decision is minimal. This is further complicated by the fact that companies are often in a situation where they need to create one document that speaks to all their users, in all social, cultural and political contexts in which the project, application, or other equipment may be used, whilst being unable to provide an exhaustive or relevant overview of the risks involved.

In the discussion the comparison to (medical) research was frequently drawn, as it also involves some risks to the participants. In such research there are strict procedures to ensure informed consent is meaningful by clearly defining and identifying potential risks, balancing these against benefits, and adhering to established ethical guidelines on the use of for instance deception. However, the issue was raised that medical ethics have become engrained in law, whereas Internet research may benefit more from self-regulatory structures, due to the flexibility needed to accommodate its innovative and dynamic environment.

Several participants emphasised that we should not underestimate the ability of individuals to make well-informed decisions. This argument was countered by other experts, however, who argued that in many of the countries that OONI runs its experiments, it is not just about the individual participant who have given informed consent, because if they are arrested by the authorities there may also be repercussions for their direct social circle (who did not give consent). Therefore, it may be a useful exercise to differentiate between activist participation, and passive providers of data when considering who the participants are likely to be.

(4) Remunerating participants

The general advice from the philosophers was to refrain from creating a financial link between sensitive research projects and participants. Even if the motivations of participants are good and just, it is hard to predict how government agencies will respond to network measurements. In the worst-case scenario, participants may be prosecuted for espionage, especially if data is transmitted to foreign countries.

Case 4) Analysing Tweets

The presenter described contemporary trends in research that involve the collection and analysis of Twitter data. Twitter posts can be collected in real time via the Twitter API²³. They typically contain identifiers of individuals (e.g. their usernames) and may also contain potentially sensitive information (e.g. political opinions, inflammatory content etc.). Although the collection without informed consent by users is likely in breach of data protection laws, the practice is so ubiquitous, that it doesn't seem to be enforced by regulatory agencies. It is in cases such as these where self-regulatory ethical standards become increasingly important for researchers engaging in this type of data collection. Where Twitter datasets are analysed quantitatively, characteristics of the data (demographic characteristics of users, content of tweets etc.) are typically reported at an anonymised, aggregated level. Qualitative analysis might involve the in-depth examination of individual tweets and researchers might therefore wish to include real tweets in publications.

The question the presenter put to the experts was: Is it unethical to publish information, specifically individual tweets, that is publically available without awareness of the people creating this information?

The presenter pointed out that Twitter's Developer Agreement²⁴ states that tweets must be reproduced in other publications in full – that is, including the username and full content without modification or translation. This appears to rule out the possibility of anonymising tweets by removing the user name and/or other identifying details.

Technologist's view

The first issue raised in response to the presentation questioned the legal status of tweets, more specifically whether such data can be considered to be part of the public domain? An interesting tension emerged between the contractual obligations of Twitter towards their users and common research standards that require informed consent. According to the Terms of Service (ToS) of Twitter, the contract between Twitter and the user, tweets are published publicly²⁵. However, because no contract exists between the user and the researcher, the public publishing of tweets does not entail that a researcher can pluck tweets from the web to use them for research.

²³ An application programming interface (API) is a set of functions and methods that can be used by an external user to perform some queries on the service, or a set of programming building blocks that can be used to develop an application.

²⁴ <https://dev.twitter.com/overview/terms/policy>

²⁵ The Twitter Terms of Service states: "Most Content you submit, post, or display through the Twitter Services is public by default and will be able to be viewed by other users and through third party services and websites." See <https://twitter.com/tos?lang=en>.

Philosopher's view

The philosophers reiterated the position that collecting and processing twitter messages for research breaches data protection laws. They considered two options: (1) publishing tweets in full, including meta-data, and (2) de-identifying the messages.

The risks of harm are highest when tweets are published in full, including meta-data like username, location, etc. Therefore, more effort is required by researchers to justify publishing the full tweets. Due to the increased risks, as well the violation of privacy by such an action, data protection laws exist in many jurisdictions to protect users from such violations. Since it may not be possible to consult a lawyer or to research the position of all relevant national data protection agencies, as well as privacy and information commissioners, the following suggestions were made.

First, researchers need to assess whether they are changing the availability, intended audience, and context in which the tweet will be read. This is most likely the case if one publishes information in a different context, such as a research journal or open access platform. Then, a balancing test²⁶ should be applied, such as suggested by the European Article 29 Working Group.²⁷ Overall, researchers need to ensure that they are fairly representing people.

However, if it is unnecessary to publish full tweets, researchers should consider conducting a content analysis for specific terms or themes and subsequently apply statistical techniques to infer data that can be published, rather than publishing tweets in full, including user names and other meta-data. However, if the full messages are important for the publication, tweets could for example be de-identified and the wording of the message could be altered, possibly by translating it to another language and back. Finally, the philosophers stressed that an opt-out is an actual opt-out, and not merely the result of Twitter users missing the message from the researchers.

Emerging themes

Three distinct themes emerged from the discussions: (1) The legal and ethical issues of collecting 'easily accessible' data, (2) balancing tests to assess the severity with which privacy is breached, (3) the information asymmetry and power imbalance between data subjects and technically literate researchers (again).

(1) Status of easily accessible data

The Internet is a technology that inherently provides a platform for vast data collections. The fact that data is (technically) easily collected, does not mean directly that researchers should collect, process, and publish this data without due consideration. Much of the data on the

²⁶ See p. 30 of the "Opinion 06/2014 on the notion of legitimate interests of the data controller under Article 7 of Directive 95/46/EC", accessible at http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2014/wp217_en.pdf.

²⁷ The Article 29 Working Party is made up of a representative from the data protection authority of each EU Member State, the European Data Protection Supervisor and the European Commission. The group gives advice to States regarding data protection.

Internet contains identifiers, such as a username within a tweet or an IP-address in other Internet measurement projects, so is therefore typically protected by privacy as well as data protection laws.

(2) Balancing test

A (legal) balancing test can assist a researcher in deciding whether they have a legitimate interest (a justifying factor in European data protection law) in collecting, processing, and disseminating information gathered on the Internet. This test includes "(a) assessing the controller's legitimate interest, (b) impact on the data subjects, (c) provisional balance and (d) additional safeguards applied by the controller to prevent any undue impact on the data subjects." The Article 29 Data Protection Working Party explains the approach to their balancing test as follows:

"Legitimate interests of the controller, when minor and not very compelling may, in general, only override the interests and rights of data subjects in cases where the impact on these rights and interests are even more trivial. On the other hand, important and compelling legitimate interests may in some cases and subject to safeguards and measures justify even significant intrusion into privacy or other significant impact on the interests or rights of the data subjects."

The philosophers present at the workshop added the following criteria, which should be taken into consideration:

- Consideration of alternative, less intrusive methods of publishing,
- Taking into account relevant law, regulation, and interpretation by authoritative agencies,
- Offer meaningful opt-out procedures where applicable,
- Assess the change in availability and context of the information,
- Considerations of compatible uses according to previous informed consent documents (e.g. a privacy policy),
- Consideration of the power imbalance between actors,
- Consideration whether identified data subjects are represented fairly.

(3) Power imbalance

Finally, technical researchers need to be aware of their ability to collect and analyse data that average persons are typically not aware can be collected about them. This theme will be discussed in more detail in the following case. However, for this case it is relevant to understand to what extent it is ethical to anonymously observe, collect and process online data without interacting with the internet users. Some scholars suggest that such methods can be equated to cyber espionage and surveillance (Kenneally, 2015).

Case 5) Linking wireless devices using Wi-Fi probe requests

Mobile phones make use of the Wi-Fi protocol (IEEE 802.11)²⁸ to access wireless routers as one way to connect to the Internet. Most smartphones send probes for SSIDs²⁹ to which they

²⁸ The IEEE 802.11 protocol is a set of specifications for implementing wireless local area network (WLAN) computer communication. This is commonly referred to as Wi-Fi.

have previously connected, as well as the phone's MAC address,³⁰ which is a unique identifier of the device. Commercial or public parties can collect this data to monitor and profile mobile phone users, for example to infer their shopping patterns, as opposed to the tracking already in place in the digital/online world.

In a series of papers, researchers in France and Australia have attempted to understand the extent to which private information leakage for smartphone users when this transmitted data is collected, with the aim to shed some light on the emerging privacy issues (Cunche et al., 2012). To carry out this project, the researchers needed to engage in the unethical practices that they wanted to prevent. Ethics boards struggle with the highly consequentialist justifications given by researchers for such projects, whereby the researchers hope to justify the highly invasive *means* with the value of the *end* to stop such practices.

Technologist's view

The researcher's objective is to find out whether they can understand the phenomenon of collecting Wi-Fi signals in commercial settings, such as shops or department stores, and develop tools to counteract this data collection. From their point of view, individuals would benefit from such research, because it gives them more control over the data emitted by their mobile phone, and thus their privacy. However, the researchers quickly found themselves in a deontological "Catch 22" situation where ethics boards were contemplating whether those methods can be justified from an ethical point of view.

Some technologists discussed whether data collection via mobile phones could be considered to be a digital equivalent of observational research³¹ in the physical world. Some agreed, while others highlighted that the data emitted by mobile phones reveals much more about a person than could be merely observed and inferred by looking at people in a street. Without a consensus on this point, the discussion took a turn to reasonable expectations of privacy when a mobile phone user switches on their Wi-Fi or Bluetooth connection. Some technologists argued that the user thereby agrees that the emitted data from their mobile phone can be collected, because that's how the technology works.

Philosopher's view

Similarly to the discussion on the Carna Botnet, the philosophers highlighted that even if the research aims may be beneficial, research methods need to be published and thus will set a standard or precedent. Publishing the methods of how to conduct surveillance on mobile phones raises the likelihood of others learning and applying these methods, possibly for malevolent purposes, ranging from more commercial applications to widespread and unheralded surveillance. Questionable methods for benign goals can be misused for

²⁹ The Service Set Identifier (SSID) is the name of a WLAN/Wi-Fi network.

³⁰ A media access control address (MAC address) is a unique identifier programmed into a read-only section of the device's memory.

³¹ Observational research is a methodology whereby the researcher directly observes behaviour. For more information, see: <http://atlasti.com/observational-research/>.

destructive uses. It is therefore important that researchers engage actively with the fact that their methods may be misused, and find ways to mitigate risks and harms.

The philosophers took sides with the engineers who argued that passively collecting data emitted by mobile phones is not equivalent to observational research. Their argument was that most people, including even some of the technical researchers present, were not aware of the amount of information that can be collected via mobile phones. In the physical world, people do understand how their presence and behaviours can be observed and interpreted by others to some extent, including parties conducting observational research. Conversely, the average mobile phone user will not be aware of the extent inferences can be drawn about his or her life based on the data that he or she wasn't aware was being emitted.

Therefore, the philosophers suggest that currently two planes exist between which must be differentiated: the physical environment, and the digital plane. Furthering this line of thought, as well as the lack of comprehension or awareness of the digitally mediated environment and the technologies that support it, the philosophers argue that technologists cannot assume that digital data collection from mobile devices is ethically acceptable without any informed consent. It is unrealistic, unreasonable, and wrong to expect people to agree to the further processing of the data emitted by their mobile device when they switch on the antennae, such as for Wi-Fi or Bluetooth. Expecting mobile phone users to be aware of this would lead to a cognitive overload, which may for example lead to the average user less inclined to act on the knowledge that using Wi-Fi or Bluetooth can be contrary to their best interests.

A final consideration, which is counter to the considerations outlined above, is an assessment of possible benefit of research using this methodology, similar to the fields of information security and cryptology research. If the research in question is going to be used to protect people precisely from the kinds of harms they might be victims of on account of the vulnerabilities of their smartphones that permitted the research in the first place, the benefits may very likely outweigh the risks. The research might be considered ethical particularly if care is taken to address the other concerns – if consent is acquired or if risks are minimised to levels of ordinary life so that consent may not be an absolute requirement, and if details of the methodology and inference mechanisms are not published in ways that might make it likely to be misused.

The philosophers highlighted – again – that technologists need to be aware of their disproportionate amount of power on the digital plane. It is only a small group of people who are aware of the possibilities for inference of the data collected via mobile devices, and an even smaller group who have the technical ability to collect this data. With such power, it becomes important for researchers to adhere to some of the principles of virtue ethics, where consequentialist or deontological justifications should be supported by the demonstrated virtues of the researcher and how these are technically manifested.

Emerging themes

The discussion led to emerging themes about (1) values such as the concept of privacy, (2) the duality of the physical world and the digital plane and the collection of data that some

consider to be "freely available", and (3) the power imbalance between technical researchers and average mobile phone users.

(1) The concept of privacy

While privacy is protected in several human rights statutes, it remains a vague concept to define exactly. The discussions showed that the subset of information privacy, commonly protected by data protection regulation (Greenleaf, 2012), is also a concept that evolves with technological advancements in information and communication technologies (Solove, 2012). Tene and Polonetsky describe how at present the ubiquity of data and technically passive collection methods have shifted the focus from 'acceptable uses' to relying more on ethical judgements (Tene and Polonetsky, 2012). For a comprehensive understanding of information privacy, we refer to papers by such as by Solove, Nissenbaum, Cohen, Ess, as well as many others (see for an overview Smith et al., 2011).

(2) Data collection on the digital plane

The duality of the physical world and the digital plane infused by digital technologies raises the issue of using analogies in the physical world to justify digital data collections. Analogies are prone to misinterpretation that may be intentional or merely misguided. Further, data collection methodologies, for example data emitted from Wi-Fi signals or through Application Program Interfaces (APIs), may be common knowledge and practice for a technical researcher, whilst this may be alien knowledge for most of the population. When collecting such data that is technically easily accessible, researchers should use a balancing test, to weigh the legitimacy of the data collection.

(3) Power imbalance

Finally, the power imbalances created by a lack of technological awareness by the majority of mobile device users must be understood and taken into account. Even if some projects are ethically justifiable, the methods set standards and precedents that equally technically powerful but malevolent actors may adopt. This is slippery slope can lead to unintended consequences.

Analysis

The cases described in this paper, as well as the reflections by different disciplines, provide a set of emerging themes. Below we collect and describe these themes. This section is not meant to establish a definitive set of relevant themes for ethics guidelines for the Internet measurement discipline, but these themes will feed into the broader work of the ENSR project. Each part likely raises more questions than are answered in this paper.

Internet as socio-technical system

The Internet has become an important backbone and central nervous system for many diverging parts of society, commerce, culture, and government. Therefore, the Internet needs to be considered not as merely a technical system, but as a sociotechnical system, in which humans and technical artefacts interact in a complex and dynamic information environment. This system is technically designed and mediated by a relatively homogeneous

group (technical, highly educated, male, Caucasian, from economically developed countries), where design decisions likely embody the ethic of these groups, at least to some extent.

Access to the Internet has been democratised and is more or less available to anyone with a compatible device. Current estimates of the Internet user base exceed 3 billion people (Statista, 2015). It is unlikely that the group responsible for the technical design, subsequent experimentation, and maintenance of the Internet fully understand the relevant social norms, social rules (e.g. laws and regulation), political contexts and its sensitivities of the diverse Internet user base. Through a set of targeted questions, guidelines should make engineers aware of divergences in social contexts, as well as acknowledge their own shortcomings in the knowledge of given contexts, in order to entice them to gain more relevant local knowledge to adequately assess the expected impact in potential target regions.

Responsibilities resulting from power imbalances

In the technically mediated information and communication environment of the Internet, persons who have the skills and knowledge to technically alter the environment or collect data from users can be considered to be relatively more powerful than the average Internet user. This results in a power imbalance between the technically knowledgeable few and most of the population. While a power imbalance is not necessarily and by definition 'bad', engineers and technical researchers, as well as stakeholders³² they engage with, should be aware of this imbalance, and take it into account when making ethical decisions on research projects.

It would be easy to argue that virtue ethics should be applied to Internet research and engineering – where the technical persons must fulfil the character traits of the 'virtuous agent' – research subjects and Internet users will likely not want to rely merely on such self-governance. Responsible researchers have many more duties, such as to inform their data subjects and users about the risks and benefits of a system. Since they are more aware of how data behaves on digital devices and the Internet than the majority of people using them, they cannot assume similar levels of comprehension (see for example Mortier et al., 2014).

Meaningful informed consent

Informed consent is not necessary for every type of research or experiment, since there are several other grounds on which research – even human-subject research – can be conducted. However, when informed consent is necessary, research must make it meaningful to the data subjects. It is commonly agreed that current commercial practices, for example, are not particularly meaningful, when users click on a tickbox to agree with many – sometimes dozens – of pages of dense legalise. A challenge exists for researchers to adequately inform all (possibly hugely diverse) participants about the expected risks and harms, as well as the benefits and the processes of the research in a meaningful way, and in a single document that is relevant for all potential users in their specific contexts. The use of

³² It should be noted here that funders of Internet research, for example, also have more power than the average user, since they can shape the direction of research and possibly even guide the research questions that are being asked.

easily accessible metaphors should be discouraged, since they will unlikely convey the technical complexities adequately and gives researchers the possibility to steer the decision through oversimplification.

Weighing risks, benefits and values for an ethical analysis

During the research design, as well as when drafting the informed consent sheet, the researcher and other stakeholder need to identify the likely risks and harms to participants, as well as their probability and likelihood of occurrence. Before risks and harms can be identified, they must first be defined. Due to the complex, dynamic, and innovative nature of the Internet, it is difficult to define the harms concretely, or to set standards that stand – even a short – test of time. What may be considered harmless today, may become a much larger threat in future when (currently) protected databases are fused, computing power increases to enable new inferences, or new data collection methods enable new ways of identification.

Further, the socio-political contexts of users are prone to sudden or more gradual shifts that cannot be predicted precisely. For example, a war, revolution or a coup in country is difficult to predict years in advance, but equally important to consider are gradual shifts in the attitudes of governments towards data collection, privacy and data protection. Governments are already showing tendencies to use available information to make decisions about citizens in several new areas that were unheard of just a few years ago.

Researchers should articulate the values they wish to protect and how this will be achieved. It is the task of guidelines to assist in defining and identifying risk, harms, and values, and to provide a model or standard to aid in the ethical trade-off between facts and values.³³ Ultimately, however, researchers and stakeholders will need to come to conclusions on a per case basis, but being able to refer to a methodology would help create some certainty in this process.

Status of easily accessible data

Several of the cases addressed during this workshop relied on the collection of data that is technically easily accessible, such as Wi-Fi signals, Bluetooth signals, and Tweets. However, from a legal perspective this information is not freely available *per se*. Usually, this information will be protected by some form privacy law, or more specifically through data protection regulations. In most cases, informed consent will be necessary from the data subjects. However, due to the ubiquity and availability of this data, as well as some potential benefits resulting from its collection, it is near impossible to enforce these laws on all collections of such data. Researchers need to be aware of the limitations to such data collection, the legal obligations, as well as the ethical implications. Methodologies to assess

³³ It should be added that only recently has research into the negative consequences of Internet research been appearing. For example, the study “Experimental evidence of massive-scale emotional contagion through social networks” Kramer et al., 2014 has sparked much debate, but not yet empirical evidence about the negative consequences. Further, it remains an open and broader question how to measure the impact and consequences knowledge that has been created through research. The question itself is not new, but applying this question to Internet research may lead to new issues to be considered.

the desirability to collect technically freely available information exist and should be modified appropriately for the reasoning to make sense in the Internet measurement discipline.

Condoning potentially unethical research methods

Some of the cases presented during the workshop used methods that may be considered unethical or warrant strong scrutiny, as well as redesign in some aspects from a deontological, norms- and duties-based perspective. However, the technologists involved in these cases argued from a consequentialist perspective that the outcomes of the project would be beneficial, particularly when the research could help mitigate the same methods used by others. Further, they argued that commercial actors had less ethical obligations and restrictions when collecting data, so such research could expose the ethical shortcoming of the private sector. Whether a consequentialist justification – whereby the ends justify the means – is sufficient for ethical approval depends on each separate case. However, guidelines could offer some reasoning steps to take. It is ultimately the responsibility of individual researchers, in dialogue with ethics boards and other stakeholders in a specific project, to agree on the limitations, based on a thorough understanding of the project. It is important to consider here the precedent and standard setting effects of condoning ethically improper research during the ethical weighing of risks and benefits.

Further research

This workshop was the first of many that are organised as part of the Ethics in Networked Systems Research project. The project aims to establish a set of ethical guidelines for conducting Internet research. This paper presents observations of the discussions, focussed especially on the arguments made by different disciplines during the day. The aim is to let the themes that emerged in this workshop, as well as other workshops, influence the themes that are addressed in the resulting guidelines.

Each emerging theme warrants a full multi-disciplinary analysis in a separate academic paper, since they raise many difficult issues. As this paper is merely a(n extended) workshop report, we have only scratched the surface of these topics. The next steps include a fuller analysis of the emerging themes, as well as a study into the extent to which they are covered in existing guidelines for Internet research. Finally, the project will formulate a set of questions, which aims to trigger a self-realisation of these themes with technical researchers, rather than provide a checklist of issues to be considered when designing Internet research projects.

Conclusions

The primary aim of the workshop was to shed light on reasoning in ethical justifications by different disciplines involved in technical Internet research, mainly in the field of Internet measurement. In this paper we have outlined how such reasoning by participants from distinctly different academic backgrounds varied, diverged, but also converged when confronted with specific cases. Most notable was the degree in which participants reasoning

from either more consequentialist or deontological approaches were able to come to a compromise that seemed acceptable to all through informed discussion.

The discussions where cross-disciplinary awareness was weakest, or most interesting, have been outlined in this paper. The ENSR project will use these themes to inform the agenda of future workshops, as well as to guide the focus of the resulting guidelines.

The guidelines will be developed specifically for research in the networked systems disciplines, such as Internet measurements. However, the lessons drawn can be applied more widely, for example for online commercial projects and government services via the Internet. Most projects, services, applications, or other ventures on the Internet, the persons designing the underlying technologies and architectures will need to understand the social impact of their creations as well as make decisions together, or come to consensus. Such processes will be informed by several disciplines, or ways of thinking about technology and the Internet, whereby an understanding of different types of reasoning and guidance how to approach such interdisciplinary projects will likely be useful.

References

- ABBATE, JANET ELLEN (1994), 'FROM ARPANET TO INTERNET: A HISTORY OF ARPA-SPONSORED COMPUTER NETWORKS, 1966--1988',
- ANONYMOUS (2012), 'INTERNET CENSUS 2012: PORT SCANNING/0 USING INSECURE EMBEDDED DEVICES',
- ARRAS, JOHN D (1990), 'NONCOMPLIANCE IN AIDS RESEARCH', HASTINGS CENTER REPORT, 24-32.
- BEAUCHAMP, T AND JH CHILDRESS (1979), 'MORALITY AND ETHICAL THEORY', PRINCIPLES OF BIOMEDICAL ETHICS,
- BERNERS-LEE, TIM, ROY FIELDING, AND HENRIK FRYSTYK (1996), 'HYPERTEXT TRANSFER PROTOCOL--HTTP/1.0',
- BYNUM, TERRELL WARD (2008), 'NORBERT WIENER AND THE RISE OF INFORMATION ETHICS', INFORMATION TECHNOLOGY AND MORAL PHILOSOPHY, 8.
- CLOUSER, K DANNER AND BERNARD GERT (1990), 'A CRITIQUE OF PRINCIPLISM', JOURNAL OF MEDICINE AND PHILOSOPHY, 15 (2), 219-36.
- (1994), MORALITY VS. PRINCIPLISM, (JOHN WILEY AND SONS, NEW YORK).
- CUNCHE, MATHIEU, MOHAMED ALI KAAFAR, AND ROKSANA BORELI (2012), 'I KNOW WHO YOU WILL MEET THIS EVENING! LINKING WIRELESS DEVICES USING WI-FI PROBE REQUESTS', WORLD OF WIRELESS, MOBILE AND MULTIMEDIA NETWORKS (WoWMoM), 1-9.
- DANIELS, NORMAN (1996), JUSTICE AND JUSTIFICATION: REFLECTIVE EQUILIBRIUM IN THEORY AND PRACTICE, (22; CAMBRIDGE UNIV PRESS).
- (2007), JUST HEALTH: MEETING HEALTH NEEDS FAIRLY, (CAMBRIDGE UNIVERSITY PRESS).
- DARWALL, STEPHEN L (2003), 'VIRTUE ETHICS',
- DITTRICH, DAVID, ET AL. (2011), 'THE MENLO REPORT: ETHICAL PRINCIPLES GUIDING INFORMATION AND COMMUNICATION TECHNOLOGY RESEARCH', US DEPARTMENT OF HOMELAND SECURITY,
- EDGAR, A. (2005), 'THE EXPERT PATIENT: ILLNESS AS PRACTICE', MEDICINE, HEALTH CARE AND PHILOSOPHY, 8 165-67.
- COMMISSION, FEDERAL COMMUNICATIONS (2015), 'PROTECTING AND PROMOTING THE OPEN INTERNET, FCC 15-24',
- FILASTO, ARTURO AND JACOB APPELBAUM (2012), 'OONI: OPEN OBSERVATORY OF NETWORK INTERFERENCE', *USENIX FOCI*,
- FREY, RAYMOND GILLESPIE AND CHRISTOPHER HEATH WELLMAN (2008), *A COMPANION TO APPLIED ETHICS*, (JOHN WILEY & SONS).
- FRIEDEWALD, MICHAEL, ET AL. (2005), 'PERSPECTIVES OF AMBIENT INTELLIGENCE IN THE HOME ENVIRONMENT', *TELEMATICS AND INFORMATICS*, 22 (3), 221-38.
- GREENLEAF, GRAHAM (2012), 'GLOBAL DATA PRIVACY LAWS: 89 COUNTRIES, AND ACCELERATING', *PRIVACY LAWS & BUSINESS INTERNATIONAL REPORT*, (115),
- HADDADI, HAMED, ET AL. (2008), 'MODELING INTERNET TOPOLOGY DYNAMICS', *ACM SIGCOMM COMPUTER COMMUNICATION REVIEW*, 38 (2), 65-68.
- HARRIS, JOHN (2003), 'IN PRAISE OF UNPRINCIPLED ETHICS', *JOURNAL OF MEDICAL ETHICS*, 29 (5), 303-6.
- HUITEMA, CHRISTIAN (2000), *ROUTING IN THE INTERNET*, (UPPER SADDLE RIVER, NJ: UPPER SADDLE RIVER, NJ : PRENTICE HALL PTR).
- JONSEN, ALBERT R AND STEPHEN EDELSTON TOULMIN (1988), *THE ABUSE OF CASUISTRY: A HISTORY OF MORAL REASONING*, (UNIV OF CALIFORNIA PRESS).

- KENNEALLY, ERIN (2015), 'HOW TO THROW THE RACE TO THE BOTTOM: REVISITING SIGNALS FOR ETHICAL AND LEGAL RESEARCH USING ONLINE DATA', *ACM SIGCAS COMPUTERS AND SOCIETY*, 45 (1), 4-10.
- KRAMER, ADAM DI, JAMIE E GUILLORY, AND JEFFREY T HANCOCK (2014), 'EXPERIMENTAL EVIDENCE OF MASSIVE-SCALE EMOTIONAL CONTAGION THROUGH SOCIAL NETWORKS', *PROCEEDINGS OF THE NATIONAL ACADEMY OF SCIENCES*, 111 (24), 8788-90.
- LALLY, VIC, ET AL. (2012), 'RESEARCHING THE ETHICAL DIMENSIONS OF MOBILE, UBIQUITOUS AND IMMERSIVE TECHNOLOGY ENHANCED LEARNING (MUITEL): A THEMATIC REVIEW AND DIALOGUE', *INTERACTIVE LEARNING ENVIRONMENTS*, 20 (3), 217-38.
- MACINTYRE, ALASDAIR (1985), *AFTER VIRTUE: A STUDY IN MORAL THEORY*, (LONDON: DUCKWORTH).
- MOOR, JAMES H (1985), 'WHAT IS COMPUTER ETHICS', *METAPHILOSOPHY*, 16 (4), 266-75.
- (2005), 'WHY WE NEED BETTER ETHICS FOR EMERGING TECHNOLOGIES', *ETHICS AND INFORMATION TECHNOLOGY*, 7 (3), 111-19.
- MORTIER, RICHARD, ET AL. (2014), 'HUMAN-DATA INTERACTION: THE HUMAN FACE OF THE DATA-DRIVEN SOCIETY', *AVAILABLE AT SSRN 2508051*,
- NARAYANAN, ARVIND, JOANNA HUEY, AND EDWARD W FELTEN (2015), 'A PRECAUTIONARY APPROACH TO BIG DATA PRIVACY',
- NATIONAL COMMISSION FOR THE PROTECTION OF HUMAN SUBJECTS OF BIOMEDICAL AND BEHAVIORAL RESEARCH, BETHESDA, MD. (1978), *THE BELMONT REPORT: ETHICAL PRINCIPLES AND GUIDELINES FOR THE PROTECTION OF HUMAN SUBJECTS OF RESEARCH*, (ERIC CLEARINGHOUSE).
- OAKLEY, JUSTIN (2007), 'VIRTUE THEORY', *PRINCIPLES OF HEALTH CARE ETHICS, SECOND EDITION*, 87-91.
- PELLEGRINO, ED AND DC THOMASMA (1993), *THE VIRTUES IN MEDICAL PRACTICE*, (NEW YORK: OXFORD UNIVERSITY PRESS).
- QUILICI-GONZALEZ, JA, ET AL. (2012), 'UBIQUITOUS COMPUTING: ANY ETHICAL IMPLICATIONS', *ETHICAL IMPACT OF TECHNOLOGICAL ADVANCEMENTS AND APPLICATIONS IN SOCIETY*, 47.
- RAWLS, JOHN (2009), *A THEORY OF JUSTICE*, (HARVARD UNIVERSITY PRESS).
- SADRI, F (2011), 'AMBIENT INTELLIGENCE: A SURVEY', *ACM COMPUTING SURVEYS (CSUR)*,
- SCANLON, THOMAS M (2003), *THE DIFFICULTY OF TOLERANCE: ESSAYS IN POLITICAL PHILOSOPHY*, (CAMBRIDGE UNIVERSITY PRESS).
- SINGER, PETER (2011), *PRACTICAL ETHICS*, (CAMBRIDGE UNIVERSITY PRESS).
- SMITH, H JEFF, TAMARA DINEV, AND HENG XU (2011), 'INFORMATION PRIVACY RESEARCH: AN INTERDISCIPLINARY REVIEW', *MIS QUARTERLY*, 35 (4), 989-1016.
- SOLOVE, DANIEL J (2012), 'INTRODUCTION: PRIVACY SELF-MANAGEMENT AND THE CONSENT DILEMMA', *HARV. L. REV.*, 126 1880.
- STATISCA (2015), 'NUMBER OF WORLDWIDE INTERNET USERS FROM 2000 TO 2015',
- STONE-GROSS, BRETT, ET AL. (2009), 'YOUR BOTNET IS MY BOTNET: ANALYSIS OF A BOTNET TAKEOVER', *PROCEEDINGS OF THE 16TH ACM CONFERENCE ON COMPUTER AND COMMUNICATIONS SECURITY*, 635-47.
- TENE, OMER AND JULES POLONETSKY (2012), 'BIG DATA FOR ALL: PRIVACY AND USER CONTROL IN THE AGE OF ANALYTICS', *NW. J. TECH. & INTELL. PROP.*, 11 xxvii.
- TOULMIN, STEPHEN (1982), 'THE CONSTRUAL OF REALITY: CRITICISM IN MODERN AND POSTMODERN SCIENCE', *CRITICAL INQUIRY*, 93-111.
- WIENER, NORBERT (1954), 'CYBERNETICS IN HISTORY',
- (1964), *GOD AND GOLEM, INC: A COMMENT ON CERTAIN POINTS WHERE CYBERNETICS IMPINGES ON RELIGION*, (42; MIT PRESS).

- WRIGHT, DAVID, ET AL. (2008), *SAFEGUARDS IN A WORLD OF AMBIENT INTELLIGENCE*, (1; SPRINGER SCIENCE & BUSINESS MEDIA).
- YONEKI, E, P HUI, AND J CROWCROFT (2007), 'VISUALIZING COMMUNITY DETECTION IN OPPORTUNISTIC NETWORKS', ... *OF THE SECOND ACM WORKSHOP ON ...*,
- YONEKI, EIKO (2011), 'FLUPHONE STUDY: VIRTUAL DISEASE SPREAD USING HAGGLE', PROCEEDINGS OF THE 6TH ACM WORKSHOP ON CHALLENGED NETWORKS 65-66.