A Theory of Social Machines



Rob Walton
St Anne's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2023

This thesis is dedicated to
Katharine
for being unreasonably supportive throughout, and to
Edward and Joseph
for being most excellent

Acknowledgements

I would not have started or finished this thesis without the support of many great people and institutions. First and foremost, I would like to acknowledge my advisor David De Roure. Thanks Dave, for the good company, great advice and discussions, the occasional pep talk and helping hand, and of course for the offer that day, when I came looking for someone to work with, of the "ecosystem you have to study". I will be always grateful. Thanks also to Justin Coon for the direction and great questions, and for keeping me on a good track while co-advising me. I thank Max Van Kleek and Kieron O'Hara for examining this thesis and for the constructive conversation and questions—your suggestions lead me through the thesis with new eyes—anyone reading this thesis should probably thank them too.

I thank everyone at the Oxford e-Research Centre for the fun times and for providing a creative environment that attracts amazing people from home and far and wide. Thanks to Judy Dendy for everything; to Wes Armour for making the centre a great place to work and an early conversation about experience; and to Ségolène Tarte for early discussions; and of course to Andy, Ania, Ben, David, Fred, Jake, Jamie, Jan, Janet, Jude, Karel, Leia, Min, Philippe and Stef for the many friendly coffee breaks and chats. Thanks to all in Dave's and Kevin Page's research group: David Weigl, David Lewis, John, Toby, Matt, Iain and Athanasios—and particularly to Graham for conversations about knowledge representation and for your persistent interest; and to Petar for involving me in the writing of many interesting papers; and Kevin for the paper on social machines advancing two goals. Finally, thanks to Anna Novotná and Jan for very warmly hosting me in Opava and for the introduction to running experiments with students. Thanks also Anna for you keen engagement and for introducing me to Carol Kuhlthau's work on the information search process.

Thanks to the universally supportive, ambitious, and interesting people of St. Anne's College. You formed a warmness in Oxford; particularly thanks to Alex Rogers, Jakob Foerster, David Murray and Shannon McKella for good conversations and questions.

Thanks to the web science community for your supportive nature and wise feedback, and for caring about the effect technology has on us. Thanks to those

at the research infrastructure for web science workshop, for a great day; and particularly to Jenny Bunn for introducing me to works around sense making.

Thanks to everyone I spoke to about this at Diamond Light Source. Thanks to Jon, Giles, Colin, Tobias, Chris, Steve, Mark, Richard W, and Guy for many good on-project chats; Dave Stuart for this and for encouraging me in my project; and to Andrew for introducing me to Dave De Roure; and to all, thanks for your patience. Thanks particularly to Kaz Wanelik for the chance conversation in a pub where we discovered our mutual interest in Alfred Whitehead; the conviction gained here and in our later conversations about what an engineering Ph.D. about Whitehead's metaphysics might look like, were instrumental in my decision to take this project further.

Thanks to anyone who found something encouraging to say about my interest in this work's earlier roots; particularly Ed Nowicki and Christopher Watson. Thanks to everyone at George Squires' TRLabs and to old friends in Calgary more broadly who supported me while this interest blew me off course the first time around—if you are reading this then you know who you are.

Many people have helped me and my family, whose lives this project has now been a big part of for some time. Thank you to all of you. I greatly appreciate my friends in the village and about; Alan, Richard, Loraine, Mark, Sarah, Kate, Rich, Dave, Val, Julie, Magdalena, Dorian, Barb, Caroline, Dave, Nat, Simon and Nicholas—thanks, especially for help filling gaps I sometimes left in family life. Thanks also to old friends and new on Conrad's well-timed adventure—particularly for help blowing off steam, and for conversations about social Turing machines and the first-person perspective. Thanks to AY for many wonderful open conversations; to Catherine Alexander; and in particular to Lysbeth, who made the idea of turning my interests into a Ph.D. real. More recently, thanks to the Allstar Coders for working together so well.

Thanks Mum for your unfaltering confidence; and Keith for many good walks and for time spent together in the cathartic pleasure of hitting rocks into the Salish Sea. Thanks to Francis for some excellent retreats and supportive times together; and to James and Alice for the frequent use of your bat house. Thank you Pip for your always reassuring words and Kai for a great evening spent on the whiteboard. Finally, thank you to my awesome boys Edward and Joseph: you have been much more helpful than I sometimes made you feel. And thank you Katharine, for being who you are to our family and to us and to me, while my focus on this project took hold—you are wonderful.

Abstract

The modern web is built on services that foster the expression of our social nature to give rise to an ecosystem of overlapping communities. There is limited fundamental theory useful for building these communities, or for modelling and reasoning about the large-scale aggregate effect of our interactions in this ecosystem. Our aim in this thesis is to develop a theoretical model capturing the way that short and long-lived social situations arise and can combine across multiple scales to produce such ecosystems. The inquiry is focused by looking particularly at social machines—as phenomena, both naturally arising within these structures and deliberately deployed into them; and as an abstraction like a Turing machine, but for groups of people. We develop an ontology to describe these ecosystems, and a generative model to represent them and their evolution; and then use these to support a theory of what phenomena together constitute a social machine able to exist in such ecosystems.

The ontology includes overlapping structures of interacting strands of activities that include the forming of shared human experiences. It was developed to support communal attention and shared evolving state; and the emergence of single strands of activity among the interactions of many. We apply it to represent interactions exhibiting elements of shared intentionality, and the evolving structure of social machines that helped discover a new class of galaxy.

The model based on this ontology is implemented so that scenarios can be represented as coloured stochastic Petri nets—and all the tools associated with this representation applied. We model some scenarios that (a) capture, well studied, observable activities shown to result in increased normative commitment to a community; (b) recreate Axelrod's model of the dissemination of culture; and (c) exhibit dynamics of the HTP model, an extension to actor network theory that conceptualises the shifting of connected networks between stable phases of activity as they affect each other's overlapping agendas.

The resulting theory of social machines describes a system that is self-similar, and that not just depends on the social nature of its participants, but that may express this social nature itself. Participants, driven by a desire to form shared experiences, and with a tendency to bind to those that they share experiences with, form structure. Social machines can exist among communities where parts of this structure sustain. This occurs around ongoing processes of sharing experience such as advancing a common purpose. In this way social machines can both coalesce around a purpose, and support the coalescence of their participants around further purposes to form further social machines.

Contents

Notation used **vii**

1 Introduction **1**
1.1 Social machines . . . 2
1.2 Developing a theory of social machines . . . 5
1.3 Theory and model domain and assumptions . . . 8
1.4 Social machines and Turing machines . . . 15
1.5 Modelling approach and thesis method . . . 16
1.6 Argument and presentation . . . 19
1.7 Contributions . . . 25

2 Social machines background **28**
2.1 Definition and situation . . . 28
2.1.1 Researchers and groups . . . 29
2.1.2 Working definition . . . 30
2.1.3 Situation within the social web . . . 32
2.1.4 Situation within the academic disciplines . . . 32
2.2 Lens . . . 34
2.2.1 Types of social machines . . . 35
2.2.2 Ecosystem and lifecycle . . . 35
2.2.3 Purpose, platform and trajectories . . . 36
2.3 Technical aspect . . . 37
2.3.1 Software and hardware stack . . . 37
2.3.2 Abstraction of technical administration . . . 37
2.4 Conclusion . . . 39

3 Background for modelling **40**
3.1 Interaction within a community . . . 41
3.1.1 Sharing states of mind through sharing focus on content . . . 41
3.1.2 Shared intentionality . . . 43
3.1.3 Aggregate social interaction and forming consensus . . . 45
3.2 Emergence of structures of communities . . . 46
3.2.1 Self-similar structure of communities . . . 46
3.2.2 Axelrod's model of dissemination of culture . . . 47
3.2.3 Dynamics of community growth . . . 48
3.3 Petri net modelling . . . 49
3.3.1 Petri net definition . . . 49
3.3.2 Stochastic Petri nets . . . 51
3.3.3 Coloured Petri nets . . . 52

3.4 Whitehead's process metaphysics . . . 53
3.4.1 Introduction: Moments of experience as real things . . . 54
3.4.2 Process philosophy . . . 56
3.4.3 Method of extensive abstraction . . . 57
3.4.4 Philosophy of Organism . . . 59
3.5 Conclusion . . . 60

4 Strands of experience 62
4.1 Ontology . . . 63
4.2 Experience as content integration into states . . . 67
4.3 Simple experience model(G^1) . . . 68
4.4 Composite experience model (G^c) . . . 69
4.5 Extended models . . . 71
4.5.1 Prerequisite experience model . . . 71
4.5.2 Trust or slow experience model . . . 72
4.5.3 Labelled composite content and emergent strands . . . 73
4.6 Discussion . . . 74
4.7 Conclusion . . . 74

5 Collective experience in forums of strands and the emergence of strands 75
5.1 Ontology . . . 76
5.1.1 Overlapping forums of strands . . . 77
5.1.2 Dynamics of strand interaction in forums . . . 79
5.1.3 Dynamics of aggregate strand formation . . . 83
5.2 Examples . . . 84
5.2.1 Shared intentionality . . . 84
5.2.2 Discovery of peas galaxies . . . 85
5.2.3 Social machine frameworks . . . 88
5.3 Forum model . . . 91
5.3.1 Simple experience . . . 92
5.3.2 Composite experience . . . 94
5.4 Experiments and results . . . 95
5.4.1 Flock formation . . . 96
5.4.2 Flock life-cycle and synchronisation . . . 96
5.5 Discussion of results . . . 99
5.5.1 Flock formation . . . 99
5.5.2 Flock life-cycle and synchronisation . . . 101
5.6 Conclusion . . . 103

6 Forums as static networks of forums 105
6.1 Single layer networked forums model . . . 106
6.2 Analysis and application . . . 108
6.2.1 Simplified form . . . 108
6.2.2 Proof of self-similar dynamics . . . 109
6.2.3 Scenario: Two forums connected by a strand . . . 111
6.3 Experiments and results . . . 111
6.3.1 Fully connected network of forums . . . 112
6.3.2 Synchronisation on a grid . . . 112
6.4 Conclusion . . . 114

7 Dynamic networks of forums **115**
7.1 Modelling dynamic binding between forums . . . 116
7.2 Axelrod's model of dissemination of culture . . . 119
7.2.1 Representation with our model . . . 119
7.3 Experiments and results . . . 122
7.4 Discussion and conclusion . . . 125

8 Higher order overlapping forums model **126**
8.1 Model . . . 127
8.2 Visualising overlapping forums of strands . . . 129
8.3 Conclusion . . . 130

9 Persistence through ongoing binding **131**
9.1 Ontology and example application . . . 133
9.1.1 Peas galaxy example . . . 134
9.2 Stable solutions in a dispersive environment . . . 134
9.3 Model: Explicitly purposeful forums . . . 135
9.3.1 Experiencing purpose . . . 136
9.3.2 Purpose filtered forums . . . 137
9.4 Ecosystem boundary scenario . . . 138
9.5 Purpose as recurrent component of composite experience . . . 140
9.6 Model: binding to source of motivation . . . 142
9.7 Example: normative commitment as binding . . . 143
9.8 Conclusion . . . 144

10 Life cycle of persisting forums **146**
10.1 Ontology . . . 148
10.2 Spawning new forums . . . 148
10.3 Coalescence of persisting forums . . . 151
10.4 Future: long and short term binding . . . 152
10.5 Conclusion . . . 152

11 Theory and model of social machines **154**
11.1 Theory of social machines . . . 155
11.2 Idealised social machine model . . . 156
11.2.1 Model operation . . . 156
11.2.2 Topology of social machines . . . 158
11.2.3 Executable models . . . 159
11.3 Model agreement . . . 160
11.3.1 Channel for sharing excitement about content in a forum . . . 160
11.3.2 Channel for sharing each state of mind in a forum . . . 161
11.3.3 Private forum for reacting to content . . . 162
11.4 Model limitations . . . 163
11.4.1 Social machines may sustain a flow of participants . . . 164
11.5 Model implications . . . 165
11.5.1 Social machine ecosystem abstraction from medium . . . 165
11.5.2 A social machine needs a habitat . . . 166
11.5.3 Archives and strand state are inseparable . . . 166
11.6 Conclusion . . . 166

12 Discussion **168**
12.1 Comparison with other conceptions of a social machine 169
12.2 Initial theory and adjustments made during research 172
12.3 Social interaction facilitated network technology 175
12.3.1 Representing a social interaction with the ontology 176
12.3.2 Personal difference in network-facilitated interaction 177
12.3.3 Effects on community structure. 179
12.4 Applications of theory 180

13 Conclusion and future work **183**
13.1 Conclusion 183
13.2 Future work 189

References **193**

A Modelling software **205**

B Log of changes made during corrections and their rationale **213**

List of Figures

1.1 Visualisation of a cross-section through a eukaryotic cell 8

2.1 Software stack supporting social machines on the web 38

3.1 A simple Petri net 50
3.2 Common Petri net elements 50
3.3 A coloured Petri net as displayed in the software package Snoopy 53
3.4 The coloured Petri net of figure 3.3 unfolded onto a graph 53
3.5 Sketch showing a method for representing a trace of occasions of experience 55

4.1 An experience graph showing two components coupled through overlapping content 69
4.2 A simple example of a second order experience generator 70
4.3 A third order experience graph and an approximate form made using social guards 72

5.1 A coloured Petri net describing a fully-connected community, or forum, model 92
5.2 A Petri net model encoding a forum supporting composite experience . . . 95
5.3 Traces showing the emergence of a flock of strands in a forum 97
5.4 Simulations showing the breaking of three flocks and the formation of a single flock 98

6.1 A Petri net model encoding a network of forums 107
6.2 Equivalence between strands on a network and strands in forums 110
6.3 A scenario where one person has strands overlapping with two forums . . . 111
6.4 Traces showing synchronisation across a network of fully connected forums 113
6.5 Traces showing synchronisation across a 3 by 4 grid of mutually connected forums 114

7.1 Petri net model for forum supporting trust or excitement toward neighbouring forums 117
7.2 The experience graph for one feature of our representation of the Axelrod model 121
7.3 Traces showing one forum's changing culture and corresponding chatter . . 123
7.4 Traces showing the changing culture across a grid of forums during an Axlerod model run 124
7.5 Cultures at three times during an Axelrod model run 124

8.1 Scenario showing overlapping forums 127
8.2 Three representations of a scenario with overlapping forums 128

9.1 Graph of purpose for ecosystem boundary scenario 138
9.2 Experience graph for ecosystem boundary scenario 138
9.3 Overlapping structure for ecosystem boundary scenario 139
9.4 Execution traces for ecosystem boundary scenario 139
9.5 Experience graph for ecosystem boundary scenario 140
9.6 Spatial view of hypergraph for ecosystem boundary scenario 141
9.7 Connectivity of hypergraph for ecosystem boundary scenario 141
9.8 A sketch showing a method for representing the binding of strands to their source of motivation when this source is a forum unknown to them 143

10.1 A sketch showing a method for representing the formation of a new forum in response to strands forming a composite experience 150

A.1 Static image of a trace created with the occmodeler software 207
A.2 Screenshot of web application showing the experience generator of a model selected from the library 209
A.3 Screenshot of web application showing model tuning and execution 210

Notation used

Experience generators (chapter 4)

G	The class of experience generators describing the possible paths of experience of strands
G^1	The class of first order experience generators describing the possible paths of experience of a single strand
G^1_{ts}	A first order experience generator built around a labelled transition system, see equation (4.1)
X	A set of internal states resulting from integrating shared experiences, see equation (4.1)
Λ	An alphabet of shareable experiences used to label edges between states, see equation (4.1)
E	A set of transitions $\subseteq X \times \Lambda \times X$ describing the intergration of a sharable experience while moving from one state to another, see equation (4.1)
G^c	The class of cth order experience generators supporting the contribution of up c strands toward a single composite experience
E^2	Pairs of edges taken from E

Collective experience in single forum (chapter 5)

N	The number of strands in the forum, see equation (5.5)
$X^{\text{e}}[q]$	The number of strands excited about the state q, see equation (5.1)
$X[p]$	The number of strands in state p, see equation (5.2)
$\Lambda^{\text{e}}[\alpha]$	The number of strands excited by having integrated the content α, see equation (5.5)
w^{pull}	Weight modifying rate that strands will be pulled along edges by content, see equation (5.5)
w^{self}	Weight modifying rate that strands will spontaneously move along an edge, see equation (5.5)
w^{stick}	Weight modifying the pull that holds strands together in a state, see equation (5.5)

δ Rate constant indicating how quickly excitement is lost, see equation (5.1)

Static networks of forums (chapter 6)

F Connection network with directed edges taken from $S \times S$; a pair (r, s) indicates a connection from forum r to s

F^c Connection weights over edges in F

N_s The number of strands in forum s, see equation (6.2)

$w^{\text{inter-pull}}$ Weight determining how hard neighbouring forums' excitement over content is felt

$X[s, p]$ The number of strands in forum s in the state p, see equation (6.1)

$X^e[s, q]$ The number of strands in forum s excited about the state q, see equation (6.2)

$\Lambda^e[s, \alpha]$ The number of strands in forum s excited about the content α, see equation (6.2)

Dynamic networks of forums (chapter 7)

$T[r, s]$ A measure of trust, or in later models of excitement, that forum s has about r

w^{trust} Weight governing influence of trusted neighbouring forums

Higher order overlapping forums (chapter 8)

F Strand participation in forums; taken from $I \times S$, a pair (i, s) indicates that strand i is in forum s; derived from $F^{hypergraph}$

$F^{hypergraph}$ A hypergraph with vertices representing strands and edges forums, see equation (8.0)

$X[i, p]$ The number of strands i (zero or one in the text) in state p, see equation (8.0)

$X^e[i, q]$ The number of strands i (zero or one in the text) excited about the state q, see equation (8.0)

$\Lambda^e[i, \alpha]$ The number of strands i (zero or one in the text) excited about the content α, see equation (8.0)

Persistence through binding on forums (chapter 9)

$\phi(s)$ A function mapping purpose from Φ onto each forum s, see equation (9.0)

$G^1_{purpose}$ A first order generator with purpose, see equation (9.0)

Ω Directed graph of purpose, see equation (9.0)

Φ Set of purposes, see equation (9.0)

Δ Pairs of purpose with the second advancing the first, see equation (9.0)

Chapter 1

Introduction

1.1 Social machines 2
1.2 Developing a theory of social machines 5
1.3 Theory and model domain and assumptions 8
1.4 Social machines and Turing machines 15
1.5 Modelling approach and thesis method 16
1.6 Argument and presentation 19
1.7 Contributions 25

It is such a natural thing to sit with someone to talk, or to engage briefly while negotiating how to pass in a corridor, that we don't give it second thought. Social situations like these and the shared experiences they encompass come and go continuously throughout our day. They often form part of more permanent social situations in which shared experiences form part of longer-lived stories or advance ongoing processes. Shared experiences arise for example in the situation of a drinks party, or of a family's life, or of a working team or active community that has formed around collective interests. Large scale shared experiences arise at organised situations like church gatherings or rallies; or increasingly in self organising forums on the web, including very large scale politically or ideologically driven communities. These social situations form part of a larger fabric or ecosystem of longer lived social systems—with parts working together or against each other to advance a range of sometimes shared and sometimes opposing interests.

In this thesis we ask how to best represent these ecosystems of short and long-lived social situations in a way that we can reason about their operation. A representation must be broadly applicable, while still specific enough that it leads the modeller to think about the world in a way that helps them use the representation. We focus this inquiry by looking particularly at phenomena of social interaction arising within the social machines described by Prof. Sir Tim Berners-Lee. Many interesting social machines are deliberately crafted. Some consider the concept to be more generally applicable, however. Prof. Dame Wendy Hall stated during the SOCIAM project, introduced in section 2.1.1, that "Social machines are to the web as Turing machines are to computer science"[1]. With this, more abstract, view, the concept scales well to social situations in which many people interact, driven perhaps by more implicit purposes or other drives, to represent communities that form with the help of the web. There is no formal mathematical definition of a social machine akin to the Turing machine. To appeal to users of both of these concepts, we work toward a theory of a more generally applicable 'abstract social machine', and then identify refining characteristics that are useful when applying the concept to understand real phenomena.

An existing approach for representing social situations is actor-network theory [47]. This is a mature system and there are case studies that use it to represent particular situations or events on the web including [41]. We are not aware however of its use to describe parts that fit together to form ecosystems. Couldry, writing in [33, Ch. 6] notices that actor-network theory has not been used that much to describe media and communication technologies in contemporary societies. He argues that this may in part be to actor-network theory's focus on describing static situations rather then those that evolve with time.

1.1 Social machines

In what has proven to be an understatement Prof. Sir T. Berners-Lee introduced the concept of a social machine in 1999 [8, pp. 172–5]:

[1] Conversation with Prof. Dave De Roure

> Real life is and must be full of all kinds of social constraint—the very processes from which 'society' arises. Computers can help if we use them to create abstract social machines on the web: processes in which people do the creative work and the machine does the administration. [...] The stage is set for an evolutionary growth of new social engines. The ability to create new forms of social processes would be given to the world at large, and development would be rapid [...].

There have been a number of concerted attempts to study social machines including, for example, focus on particular machines [6], approaches to classify machines [77], theoretical approaches for observing them [16] and large scale attempts to really observe them [18]. However, arguably no general theory of social machines has emerged or is in common application.

In building a theory of social machines De Roure et al. [16] proposed that in lieu of an observable boundary around a social machine we delineate it from its neighbours based on its common purpose or purposes. These may change with time, and perhaps we may imagine a core collection of such aims cycling in and out to form a longer-lived, possibly implicit or emergent purpose: that of propagating the social machine itself; or a purpose toward cohesion.

There is no doubt that within the social machine community, particularly the web science community, there is scope to move from the working definitions in use, and from the variety of lenses through which we may view social machines to a stronger formalism. Without this, progress in understanding social machines is piecemeal and unsystematic. A formalism would enable progress towards better theory and practice in both understanding and designing social machines. A social machine formalism might for example:

- form a theoretical basis with which to underpin analytical methods in web science and further afield;
- be used to synthesise data with which to test web science methods;

- act as a basis of techniques with which to infer the shape of overlapping social structures and to understand the shared experiences that bind them;
- better understand the systems that result from the large-scale aggregate behaviour of people interacting on the web, for example with respect to their natures as an audiences;
- form a language for describing the operation of social media platforms and the communities they support, which might be useful in legal regulation and to reduce harmful discourse; and
- help us 'break' echo chambers preventing productive discourse.

From more of an engineering perspective a social machine formalism would help answer questions such as:

- How might a social machine or platform better compete for attention?
- How does perception, adaptation and consensus forming occur in a social machine? Is this applicable to distributed AI?
- How might an organisation of people and technology, schedule tasks and meetings between dynamically changing parts of itself to advance ensembles of shifting objectives?
- How might we better push an idea into an individual or parts of a population—and how might we understand when this is happening and, if we choose to, defend against it?

The term *participant* will be used rather than *user* throughout this thesis to describe what was initially conceived as a social machine's human elements[82]. This draws attention both to the fact that these elements are part of the social machine, and that with advances in technology many of these elements contributing to some social machines' operation are technically based. We may vary the ratio of human participants to artificial participants, and also the ratio of work performed by these individuals with that of the administrative functioning of a machine:

- Successfully taking the limit as all human individuals in a social machine are replaced with artificial bots or agents results in an AI agency.
- Going in the other direction by taking the limit as we remove the technical component and afford a 'real world' way of interaction we end up with a purely social system.
- Taking the limit as we remove the ability of two way communication turns a social media machine into a broadcast network.

We may thus consider some social machines as exemplars of other systems and may i) apply insight from studying these other systems to social machines, and ii) better understand these systems by studying social machines. In the latter case we may for example consider a social machine on the web as exemplifying a conventional social system but flayed out and spread apart so that all messages can be recorded.

1.2 Developing a theory of social machines

For the reasons just given, we seek a theory that encapsulates the phenomena of ecosystems of overlapping and sometimes interdependent communities—some or all of which we might choose to identify as social machines. This theory should capture the nature of social interaction facilitated by network technology, that many users of the term "social machine" take as definitional. To this end, we will first choose or propose a theoretical model to describe the kinds of phenomenon that are present in ecosystems of interacting communities. We will then use this theoretical model to precisely describe scenarios that can be viewed as containing such ecosystems; and then, by looking at these models, we can propose a theory of what phenomena together constitute a social machine—and in turn characterise the nature of social machines able to exist in such ecosystems.

Ontology of model of ecosystems of overlapping communities. We need to define an ontology before building a model. An ontology comprises categories of entities and activities with which to describe a scenario and the relationships between these. This need not be an ontology describing all aspects of reality, but one that substantiates just

the domain of our initial subject: ecosystems of overlapping communities. This ontology will include, if not the low-level minutiae driving, for example, social interaction and community cohesion, then the fact that we see these phenomena in reality and can argue that they are present where we see such an ecosystem. Some categories of parts will describe the physical connectivity of this ecosystem and the parts that comprise it; while others will describe the activities that arise on this structure, including social interaction and the sharing of ultimately subjective experiences.

We will start identifying the breadth of the domain and assumptions about this in section 1.3—but the domain should include, for example, shared intentionality (see section 3.1.2) that we assume is core to the functioning of at least some interesting social machines. Each chapter that adds to, or refines, this ontology, starts with a section where we make this explicit, and includes examples of the phenomena categorised.

The resulting ontology forms a basis for building descriptive models of ecosystems of overlapping communities. We will use it in this way during the thesis—and indeed we develop it iteratively alongside these models. It might also, however, be used in the other direction; to design, and ultimately help us build communities, or the architectures of overlapping communities that together form an organisation and its intersection with its environment.

Generative and predictive model of ecosystems of overlapping communities. With an ontology with which to analyse structures of overlapping communities into parts, and to identify traces of activity on these parts, comes a way to make sense of the phenomena of overlapping communities. The models we build can be used in this analytical way. However, they are also executable, in that, a model of a scenario capturing certain structure and past traces of activities, can be used to synthesise activities and changes to this structure. This is quite a normal feature of models of physical systems.

In this thesis, we will limit ourselves to contriving such models by hand—but clearly, in the modern context of increasingly ubiquitous machine learning and other, more tailored, computational approaches to model building, such modelling can be automated. These models, contrived by hand, or trained more automatically, are generative, in that the

changing community structure they exhibit and the activities that arise on this structure will not always be completely predictable by the model maker, or merely regurgitated from training data. Where an instance of this theoretical model is trained to reflect reality, this output will have predictive qualities. It might predict the reactions in a neighborhood of communities to certain kinds of event in the news, for example, with these predictions including activities of discourse that in turn shape community structure. Kraut et al. suggest an approach to handling the ethics of working with this kind of technology[46].

For the primary purpose of this thesis however, which is to develop a theory of social machines, we build these generative models to exercise the ontology before using it as the basis of a theory of social machines.

Theory of social machines. If we choose to consider a social machine to be simply a community, then this theoretical model that describes the nature of overlapping communities *contains* a theory of social machines. The theory of overlapping communities just discussed would tell us how to analyse an ecosystem of overlapping communities into a model, and we would be able to identify regions of this model as social machines.

However, we may choose to refine the definition or characteristics of a social machine as a class of community within this theory. As such, the theory of overlapping communities actually provides a class of possible theories of social machines. Some refinement of the definition may make the theory more applicable, for example, to web scientists. We may, instead, choose to say that a social machine is actually a more generic concept than a community and seek to apply it more broadly. Generalising the definition in this way would make a theory more applicable to computer scientists or engineers building systems that might not include people, or indeed to scientists studying ecosystems with overlapping parts more broadly. It is possible, for example, that the concept of a social machine defined in this broader way, might apply to biological systems such as a Meerkat colony or the cell shown in figure 1.1. We formalise this approach in section 1.3.

[2] https://gaelmcgill.artstation.com/projects/Pm0JL1

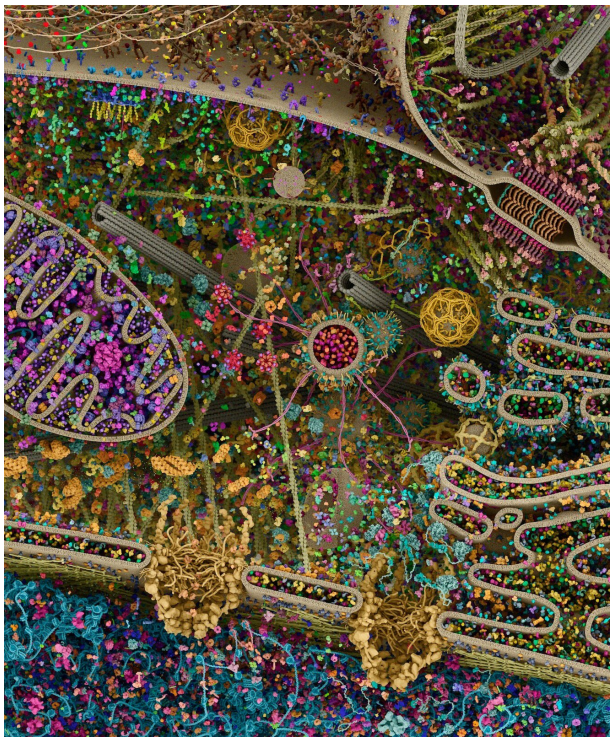


Figure 1.1: Visualisation of a cross-section through a eukaryotic cell. By Evan Ingersoll & Gael McGil, obtained using x-ray, NMR and cryogenic electron microscopy[2]—looking quite social.

1.3 Theory and model domain and assumptions

We grow our ontology and theoretical model throughout the main chapters of this thesis. The direction of travel was chosen based on certain assumptions about the nature of social machines; and some details based on assumptions about the nature of social interaction and community formation. These assumptions and axioms are backed up by work in the two background chapters and throughout the text. We gather them together here.

This work builds on a few layers of assumptions:

A. **Model basis.** These assumptions capture the basis of a theoretical model that we will use to represent the phenomena necessary to support the formation and operation of social machines within an ecosystem of overlapping communities. *These are presented last as they are specified with reference to assumptions about the nature of social machines and their host environment described in layers B and C.*

B. **Generic nature of abstract social machines.** These assumptions capture aspects of the nature of social machines taken from the review of social machines presented

in chapter 2. These are chosen to allow for a more generally applicable conception of a social machine than some authors hold.

C. **Additional assumptions about social machines.** We make our own assumptions about the nature of some social machines that we believe fill in missing elements needed to form a cohesive conceptualisation.

D. **Web Science constraints on social machines.** We separate from (B) some assumptions about social machines that narrow down the nature of social machines to increase the concept's utility to web scientists.

We won't claim that all social machines exhibit all of these qualities, but, as many interesting ones do, a successful theory should ultimately support them. We expand on these four layers of assumptions in the following paragraphs.

B. Generic nature of social machines. These assumptions capture aspects of the nature of social machines taken from the review of social machines presented in chapter 2. These allow for a more generally applicable conception of a social machine than some authors hold.

B1. **Collaborative.** Participants in a social machine may collaborate toward a shared purpose or purposes. This is implicit in much of the literature on social machines.

B2. **Personal or large scale.** Participants in a social machine can socially interact in both small groups and at scales facilitated by network technology. Some large-scale, anonymous, collections of participants may not meet socially in all machines. (See, for example, [74, pp. 7–8].)

B3. **Distributed yet coherent.** A social machine has coherence despite comprising parts that act concurrently (see, for example, [74, pp. 7–8]).

B4. **Part of ecosystem.** A social machine is part of an ecosystem of participants and other social machines. Social machines may be emergent or crafted more deliberately in this ecosystem; and participants tend to be free to roam through parts of the

ecosystem they find accessible. Social machines may be joined from this ecosystem or from the environment outside it—unless perhaps the next assumption, item (B5), is proven to afford no place outside the ecosystem. (See, for example, [16].)

B5. **Abstract.** A social machine is an abstraction applicable to any situation it describes—and not, for example, necessarily existing only on the web or always including human participants. Prof. Dame Wendy Hall stated during the SOCIAM project, introduced in section 2.1.1, that "Social machines are to the web as Turing machines are to computer science"[3].

The last assumption, item (B5) is important as it allows a social machine ecosystem to blend from the web into people and into a broader society. It makes the concept of a social machine applicable to building purely technical systems—and to building architectures of social machines, in which roles undertaken by human participants can be automated by purely technical social machines, and in which the same underlying theoretical model can be used to describe both. To return to the more traditional definition of a social machine, we may later add the assumptions of layer D.

C. Our additional assumptions about social machines. To form a theory of social machines, we assume some additional characteristics.

C1. **Defined by activity or causal process.** A social machine exists where we wish to attribute it as supporting and causing some strand of activities—or process. For example, if a community can plough, then we infer the presence of a social machine that can plough. Just as it would be difficult to draw a boundary around the particular neurons involved in a certain state of activity maintained by a single person, so it would be difficult to section out the neurons, or even the physical people always, involved in a state of ongoing activity held by a community. We will likely find the physical embodiment of a social machine to be open to the environment and its edge hard to pin down due to dependence on the broader physical and

[3]Conversation with Prof. Dave De Roure

social worlds—this leads in part to the following assumption, item (C2), that social machines can be self-similar.

This approach draws our attention away from the assumption that a social machine has a physical boundary at all and abstracts it from its embodiment in particular participants. These types of strands of activity were described by Russel in his work on causal lines [99]. This approach was also taken by Whitehead in his work on process metaphysics introduced in section 3.4. In his system, in which events are more primary than objects, he identifies an enduring object as a collection of events with personal order: none occur contemporaneously and they form a thread of continuous inheritance from one to the next [98, p. 50]. Where some of these events are recurrent, they allow strands of activities, or processes, to carry attributes we would normally associate with objects. Bohm also tried this approach, experimenting with what he called the Rheomode, a way of speaking about reality based around verbs[10]. He would say, for example, that we are aware of something only while it is actively rising into awareness. As an aside, he hoped that this mode of thinking might, if more more broadly applied, act to reduce fragmentation he saw present in individuals and in society.

C2. **Self-similar or interdependent.** The social machine conception should be self-similar in that a machine may form strands of activity that might themselves be attributable to, or at least, contribute to, another social machine; and, in turn, a strand of activity attributable to a social machine may also contribute to another.

For example, following on from the example of assumption item (C1), where a community can plough, we infer the presence of a social machine that can plough. This may depend on a strand of activity supported by a black smith, which we might identify as forming part of a smithing machine. This smithing machine may in turn also be part of an animal husbandry machine that makes—or uses—shoes and nails. There are many ways to cut reality up. When saying that an abstract social machine should be "self-similar", we capture this idea that they may be composed to create something that cannot be done independently. This approach allows researchers

to identify social machines at different levels of granularity—to apply the social machine lens to strands of activity that, were they to step back, clearly form part of a larger strand of activity perhaps associated with larger social machines.

In section 3.2.1, we review evidence supporting the existence of some self-similar structures, including the web, the semantic web, Reddit, and the internal structure of the mind. To back this assumption up further, as described in section 3.1.2, and as discovered by cyberneticians, "intelligent" systems maintain control loops in which they (i) act to bring about a goal in the environment, by (ii) influencing the environment, while (iii) perceiving the environment so see if its state is changing toward the goal state. Many goals can be broken into subgoals, and we may view each such subgoal-driven loop as causing a strand of activities for which we have the choice to infer the presence of a social machine.

This conception of a social machine situates it explicitly in ecosystems of overlapping and possibly interdependent communities, some or all of which may be social machines, depending on how these are defined. A social machine then is not just something that depends on the social nature of its participants but something that is social in itself.

C3. **Coherence implies causal ordering.** For a social machine to be coherent (assumption (B3)) implies that there is some thread of activity that is common across the machine as a whole. For example, social machine frameworks manage interactions in such a way that there is a definite ordering that allows conversation. Where a community maintains a common set of beliefs or shared history, this, too, implies causal ordering—at least where these change over time.

C4. **Self-organising or emergent.** Social machines tend to exhibit self-organisation in that their components are free to form and act on their own intentions and thus effect the purpose of the machine as a whole. (See, for example, [74, pp. 7–8].)

C5. **Ecosystem is dispersive.** As participants are generally free to move their attention between social machines, we assume that if a social machine provides no

stimulus, then participants will shift their attention elsewhere. (See, for example, [16].)

D. Broadly accepted assumptions about social machines that we defer until later in the development. To avoid questions about the definition of social machines, we discuss communities throughout the bulk of our development. The literature describes some characteristics of social machines that we chose to leave out from the more generic assumptions itemised in (B); these are well summarised in [74, pp. 7–8]. Depending on which refining features we choose from below, all such communities may be social machines or just some of them. These assumptions are:

D1. **Has output or does work.** This may in fact be implicit in our other assumptions: where a strand of activity is driven by a goal or purpose, then some aspect of that strand may be considered an output.

D2. **On the web and with some human participants.** It's generally accepted that a social machine has a material realisation on the web. The nature of communities and social machines on the web can be quite different from those off the web but we did not focus on this during this development. See chapter 12 for a discussion of what we learnt during this work about social interaction facilitated by network technology. Keeping with a broader definition of a social machine, however, would allow it to be defined recursively without the definition failing when it hits phenomena wholly within a single person, or not wholly within the web.

D3. **Created deliberately.** For those deliberately designing a social machine or a social machine architecture, it is natural to describe a social machine as something made deliberately. However, if a new social machine forms in an ecosystem of social machines, then, even if this requires deliberate action at a low level, this may also be viewed at a broader scale of granularity as naturally arising.

A. Assumptions behind our theoretical model of collaboration in ecosystems of overlapping communities. In the bulk of this thesis, we develop a theoretical model of

collaboration in ecosystems of overlapping communities. The following assumptions, chosen for compatibility with those itemised in (B) and (C), form the basis for this development:

A1. **The fullest levels of collaboration require shared intentionality**. To support assumption (B1), that participants may work collaboratively, we assume that they must share intentionality, as introduced in section 3.1.2 [87]. We cannot say that a social machine must support shared intentionality—although it is possible that this would make a good defining characteristic, that is always present, at least in some degenerate form—but we can say that our theoretical model, if it supports the fullest potential for collaboration, must be capable of representing social machines that do.

A2. **Activity is primary to structure.** To support (C1) we focus on describing processes, or strands, of activity rather than the material-like systems that further these activities. This also extends to a strand's external environment in that, where a strand of activity related to a social machine must interact with the wider environment, it does so with strands of activity in the environment. See assumption (C1) for existing knowledge supportive of this idea.

A3. **Experiences are discrete and are affected by and affect state of mind**. This allows us to start reasoning with very simple models. See section 3.4.2 for evidence supporting this approach, particularly the discussion about James' view of experience coming to us like droplets of perception in which our understanding of the present grows like buds on a tree; and Whitehead's view that experiences are pulled into our model of the present as much as pushed into it from outside.

A4. **Strands of activity may be self-similar.** To support (C2) a collection of strands of activity can act coherently within another collection as if it were a single strand.

A5. **Participants need not be human.** Modelling approaches suitable for capturing human behaviours and interactions should be useful to capture at least the interaction with machine participants if not their full workings. This choice is pragmatic and allows us to include "participants" in a social machine that, for example, act

to maintain collections of images or archives of discussions (see the examples of chapter 5). This may also support the application of the theoretical model to human-machine interaction, or as a basis for incrementally substituting activities performed by people with automation in an organisation.

We have assumed that the same formalism should describe both communal and individual-person level experience. This idea, that groups of individuals can have communal moments of thought-like experience, is synonymous with the idea that somehow an individual comprised of a distributed system of non-sentient neurons can have a sequence of experiences attributable to the individual as a whole. The view "that mentality is fundamental and ubiquitous in the natural world" is panpsychism—and is a topic that is well explored [29].

1.4 Social machines and Turing machines

Alan Turing came up with his abstract model of computing in a time before there were computers in the modern sense. The Turing machine is not a computer or even a model of a computer but an abstraction of computation. It can be used to apply the techniques of mathematical proof to answer questions about the limits of computation—answering questions for example that would otherwise require us to run an infinite number of programs and test them all. One repercussion of this model was very practical and has far reaching impact to this day. At the time it was assumed that we'd need to build machines customised for each kind of computational task. Turing showed otherwise and gave rise to the general purpose computer that is now so ubiquitous.

Given our assumptions just presented, our social machine abstraction will differ from that of the Turing machine in two very fundamental ways:

1. The social machine abstraction will operate in some formal representation of, potentially human, experience rather than the unrestricted grammar of the Turing machine.

2. The social machine abstraction, or interacting abstract social machines, must support concurrent activities while the Turing machine operates serially.

The subjective world of internal experience is the domain of metaphysics; but scientists have ways of tackling it and our approach must draw on these. At the same time this representation should be something that in some limiting situations, could encode Turing's approach.

As just introduced, a social machine might be delineated by possibly changing purpose. Shadbolt, O'Hara et al. drew together a book of work performed during the SOCIAM project [74]. The authors add the idea that social machines with a purpose are hosted on a social platform that is, at least in theory, without purpose or agenda. This view seems quite reasonable if we think about web-based social machines or work teams. However this is not how we think about computers. These machines are general purpose and it is the software that runs on them that adds purpose. In fact it is the social platform that seems more akin to a general purpose computer. We believe that ratifying these differing views is an important part of formalising the concept of a social machine and hope that this thesis does just this.

1.5 Modelling approach and thesis method

Early work to understand and learn to optimise socio-technical systems such as ours focused on organisations at the scale of businesses. The web and the social structures on it may be considered instances of 'Large-Scale Socio-Technical Systems' identified by Thomas Hughes' system theory [79] and described by Igor Nikolic as "a class of systems that span technical artefacts embedded in a social network, by which a large-scale, complex socio-technical artefact emerges" [59]. At this scale, a lot of research has gone into modelling electricity systems all the way from power generation and distribution up to regulatory frameworks. Modelling these systems is complex in that more than one formalism, possibly described with different languages from different fields of study, is required to capture all relevant aspects of the system. Bringing these together requires simulation in the form of agent

based modelling.

John Holland identifies some agent based modelling as *exploratory modelling* which he describes as: "starting with a designated set of of mechanisms, such as the various bonds between amino acids, with the objective of finding out what can happen when these mechanisms interact" [36]. For example, Schelling showed in his study of segregation that for spatially segregated subpopulations to form there is no need for individuals to avoid others: it is sufficient that individuals have a weak preference to be surrounded by those like themselves [71]. More recently and with access to computers Starnini et al. asked which behaviours are sufficient for metapopulations and echo chambers to emerge in populations of mobile agents [81]. They studied the interplay between mobility, Schelling's homophily and social influence in models of opinion dynamics.

Such exploratory modelling is powerful but models with analytically derivable or empirically measurable properties have historically had more impact. Graph growth algorithms based on properties of generic networks, such as Watts and Strogatz's small-world property [95] or power-law degree distributions, are in common use. Girvan and Newman identified a basic property of community structure in social and biological networks [27]. Such models are powerful, as unlike exploratory models, they provide mathematically solid foundations for further work. For example, algorithms for identifying community structure in networks are often based on properties of classes of randomly growable graphs. In addition to these idealised physical or behavioural based models, statistical models such as those describing the recurrence patterns of words or topics in individual writing or collective discussion [3] are derived from empirical measurement. Holland describes exploratory modelling as a first step toward building such analytical models.

Choice of modelling tools. To ensure that we can make Holland's transition from exploratory modelling to analytical modelling we must use well defined mathematical objects. Although we will use a collection of techniques for capturing a model we will ensure that these models can always be expanded into Petri nets, described in section 3.3. Petri nets, or place/transition nets, are a graphical tool used for the description and

analysis of systems with concurrent actions. Invented by Carl Petri in 1939 for describing chemical processes [64] they were used a lot during the development of early computers. Their representation in graph form provides an intuitive formalism for describing the evolution of systems with distributed state. Beneath this they may be represented as an incidence matrix around which have developed a collection of mathematical techniques for qualitative and quantitative analysis. On the qualitative side there are many approaches for executing a net to result in a particular trace of events. An appealing quantitative approach, applicable for the particular Petri nets we used, is that in the limit as we add more identical parts to a model its dynamics becomes governed by differential equations rather than stochastic ones.

Method. We made a number of assumptions above about the nature of social machines and the ecosystem of overlapping communities these exist within. Throughout the thesis we were informed by the former and drew on the latter while growing an ontology of parts that we propose is sufficient for understanding these phenomena and in turn for modelling them. Our conception of idealised social machines, as emergent structures within these models, grew while doing this. In some cases the behaviours required of these structures seemed inescapable to supporting an ecosystem of overlapping social machines, and in these cases we added them to the characteristics required of a social machine in our theory.

Initial model development was driven by the goal of supporting the emergence of a part with properties like the parts that comprise it. Once we were able to do this we set our first milestone. This was, rather than showing that this model can explicitly fit real world data, to show that the model can recreate the behaviours described by a well known model that certainly does: Axelrod's model for the dissemination of culture (introduced below). Our second milestone was to advance the model until we could model scenarios that captured: first the static structure of parts of the social machine ecosystem, and then their life cycle. To do this we picked out phenomena that occur in the ecosystem of communities found on the web that we felt were key to to its operation; and then sought to represent these in our model, extending it where necessary. We were not able to represent all of these new features with Petri nets however, as the these grew too quickly

to execute using our software. As well as checking that the model can represent Axelrod's models we will have confidence in the model if at end of the development it has familiar behaviours from common social media platforms—some dominant platforms are suggested by Murray-Rust et al. in [57].

Murray-Rust et al. also explain why it is useful to develop executable representations of social machines. The Petri nets we built to implement these models are executable and our goal was that some structures that arise in them should meet the characteristics of a social machine. Given that we've asked of our theory that it might represent a social machine in any media and comprising any form of participant then we could claim that these structures actually *are* social machines. In some sense then we were just trying to build a social machine—using something akin to the iterative model of software development but working with mathematical representations rather than code. In this sense the characteristics required of a social machine by our theory might be viewed by an engineer as a set of requirements or design principles.

1.6 Argument and presentation

The work is presented across a number of chapters, starting with a review of existing work. We briefly describe each chapter here and highlight the main argument of the thesis.

We start by reviewing the social machine literature, focusing on suggestions of definitions, accepted behaviours, and deeper formalisms.

> **Chapter 2. Social machines background.** We conclude that the existing literature forms a solid place to work from, with some of the core concepts and behaviours attributable to social machines agreed upon and backed up by case studies. We form a view that social machines are web-based socio-technical systems in which human and technical parts participate together to realize some process. This process advances a purpose and differing purpose may thus be used to delineate boundaries between machines. This is important as social machines are dynamic entities existing within and forming a larger, equally dynamic, ecosystem. A social machine's purpose

may change with time however suggesting that this alone is insufficient to define a machine's life cycle. Within this ecosystem social machines are deployed or emerge spontaneously, and then may co-evolve. A machine may exist over any time period: short-lived machines might form to tell and appreciate a joke or to organise a party; while longer-lived machines might become essentially permanent places for sharing and reacting to particular kinds of content or expressing cultural beliefs.

Driven by the assumptions proposed in section 1.3 and the choice of modelling tools discussed in section 1.5 we explored concepts that will support our modelling methods.

Chapter 3. Background for modelling.

We are led toward a model in which people are driven to form shared experiences and to share states of mind; and in which these experiences form as we mix new perceptions into existing states of mind. These behaviours have been shown to result in communities that combine their abilities while working together.

The structure of communities on the web was found to be self-similar in some cases and we asked how such structure arises. Two simple mechanisms, for example, tend to cause parts to become more connected and for large communities with a certain purpose to dominate smaller ones with the same purpose. Balancing against these, Axelrod describes a simple mechanism that prevents a tendency for people to converge in their beliefs, attitudes and behaviours from causing us to all end up the same. We described Petri nets in sufficient detail to understand their use in later chapters. We introduced Whitehead' process metaphysics which was the launch pad for our approach to this work—however it is not necessary to follow this to understand this thesis.

Our social machine formalism is built on a more generic self-similar modelling part. Introduced in the next chapter, a *strand* is formed when a serially ordered sequence of experiences appears to carry consistent internal state or properties. To prove the possible existence of strands within our system we will start by asserting their existence and then prove that new strands may emerge naturally as these interact. These emergent strands

may in turn interact and exist potentially at any scale in a self-similar ecosystem. Thus a strand may comprise a collection of strands and may form part of a collection of strands.

These strands interact through shared experience in which the nature of that experience may carry a connotation of directed status, but in which the mechanism of experience sharing itself does not. Nature itself appears to contain nothing physical with such permanence but the idea of an *indivisible strand* will be useful for modelling scenarios.

> **Chapter 4. Strands of experience.** We begin our formalism with a model describing what experiences may form a strand. A model determines how strands' experiences build on each other; with past experiences filtering which objectively shareable experiences raise into awareness and how this is subjectively experienced. In this model experience occurs during the process of integrating shareable content into a strand's state to result in a possibly new state. Some experiences are composite in that to form they require individuals to work together.

Having said that there are parts of our ecosystem that act as strands governed in this way we now need to show that these strands might arise naturally. We start by developing a model of a forum comprised of pre-existing strands that standing alone supports the emergence of a single coherent thread of shared experience—and that might act as a strand in other forums.

> **Chapter 5. Collective experience in forums of strands and the emergence of strands.** We define a forum as a collection of strands that are indistinguishable other than by their capacity to form certain shared experiences. The strands form a fully-connected social network in that they all hear each other and are heard by each other equally. Each strand may be excited to be in any state of mind and separately excited by any shareable content that has been integrated. Strands share excitement about content even when this moves them to incompatible states of mind. These levels of excitement are ceded to the whole such that strands cannot tell their own excitement apart from their neighbours'. We define dynamics that can cause strands to follow each other along paths of experience. A drive forward to integrate

content that the whole is excited by is balanced by a drive to stick with those excited by being in compatible states of mind that form in reaction to this content. We modelled a number of scenarios and showed that if this excitement is allowed to decay that these rules support the emergence of flocks of strands that move together. Thus we show in principle that these forums of strands can have the same strand nature as the indivisible strands they comprise, and might themselves act as strands in other forums.

We then expand the model to support a network of interacting forums. This is done in a way that closely connected forums similarly situated in a broader network may act like a single forum in this broader network. This is a useful approximation to the 'overlapping forums of strands' model introduced in the last chapter that we will return to in chapter 8.

Chapter 6. Forums as static networks of forums. We added a following network to the model allowing each forum of strands to follow another with a certain weight. Just as strands in a forum cede their excitement about their state of mind and of content recently experienced to the whole they also share a common connection to the strands in the forums their forum follows. This was done carefully such that the dynamics governing a number of forums on a fully-connected following network are the same as if we had put those forums' strands into a single forum. We prove this and also model scenarios that demonstrate it experimentally. We then show through simulation that the same flocking behaviours that arise in a single forum can occur across a network of forums.

At this point have shown that we can contrive scenarios in which a forum of pre-existing strands may may act as one strand and that a network of forums following each other may act as one forum. We now need a mechanism that evolves the network in way that these contrived structures form naturally.

Chapter 7. Dynamic networks of forums. We further expanded the model such that when a strand integrates content from an external forum its host forum becomes more strongly bound to that external forum. Just as a forum's excitement

about content or being in a certain state of mind decays so does this interest in external forums. This model assumes that the host forum already has a connection to the external forum—we address this later in chapter 9.

We now verify this binding mechanism and exercise the entire model. We need to show that forums emerge across regions of a network of forums to form a new network with fewer edges. We therefore looked for a scenario that encouraged the growth of a few larger forums—one in which differing paths of past experience did not prevent convergence while at the same time did not always converge to a single forum without differences.

Chapter 7. Dynamic networks of forums (continued). Robert Axelrod published his influential work "The Dissemination of Culture: A Model with Local Convergence and Global Polarization" in the journal of Conflict Resolution [5][4]. He proposed an answer to the question of why a tendency for people to converge in their beliefs, attitudes and behaviours does not result in peoples differences disappearing. He showed through agent based modelling that competing cultures will arise in a simple model in which people are more likely to take the beliefs of people whom are already more similar than them. We built a scenario that could be represented by Axelrod's model and showed through simulation that it tended to produce larger emergent forums.

The presence of these emergent forums may be inferred through looking at the strength of the ties on the following network. To build more complex models though we need more control than this.

Chapter 8. Higher order overlapping forums model. We ask how we might explicitly define such higher order forums in our models. We return to a model of overlapping forums introduced in section 5.1.1 by allowing the low level pre-existing strands of our models to exist in more than one explicitly defined forum. We expand on a hypergraph representation in which strands are represented as vertices and the forums that connect them as hyper edges.

[4]This work has been cited over 2300 times at the time of writing

In developing the forum model so far we needed forums to emerge naturally and chose a solution to this which requires a strand's connection to a forum to decay if not used.

Chapter 9. Persistence through ongoing binding. We observe that forums in our model require an ongoing process of binding to sustain. In turn we realise that if our model represents our theory closely that a social machine also needs an ongoing process of binding to sustain. We describe a few different strategies that forums might use to achieve this. Participants may have different experiences but if just some component of those experiences is common then participants may become excited by this and this may cause an ongoing process of new experience formation and subsequent binding. Picking up from the social machine literature, we focus on advancing a shared purpose as such a binding mechanism. We extend the binding mechanism so that when a few strands contribute to a composite experience they will bind to the forums that motivated each other even if not previously aware of them. This is important as it provides a way for strands to join new forums and in turn for forums to increase their overlap with each other. Acting as an example scenario we showed that a stable boundary may form between two forums where one advances a more specific purpose than the other.

Finally our theory needs to address the life cycle of social machines.

Chapter 10. Life cycle of persisting forums.

We add three last requirements to the theory: that a social machine may act as a platform to start new social machines; that this platform will act to aid the coalescence of social machines that share compatible binding processes; and that successful social machines will form part of an ecosystem in which separate social machines that could coalesce will tend to have components in a common platform.

To support the first of these requirements, we add one final behaviour to our model: every composite experience creates a new forum. This new forum may sustain if it gains an ongoing binding process. We observe that sustaining forums will coalesce naturally in the model due to a positive feedback loop that forms. We suggest that

natural selection in which ecosystems that bring together forums with compatible binding processes will out compete those that don't.

Our model is complete for now, and in it we identify a sustaining forum as a social machine.

Chapter 11. Theory and model of social machines. We summarise our theory by gathering up the requirements we used while developing our model and finessing them into defining characteristics of a social machine. We then summarise our model without reference to the mathematics we used to define scenarios. The self-similar definition of a forum raises some immediate questions. In addressing these we reason that a forum is comprised of other forums but that it cannot be said to fully contain them—and the corollary, that where a forum comprises other forums it cannot be said to be fully contained by them. We found that our idealised model matched the main behaviours of a few well known social media platforms but also found some limitations. Finally we look at some consequences to the idea that a social machine exists only while it maintains an ongoing binding process.

To wrap up, in chapter 12, we compare our social machine theory against the concept of a social machine developed while exploring the literature in chapter 2. We conclude in section 13.1 and make some suggestions for future work in section 13.2. Last of all, appendix A includes a brief description of the software we developed to build scenarios and to explore their behaviours.

1.7 Contributions

We believe that this work shares some new contributions:

1. **An ontology and idealised model supporting overlapping communities.** We model the broader situation in which structures with the theoretical characteristics of social machines arise. This model is presented, without the aid of mathematics, in section 11.2. An interesting property of this model is the way it addresses the

self-similar characteristic of social machines required by the theory—and the nature of the topology of the resulting structure of overlapping machines.

2. **A theory of self-similar social machines.** We propose that systems characterised by this theory are social machines capable of sustaining in ecosystems of overlapping communities. This theory is general enough to apply to a range of phenomena that need not include people or depend on the web. Key to supporting the self-similar nature required of social machines by this theory is the characteristic that they exist only while their participants are bound together by an ongoing process of forming shared experiences. This might occur while advancing a shared purpose, or while maintaining a corpus of knowledge or story fragments through sustained use. This definition limits the scope of what we might consider part of a social machine. We discuss these repercussions in section 11.5. This theory is presented in section 11.1.

We exercised this model by developing and testing a concrete implementation using a Petri net. We could not execute Petri nets large enough to test the features added to the model in chapter 8 and beyond however; we explain why we have confidence in this model despite this in section 13.1.

3. **Concrete implementation of model.** This is one of many possible implementations of the system for modelling scenarios in which social machines arise. A model of a scenario includes the following defined modelling parts: a generative model representing possible paths of individual and composite experience; a collection of pre-existing indivisible strands that through interacting in forums connect those forums into a structure with topology described in section 11.2.2; and stochastic rate equations determining how strands interact in these forums and how this effects their ties to those forums. These modelling parts are combined to produce a single stochastic coloured Petri net. This Petri net is executable and has many useful analytical properties. This implementation of the model is sufficient to recreate Axelrod's model of the dissemination of culture. This concrete model is developed in chapters 4 to 7. Were features from chapter 8 and on added to this concrete model,

then some of the stable structures that emerged while recreating Axelrod's model would not just represent social machines, but according to our theory, would actually *be* social machines.

4. **Software tools for defining and executing these models.** This software is described in appendix A. These tools provide a way to define particular models of experience and structures of forums using Python code. These models can be visualised and 'compiled' into a stochastic coloured Petri net. The tools provide a way to execute this model and to visualise traces showing the changing state of strands and their ties to forums. The software includes a web based GUI for exploring the models described in this work and some others we used during verification.

In conducting this research we have developed a diagrammatic approach for representing Whitehead's occasions of experience which has proven useful. This is illustrated in figure 3.5.

We published two full length papers during this work:

1. Robert Walton and David De Roure. "Modelling Web Based Socio-Technical Systems Through Formalising Possible Sequences of Human Experience". In: *Proceedings of the 10th ACM Conference on Web Science*. WebSci '19. New York, NY, USA: ACM, 2019, pp. 301–310. ISBN: 978-1-4503-6202-3. DOI: 10.1145/3292522.3326049.

2. Robert Walton and David De Roure. "Petri Nets for Modelling Communal Flocking along Paths of Possible Experience". In: *13th ACM Web Science Conference 2021. WebSci '21*. New York, NY, USA: Association for Computing Machinery, 2021, pp. 53–57. ISBN: 978-1-4503-8525-1. DOI: 10.1145/3462741.3466668.

These papers cover the main argument of chapters 4 to 6. We indicate when a section includes parts of these papers. The first forms an early iteration of the work presented here, but parts of it have been used in the presentation. The second is included here almost in its entirety.

Chapter 2

Social machines background

2.1 Definition and situation 28
2.1.1 Researchers and groups 29
2.1.2 Working definition 30
2.1.3 Situation within the social web 32
2.1.4 Situation within the academic disciplines 32
2.2 Lens 34
2.2.1 Types of social machines 35
2.2.2 Ecosystem and lifecycle 35
2.2.3 Purpose, platform and trajectories 36
2.3 Technical aspect 37
2.3.1 Software and hardware stack 37
2.3.2 Abstraction of technical administration 37
2.4 Conclusion 39

We start by reviewing the social machine literature, focusing on suggestions of definitions, accepted behaviours, and deeper formalisms. We must form a solid conception of a social machine to build on.

2.1 Definition and situation

Social machines are loosely considered to be comprised of people and the technical systems that draw these people together and administer, mediate or even orchestrate their

interactions. These technical components and individuals may form part of more than one social machine.

2.1.1 Researchers and groups

There is a risk that the term *social machine* acts to define a community of academic researchers as much as it does a type of entity which is thoroughly unique. The concept attracted funding from the UK's Engineering and Physical Sciences Research Council (EPSRC) to spawn the SOCIAM project[1]. The goal of this project was to advance the theory and practice of social machines. The principle investigator was Sir Nigel Shadbolt, with co-investigators at Oxford (Prof. Sir Tim Berners-Lee and Prof. David De Roure), Southampton (Prof. Dame Wendy Hall and Prof. Luc Moreau) and Edinburgh (Prof. Peter Buneman and Prof. Dave Robertson). The funding period for this project is complete, but it continues to spawn academic papers. Work generated through this project forms much of the core thread within this literature review and is summarised in a book [74].

Smart et al. defend the importance of the term social machine [78]:

> By identifying a set of mechanisms and processes that are at the core of the Social web, the notion of social machines serves as a conceptual anchor for research efforts associated with the nascent discipline of web Science (Berners-Lee et al., 2006). In addition, by focusing attention on web-based systems that are involved in the mediation or material realization of social processes, the notion of social machines serves to emphasise the socio-technical nature of the web, and it provides the basis for multidisciplinary collaboration with the social scientific community. Such collaboration is of vital importance given the increasingly significant role the web plays in the functioning of contemporary society.

MIT hosted the *Laboratory for Social Machines* from their Media Lab[2]. The group "uses natural language processing, network science, machine learning, and user experience

[1] https://sociam.org

[2] https://www.media.mit.edu/groups/social-machines

design to conduct analyses and build tools that promote deeper learning and understanding in human networks". This is reflected in their publications which are not focused on understanding social machine's per se, but on data mining to understand the human communication networks which social machines help administer, and to gain a view into society. The project was funded by Twitter. In 2021 it was merged into MIT's Center for Constructive Communication[3].

2.1.2 Working definition

There is no conclusive definition of a social machine, but most researchers would clearly agree that certain systems are social machines. They may not agree on all systems and some might identify a social machine where another may not.

As described in the introduction, Prof. Sir T. Berners-Lee and Fischetti introduced the concept of a *social machine* in 1999 [8, pp. 172–5].

> "Real life is and must be full of all kinds of social constraint—the very processes from which 'society' arises. Computers can help if we use them to create abstract social machines on the web: processes in which people do the creative work and the machine does the administration. [...] The stage is set for an evolutionary growth of new social engines. The ability to create new forms of social processes would be given to the world at large, and development would be rapid [...]."

Murray-Rust and Robertson give a great example of social constraint: "In Britain, social norms prevent people in queues from talking to each other; at Edinburgh airport, there is often a large queue for taxis, and in a small city it is very likely that many people are going to nearby locations." They then go on to discuss solutions for ride sharing [56].

Smart and Shadbolt critique Berners-Lee and Fischetti's characterisation [78]. They observe that defining what constitutes *creative work* and what is administrative work is difficult and perhaps more importantly note that with advancement of technology the idea

[3] https://www.media.mit.edu/groups/center-for-constructive-communication

that technology cannot perform creative work is limiting. In later work Smart et al. cite examples of technology performing creative work including PicBreeder[4], and Wikipedia's ever more intelligent system of machine-learning powered bots, which perform tasks such as preventing vandalism or observing edit wars and either responding or flagging it up to a trusted human editor [65]. They proposed the following working definition:

> Social machines are web-based socio-technical systems in which the human and technological elements play the role of participant machinery with respect to the mechanistic realization of system-level processes.

They have removed any distinction between human and technological elements from the original working definition.

They go on to quote from Clark's 'Supersizing the mind' to give some more detail as to how they conceive *participant machinery*:

> This notion of human and machine elements serving as forms of participant machinery ... takes its inspiration from an approach to mind and cognition that sees extra-organismic resources as (on occasion) participating in the material realization of human mental states and processes – such resources are deemed to "form part of the very machinery by means of which mind and cognition are physically realized and hence form part of the local supervenience base for various mental states and processes" [13, p. 207].

The authors observe that Clark was describing the mind of a single person extending into the environment (which in our case will comprise their immediate physical surroundings as well as that of the web) where as in a social machine there is more than one person involved. Perhaps we might consider a social system as a collections of minds joined by their mutual extension into the same environment.

We will take Smart and Shadbolt's working definition, although will not discard the original conception as this draws our attention that in all likelihood the technical elements *will* be the ones doing the administrative tasks such as keeping track of social networks,

[4] https://picbreeder.org

routing messages, and keeping a record of artefacts. We must be wary that other author's works discussed below were working off the original definition.

It is important to notice that Berners-Lee's definition runs deeper in some senses then Smart's, drawing our attention to the empowerment social machines give to society to form new social processes. Also, it is important to note, as Tarte et al. have, that the people forming part of a social machine must be considered participants rather than users, along with data, algorithms and computing or networking infrastructure [82].

2.1.3 Situation within the social web

Social machines clearly relate to what some call the social web. Halpin et al. in their work for the W3C incubator group toward *A standards-based, open and privacy-aware social web*, describe the social web as 'a set of relationships that link together people over the web' [1]. Paul Smart et al. situate the concept of social machines within other terms that have been used by researchers to describe the operation of the Social web [78]. These terms include: *crowd sourcing* [21], *human computation* [66], *collective intelligence* [48], *social computing* [62], and *the social operating system* [72].

2.1.4 Situation within the academic disciplines

Social machine research is very multidisciplinary. It is useful to understand how it fits in within various scientific disciplines. Researchers from each discipline will bring different perspectives, interests and questions about social machines. Here we will mention their work briefly.

Social science and the Humanities. Social science is the study of society and individual's operation within it and includes for example economics; and the humanities include academic disciplines such as philosophy and history. There is a lot of interest in the operation of networks of interacting individuals, for example actor-network theory [47], and in the underlying mechanism of sociality. *Socio-technical systems* including how these are constituted is also an established topic of study.

If we consider that social machine administration include managing the network through which communication occurs, the communication itself and artefact keeping and sharing, then as we take the limit as administration goes to zero we end up with individuals unable to communicate through the social machine. This is the domain of the social sciences and the humanities; this includes the study of social networks through network science providing techniques and ways of looking at things drawn in part from, mathematics and physics as well as computer science. There is a wealth of literature describing social processes. Adding the web in gives us the social web, the study of which led to the conception of social machines.

The SOCIAM project, described in section 2.1.1, pointedly included work from the social sciences and humanities. This was done to shift the focus from the technical aspects of social machine operation. For example, Tarte et al. describe how to model archetypal narratives in social machines' [82]; and Tart et al. explain the operation of social machines with the analogy of stories [83].

Computer science. The *machine* part of the term *social machine* is used by computer scientists in different but related ways. It describes a mathematical model of computation, or an automota, describing how a set of outputs are computed from a set of inputs. Examples include Turing's universal computing machine, the finite state machine, the pushdown machine, and the Petri net. It is also often used to describe a physical computer implementation of, or software that operates like one such as a virtual machine. The computer scientists' abstraction of a machine need not be fully deterministic, and, with the advent of distributed computing and the promise of quantum computing, work has been done to expand the theory into this area. A collection of people and hardware may be considered to act like a machine capable of supporting social processes. However, when we try to pin down something as fundamental as a social machine's boundaries, yet alone its state or its inputs and outputs, we quickly run into difficulty.

In traditional computing engineers and computer scientists would say that a computing machine runs computer programs in one or more processes; where these processes are hosted in a computer. A *distributed program* can be run on a *distributed system* of

networked computers running processes which communicate via messages. We have no such clear definitions for social machines, but following the suggestions in sections 2.1 and 2.2 a social machine might be considered analogous to the distributed system and in some readings the program running on this as well.

Hybrid automata are used to describe systems comprised of a discrete, machine like component, existing in an analogue environment [4].

Web science. Tiropanis et al. provide a good summary of the differences between the study of web science, internet science and network science[86]:

> Web science and internet science aim to understand the evolution of the web and the internet respectively and to inform debates about their future. These goals lead to different priorities in their research agendas even though their communities over overlap. Network science aims to understand the evolution of networks regardless of where they emerge: the internet as a network transforming and forwarding information among people and things, and the web as a network of creation and collaboration.

2.2 Lens

We must build our intuition of a social machine, and choose a lens or set of lenses through which to view them, prior to being able to describe their components or principles of operation.

Conceptual work on social machines includes an understanding of social machines as complex systems forming ecosystems where individual social machines might be delineated by purpose, and which undergo shared transitions where these align with each other's or with a user's purpose.

2.2.1 Types of social machines

Shadbolt et al, developed a taxonomic framework for social machines [75, 77]. They interviewed researchers in the field and built up a list of qualities that social machines may have; and then categorised a wide selection of systems to create a taxonomy. This list of qualities is extensive. Together they capture the nature and roles of a community, the kind of purpose it advances, the mechanics of interaction and contribution and the cohesion of these contributions. However the list of qualities are focused on individual machines. In the context of our work, it stands out that these characteristics do not capture how a social machine fits into a broader ecosystem. For example, characteristics might also have included: mechanisms for how participants find a machine; what portion of participants' time is spent in a machine compared with others; if contribution is aided by participation in other machines; and generally how connected or disconnected a machine is from its place in a broader ecosystem.

2.2.2 Ecosystem and lifecycle

In preparation for a project to build a web Science Observatory [85] a part of the larger *Theory and Practice of Social Machines* (SOCIAM) project De Roure et al. asked how social machines should be observed to understand them better. In doing so they came up with a 'lens' with which to view them. They came to an understanding of social machines as dynamic entities existing within and forming a larger, equally dynamic, ecosystem [17]. Within this ecosystem social machines are deployed or emerge spontaneously, and then may co-evolve: at once varyingly dependent on each other and competing for the attention and future attention of users. The connections between social machines may be one off encounters or more stable channels, or form more of a hierarchical relationship, with for example a Twitter feed acting as social machine largely within the Twitter platform. These machines may be as short-lived as a trending twitter hashtag or as permanent, perhaps as the web itself, or at least as the big social machine platforms like Twitter, Reddit, Wikipedia, Facebook or Uber. The authors notice that social machines are complex systems, in that

> a successful social machine, whether intentionally designed or serendipitously emerging, has all the properties of an emergent system; the fabric itself of its constituting parts mutates under their mutual influence, as do the interactions between them.

To back up this ecosystem view, in some cases social machines may form long term symbiotic relationships. For example, until May 2016 images in Reddit posts were all hosted by other websites [23]. In other cases this is more dynamic with for example social machines, probably through their individuals, observing content on another social machine and sharing it with their users, some of whom after possibly altering it might feed it back to the first. Social machines may also act to create or gather content which is then shared more broadly within the broader social machine ecosystem and out into its environment through word of mouth or other forms of media.

2.2.3 Purpose, platform and trajectories

In building a theory of social machines De Roure et al. proposed that in lieu of an observable boundary around a social machine we delineate it from its neighbours based on its common purpose or purposes. These may change with time, and perhaps we may imagine a core collection of such aims cycling in and out to form a longer-lived, possibly implicit or emergent purpose: that of propagating the social machine itself; or a purpose toward cohesion. It has been proposed that when observing and recording the operation of the social machine ecosystem one might focus on sequences, or trajectories, of interactions driven by an individual's (or another social machine's) end goal or purpose [61].

> In our previous work [17] we proposed purpose as a key mechanism in the study of social machines. We extend that notion here to one of a trajectory, each of which is defined primarily by such a purpose; that is, an end for which a path through elements of the social machine is taken.

Where an individual and a social machine of which they form a part have a common purpose at any moment then they will work together to advance that purpose. They gave

as one example reCAPTCHA which is used by websites to verify that a client is being operated by a human [92]. In this example the human may have the purpose of verifying themselves to be human and reCAPTCHA the dual purpose of both verifying the client and also solving another task such as character recognition.

Some researchers add another layer to the definition saying that social machines with a purpose are hosted on a social platform that is, at least in theory, without purpose or agenda [74]. The term platform is often used to describe services such as Twitter of Facebook. We'll break this down in the next section. For the level of abstraction we are interested in here this idea brings to mind that a platform is like a group of apes sat around relaxing. Social machines form on this platform around common purposes such as grooming or fending off predators. The participants are bound by working together and this binding applies to the platform as well as the machine.

2.3 Technical aspect

There is confusion as to whether Facebook for example is a social machine. As just described it is more useful to say that Facebook provides a platform for social machines. It helps when discussing social machines to understand the technology stack behind these platforms.

2.3.1 Software and hardware stack

Figure 2.1 shows the software stack of the part of the modern web in which social machines make their home. Social machines fit within a broader ecosystem ranging from clearly technical to clearly social layers; existing across the Services and Causes/Groups layers but dependent upon the rest.

2.3.2 Abstraction of technical administration

Noticing that social machine platforms each have their own set of hidden rules describing what messages can be sent, when and between whom, but that this is hidden from social

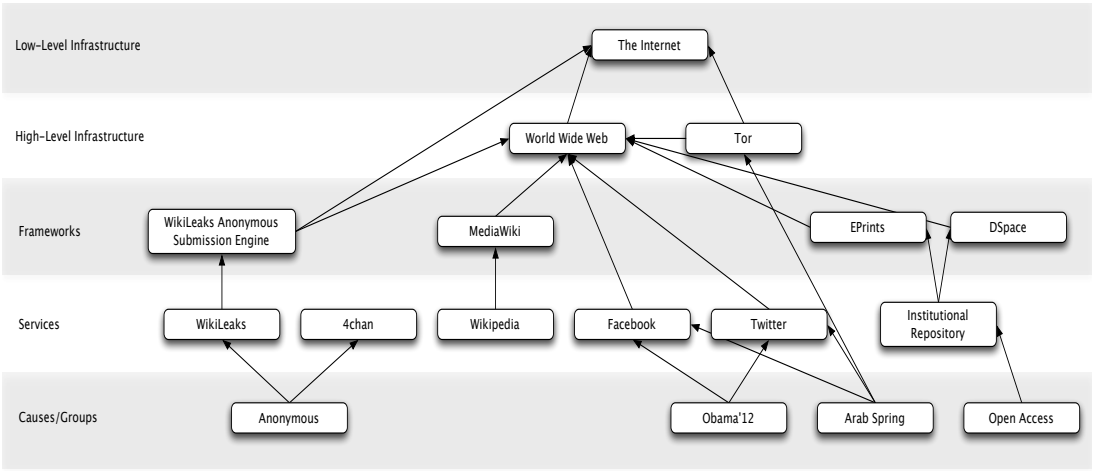


Figure 2.1: Software stack supporting social machines on the web. Social machines exist across the Services and Causes/Groups layers [77].

machine users, Murray-Rust and Robertson have developed a *Lightweight Social Calculus* (LCS) as a standard *interaction model* for describing these rules [56]. Although these models might be used to build social machine platforms from scratch, the authors propose using them to build social machines as *shadow institutions* bound onto existing social network services. *Institution* here is shorthand for the concept of an *electronic institution* introduced by d'Inverno et al. to the AI community: "Just as conventional institutions are organisational structures for coordinating the activities of multiple interacting individuals, electronic institutions provide a computational analogue for coordinating the activities of multiple interacting software agents" [15]. Murray-Rust and Robertson are thus likening a social machine to an institution. They are proposing the LCS as a lighter way of formalising an electronic organisation than that used by the AI agent community; one which is suitable for formalising human interaction rather than of software agents.

The LCS provides a useful description of the interface between the technical and human components of a social machine. LCS is an extension of LCC, a Prolog-like language. It would be interesting to ask, with the more dynamic life-cycle of social machines described below, if the emergent social machines described below are describable by an LCS specification of their host social machines. An algebraic representation might be useful for example when a social machine emerges across multiple social machines that can be represented by different LCS descriptions.

On this front Meira et al. "define a notion of social machine and envisage an algebra

that can describe networks of such" [51]. It is not clear that they had worked from, or were aware of, Berners-Lee's conception of a social machine, and it is unclear how the work they present may be applied. However, as described, their vision of an algebra is an interesting one. Nascimento et al. from the same group later went on to present the *Social Machine Architecture Description Language* [58]. This is an abstraction for describing interactions between social machines. However their conception of a social machine is of what we identified in section 2.3.1 as a *service*.

2.4 Conclusion

We conclude that the existing literature forms a solid place to work from, with some of the core concepts and behaviours attributable to social machines agreed upon and backed up by case studies, but that there is work remaining to reach a formalism strong enough to be represented mathematically.

We have formed a view that social machines are web-based socio-technical systems in which human and technical parts participate together to realize some process. This process advances a purpose and differing purpose may thus be used to delineate boundaries between machines. This is important as social machines are dynamic entities existing within and forming a larger, equally dynamic, ecosystem. A social machine's purpose may change with time however suggesting that this alone is insufficient to define a machine's life cycle. Within this ecosystem social machines are deployed or emerge spontaneously, and then may co-evolve. A machine may exist over any time period: short-lived machines might form to tell and appreciate a joke or to organise a party; while longer-lived machines might become essentially permanent places for sharing and reacting to particular kinds of content or expressing cultural beliefs.

The operation of a social machine platform may be defined by an interaction model. The software aspect of an implementation of this model can be called a framework. These frameworks are deployed as services by companies and non-profit groups who are said to provide a platform for social machines. Social machines with a purpose then are hosted on social platforms that are, at least in theory, without purpose or agenda.

Chapter 3

Background for modelling

3.1 Interaction within a community 41
3.1.1 Sharing states of mind through sharing focus on content 41
3.1.2 Shared intentionality 43
3.1.3 Aggregate social interaction and forming consensus 45
3.2 Emergence of structures of communities 46
3.2.1 Self-similar structure of communities 46
3.2.2 Axelrod's model of dissemination of culture 47
3.2.3 Dynamics of community growth 48
3.3 Petri net modelling 49
3.3.1 Petri net definition 49
3.3.2 Stochastic Petri nets 51
3.3.3 Coloured Petri nets 52
3.4 Whitehead's process metaphysics 53
3.4.1 Introduction: Moments of experience as real things 54
3.4.2 Process philosophy 56
3.4.3 Method of extensive abstraction 57
3.4.4 Philosophy of Organism 59
3.5 Conclusion 60

We open up our literature review to explore concepts that will support our modelling methods. This is driven by the initial assumptions made in section 1.3.

In section 3.1 we ask how the interaction between parts of a community might be modelled and what drives communities to form. Then in section 3.2 we ask how the

interaction of very many similar parts can result in the emergence of multiple separate communities of those parts; starting by presenting evidence that these communities can overlap to form self-similar structures. Having explained in section 1.5 why Petri nets are a good choice for formalising our model, in section 3.3 we briefly mention some other options before then describing Petri nets and the variants we use in some detail. Alfred Whitehead's process metaphysics was the launch pad for our approach to this work. We introduce this in section 3.4 although there is nothing in this presentation that is key to understanding the work presented in the rest of this thesis. We conclude in section 3.5.

3.1 Interaction within a community

In the next section we will ask what drives the formation of structured ecosystems of communities. Here we consider how to model the interactions between members of an existing community. Firstly looking at interactions between individuals and then at the aggregate effect of interactions between many similar individuals.

3.1.1 Sharing states of mind through sharing focus on content

At some level people can be said to communicate through the exchange of packets of information or messages. This says nothing though about what drives this communication or its effect. If we began modelling at this level there would be a lot more detail to account for. Given our assumptions in section 1.3 (A), we seek a representation of human experience. In this section we look at methods for doing this and then ask what drives us to share experience with one another.

We need a single way to represent a very broad range of experiences, experienced 'by' a broad range of systems. We need to represent scenarios, for example, in which you and a friend form shared experiences while deciding who will open a door, or what the title of your thesis will be. This approach must scale to situations where larger coherent groups of people experience things together. Some parts may be less visible such as echo chambers; while some maybe familiar but hard to separate: for example the parts of yourself that

believe one thing or another, or the parts that form around your work life or home life.

These scenarios are abstract enough that its hard to always identify long-lived objects that experiences happen to. Sometimes we can put this aside and understand the experiences themselves as best described as bits of process. So for example if you've left something unfinished at work there is a process that when revisited will continue more or less where it left off. These little bits of process form during individual moments of experience, but they relate. William James [38, Ch. X] describes new experiences as coming like droplets and that it's not until an experience has formed that we can go back to consider the ingredients that somehow were selected and assembled to form it:

> Either your experience is of no content, or of no change, or it is of a perceptible amount of content or change. Your acquaintance with reality grows literally by drops or buds of perception. Intellectually and on reflection you can divide these into components, but as immediately given, they come totally or not at all.

Whitehead describes a theory describing how occasions, or droplets, of experience form; which we introduce in more detail in section 3.4. He conceives no pre-existing 'screen' onto which these droplets of experience are projected. In his theory, these are instead continually and actively pulled as much as pushed into a model of the present which is always based on prior experience. In this way a person's state of mind effects the way they experience content; and their state of mind is in turn effected by this experience.

Although experience is private and subjective, philosophers and scientists do have ways to describe it and reason about it. The well studied concept of *intentionality* [37] is used to formalise the concept of a state of mind. That mental states have intentionality means that they are directed toward some object or state of affairs; where mental states include for example thoughts, beliefs, desires and hopes.

Many scientists believe that the ability to copy each other's behaviours and experiences is so key to our survival that it is in part hardwired physiologically into many animals' nervous systems. Rizzolatti and Craighero reviewed and popularised the idea that there are a particular class of neurons related to visual perception and motor control called

mirror neurons that fire both when an animal performs an action or just when it sees another animal performing the same action [70]. We discuss this more in the following section.

3.1.2 Shared intentionality

We assume in assumption (A1) that the fullest levels of collaboration require the emergence of shared intentionality between participants in a social machine. Tomasello et al. give a good review of the subject, which we summarise here[88]. The authors describe a model of intentional action based on cybernetics in which control system theory is applied to organisms that work toward goals in a closed loop by changing their environment and perceiving the effect they have. A desired result, such as an open box given a closed box, requires an internal representation as a goal that guides the organism's behaviour. An intention comprises this internally represented goal, a chosen plan of action and a commitment to achieve this goal. To choose this plan of action, an agent needs stored knowledge and skills, as well as a model of current reality. This intention drives behavioural action, which includes both manipulation of the environment and the direction of attention onto the environment. An action plan chosen as part of an intention will often require subgoals with action plans of their own. In turn, many goals are chosen to be part of a larger goal. For example, the goal of having an open box is likely a subgoal chosen as part of a larger action plan chosen to reach the goal perhaps of accessing something inside the box.

To work together to their fullest potential, Tomasello et al. believe that people need to 1) understand the intentional action of others and 2) participate in activities involving shared intentionality.

To understand the intentional action of others requires an understanding that another person is driving an action, an understanding that this person has a persistent goal they are pursuing while doing this, and to understand that the person has chosen the action plan. This ability is an important part of learning through imitation, as it allows us to recognise that the same goals can be reached by imitating an action plan used by others.

It is also necessary to understand why the people around us do certain things and if those reasons apply to us; as such, it allows us to learn the conventions of our culture.

Shared intentionality may arise when individuals who understand each other in this way as intentional agents interact socially. The authors write:

> Shared intentionality, sometimes called "we" intentionality, refers to collaborative interactions in which participants have a shared goal (shared commitment) and coordinated action roles to pursue that shared goal [...]. The activity itself may be complex (e.g., building a building, playing a symphony) or simple (e.g., taking a walk together, engaging in conversation), so long as the interactants are engaged with one another in a particular way. Specifically, the goals and intentions of each interactant must include as content something of the goals and intentions of the other.

The authors identify preconditions for joint cooperative activities to arise:

> (1) the interactants are mutually responsive to one another, (2) there is a shared goal in the sense that each participant has the goal that we (in mutual knowledge) do X together, and (3) the participants coordinate their plans of action and intentions some way down the hierarchy – which requires that both participants understand both roles of the interaction (role reversal) and so can at least potentially help the other with his role if needed.

The authors identify three levels of social engagement. At the first level, an infant may share the emotion of performing the same activity together with an adult and take turns acting, but not share any internal state. At the next level, the infant and an adult may share goals and monitor each others' perceptions (for example, by attending to the same parts of external reality) as these goals are pursued. The fullest collaborative engagement occurs when the participants share joint intentions (internal goal, decided plans of action and commitment) rather than just goals, and joint attention in the form of mutual knowledge rather than just monitoring each other perceptions. Action plans at this level may be complementary with, for example, one interactant holding a tower steady while another adds a block.

Tomasello et al's interest in shared intentionality arose while asking what allows humans to operate together in ways that other primates don't, for example, in the building of social institutions or bridges. They identify four pairs of skills [87]. Each skill has an individualistic version and one based on shared intentionality: gaze following and joint attention; social manipulation and cooperative communication; group activity and collaboration; and social learning and instructed learning. Together, they suggest, these skills allow humans to share intentionality and to form groups that work together toward shared purpose.

3.1.3 Aggregate social interaction and forming consensus

We have proposed that a social machine, despite being inherently concurrent, must potentially host a single overarching thread of experience, so that it may participate with coherence in other social machines. For example: even as culture changes across a society parts of that society will change their beliefs together; and in deciding whether to turn left or right at a junction a crowd of friends will stick together. The way we form consensus has been well studied since DeGroot introduced his influential model [20]. More directly applicable for our idealised model however are models of flocking within animals.

In 1986 Craig Reynolds made 'Boids' a computer program that models coordinated animal motion such as that of bird flocks or schools of fish [68]. The models exhibit emergent behaviour. The rules each individual follows are simple and are applied locally but the resulting flock appears able to work together to avoid obstacles for example. The rules are simple; each part must:

1. **Separation**: steer to avoid crowding local flock mates
2. **Alignment**: steer towards the average heading of local flock mates
3. **Cohesion**: steer to move toward the average position of local flock mates

Just as these rules allow animals to act coherently as one whole while moving through space, perhaps similar rules might govern the way groups of people choose paths of common experience from almost uncountable options.

3.2 Emergence of structures of communities

Having just looked at how collections of very similar parts may form a coherent whole, we must now ask how structures of such communities might form. Alan Turing asked how structure might form with parts as simple as chemical reactants. In his work on morphogenesis, he was the first to describe how two simple reactants can result in structure rather than uniformity across space [50, 89, 90]. This question is well studied when it comes to social systems.

In this section we look at evidence suggesting that the form this structure takes on the web is sometimes self-similar; and at models of the mind as structured collections of interacting processes with similar form. We then introduce Axelrod's model for the dissemination of culture which explains how, despite a tendency to follow the beliefs of those around us, society does not end up with a single uniform culture. Finally we describe two simple mechanisms that cause communities to grow.

3.2.1 Self-similar structure of communities

Self-similar community structure has been found in networks of human interactions [30] and on the web [14]. Tim Berners-Lee described the scale-free nature of the web quite early on in its development [7] and then drew attention to this property in the semantic web [9]. If communities form around particular topics or beliefs we would expect these communities to share similar structure. Olsen and Zachary used a backbone network analysis technique [73] to generate a map of Reddit communities related by topic of interest [60]. They found Reddit to exhibit scale-free, small world and modular community structure. Network scientists find that many complex networks are well understood in terms of their self-similarity. Song et al. analysed a variety of real complex networks and find that they consist of self-repeating patterns on all length scales [80] suggesting in general less tendency for hierarchy to form than in the small-world nature of Reddit.

Marvin Minsky theorises that cognitive systems occurring in nature, including in a person's mind, are formed from a society of simpler interacting processes [54]. This suggests that the self-similar structure of communities may be viewed as extending down

into individual peoples' minds. Herman et al.'s dialogical self theory [34, 35] proposes one such model. This offers a useful view of a possible internal structure to the mind particularly suggesting that this reflects external social ties.

In this model an individual self may be modelled as a number of internal *I-positions* each of which may carry its own narratives, goals, and memories; and external I-positions that represent the I-positions of people in that individual's social group. These I-positions form a social network within the individual. As they converse, certain I-positions rise from dormancy and in and out of dominance based on the activity of other I-positions and of external events. A number of methods for inferring the shape of this internal structure have been developed.

3.2.2 Axelrod's model of dissemination of culture

Robert Axelrod proposed an answer to the question of why a tendency for people to converge in their beliefs, attitudes and behaviours does not result in peoples differences disappearing. He published his influential work "The Dissemination of Culture: A Model with Local Convergence and Global Polarization" in 1997 [5]. He showed that competing cultures will arise in a simple model in which people are more likely to be influenced in their beliefs by those that they already share other beliefs with. Axelrod takes culture to be what social influence influences. He represents a culture as a list of features or dimensions; where each feature can take on one of a number of possible traits. He gives the example of belt colour as a feature with each possible colour being a trait. Where people hold opinions about politically debated questions or beliefs, those beliefs are features and the possible opinions are traits. Features also include political issues whose traits are defined positions.

He then simulated scenarios in which people interacted on a square grid of sites giving each interior site 4 neighbours. He started everyone off with randomly assigned traits. At each step in a simulation a site is chosen at random and then one of its neighbours is chosen at random. With a chance proportional to how similar their cultures are the person on the site will align on one of the features they differ on. Axelrod ran simulations

on 10 by 10 grids with varying numbers of features and traits. With five cultural features and ten traits per feature the simulation often reached stability after around 80,000 events leaving two or three regions of shared culture.

3.2.3 Dynamics of community growth

The dynamics of how social networks grow to form structure is well studied. More generally, social networks have been well studied both on and off the internet [11, 67, 28]. Jin, Girvan and Newman propose an algorithm describing the growth of simple social networks that often applies in nature [40]:

> We propose some simple models of the growth of social networks, based on three general principles:
>
> 1. meetings take place between pairs of individuals at a rate that is high if a pair has one or more mutual friends and low otherwise;
> 2. acquaintances between pairs of individuals who rarely meet decay over time; and
> 3. there is an upper limit on the number of friendships an individual can maintain.

This describes one mechanism driving the formation of communities. Another is a positive feedback mechanism that drives the emergence of a large dominant community from smaller ones sharing the same interests: The term 'network effects' was coined by Theodore Vail in explaining why Bell should be allowed a monopoly on US telephone services. 'Network effects' exist in a market, if as networks grow they become more valuable to their users. An individual may join a social network or social network platform with a goal of only gaining something, but in doing so they give back to the network as a whole making it more valuable to existing and potential members. Katz and Shapiro developed economic theory describing the network effect considering it to occur in *systems markets*: economic "markets where consumers use compatible components together to generate benefits" [42], in contrast to, for example, commodities markets.

3.3 Petri net modelling[1]

Having explained in section 1.5 why Petri nets are a good choice for formalising our model, we describe them here in sufficient detail to understand their use in later chapters. We first introduce basic Petri nets and a few useful extensions, and then describe how stochastic Petri nets may be used to model processes that are effected by time. Finally we describe coloured Petri nets which offer a succinct way to model systems with many repeated parts.

Petri nets, or place/transition nets, are a graphical tool used for the description and analysis of systems with concurrent actions. Invented by Carl Petri in 1939 for describing chemical processes [63] they were used a lot during the development of early computers. There are more general formalisms including process algebras such as Milner's π-calculus but we feel Petri nets provide a good balance between ease of use and modelling power. Their representation in graph form provides an intuitive formalism describing the evolution of distributed systems such as the large-scale socio-technical systems on the web we seek to develop methods to model. Beneath this they may be represented as an incidence matrix around which have developed a collection of mathematical techniques for qualitative and quantitative analysis.

3.3.1 Petri net definition

As shown in figure 3.1 a Petri net is a bipartite graph comprised of places, shown as circles, and transitions, shown as rectangles, connected with directed edges. Figure 3.2 show the standard elements of common Petri net definitions. The state, or *marking*, of a Petri net is represented by tokens inside the places and shown as black disks. A transition *fires* if the places feeding into it through the net's directed arcs all have tokens. This firing is a single atomic event in which a token is removed from each input place and a token is placed in each place connected by outgoing arcs. Tokens need not be conserved. If more than one transition is enabled then the order of firing is random. Basic Petri nets have no encoding of time. If the Petri net of figure 3.1 were simulated t_a would first fire causing the tokens in x_1 and x_2 to disappear as a token appears in x_a. The transition t_b would

[1] This section was published almost exactly as it is in [93]

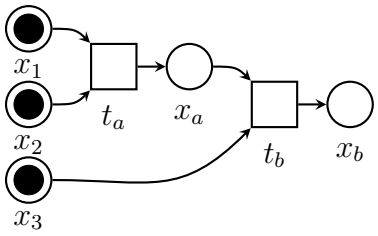


Figure 3.1: A simple Petri net. A Petri net is a bipartite graph comprised of places, shown as circles, and transitions, shown as rectangles, connected with directed edges. The black disks represent tokens.

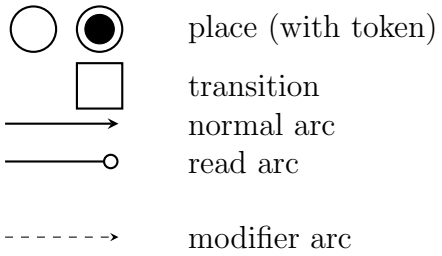


Figure 3.2: Common Petri net elements

then fire causing the tokens in x_a and x_3 to disappear as a token appears in x_b.

A Petri net PN may be represented more formally as places P, transitions T, input and output arcs A, and an initial marking M:

$$\begin{aligned} PN &= (P, T, A, M) \\ P &= \{p_1, p_2, ..., p_l\} \\ T &= \{t_1, t_2, ..., t_m\} \\ A &\subset \{P \times T\} \cup \{T \times P\} \\ M &= \{\mu_1, \mu_2, ..., \mu_n\} \end{aligned} \tag{3.1}$$

where μ_i are the integer number of tokens at a place [2].

A net may be *simulated* resulting in a possible sequence of firings. There are also many analytical techniques for testing if certain markings are reachable for example, or if a net will remain live in that it won't settle to a fixed state, or if the number of tokens at any place is bounded [2].

3.3.2 Stochastic Petri nets

Petri nets run without time, the result of a simulation being simply a sequence of changing markings or log of firings. Molloy extended the model of Equation (3.1) to include a collection of transition rates $\Lambda = \{\lambda_1, \lambda_2, ..., \lambda_m\}$ [55]. Where λ_i are, possibly marking dependent, rates for exponentially distributed transition firing times:

$$f(\tau|\lambda) = \begin{cases} \lambda e^{-\lambda\tau} & \tau \geq 0 \\ 0 & \tau < 0 \end{cases} \tag{3.2}$$

In the resulting continuous-time stochastic Petri net, when a transition becomes enabled an activation time, τ, is chosen randomly and the transition will fire only if the transition remains enabled for this time. The effect is that an enabled transition will fire continuously and independently at a constant average rate while enabled—acting as a *Poisson point process* with average rate λ. The transition is memoryless: the chance of it firing is unrelated to when it last fired.

A transition's firing rate may be specified as a function of the number of tokens in the places that feed it. The firing rate function of a transition may also depend on the number of tokens in places connected into it through *modifier arcs* [32]. These modifier arcs are shown in figure 3.2. If a transition fires then tokens in places connected to it by modifier arcs will not be removed. Also, unless the firing rate function requires there to be a token in such a place, the transition can fire even when there is not.

The transition firing rates of stochastic Petri nets ultimately only affect the ordering of transition firing of the standard Petri net, but in doing so they open up new interpretations of models. Molloy showed that with some limitations a stochastic Petri net with exponentially distributed firing rates is isomorphic to a finite Markov process. These models may also be converted to a set of ordinary differential equations which may be solved numerically, or in simple or symmetric cases, analytically. For example, stable or relatively stable states may be identified using approaches from non linear systems analysis.

3.3.3 Coloured Petri nets

Kurt Jensen introduced a generalisation of Petri nets and showed that this can be used to describe and analyse complex systems [39]. In this extension, tokens may have different *colour* or type. Places and arcs are given expressions limiting which coloured tokens may enter or traverse them. Transitions may be given optional guard functions which determine if the combination of coloured tokens at their input places meet constraints additional to these.

Coloured Petri nets then can be used to model systems succinctly with repeated structure. In the field of systems biology for example Gao et al. used Coloured Petri Nets as a method for modelling multicellular organisms across multiple scales [24]. They applied this approach to modelling planar cell polarity (PCP) signalling in a Drosophila wing. PCP is the coordinated alignment of cell polarity across a plane of tissue, describing for example the emergence of patterns on the butterfly wing or of our fingerprints.

As an example, figure 3.3 shows a coloured Petri net rendered by the Snoopy software introduced in appendix A. The 'Unit' identifier on the places is the set of colours those places may take, in this case the colours $\{1, 2, 3\}$ representing three people. The red text describes, as a multi-set, the colours of the tokens in each place: x_0 has 1 token of colour 2 and one of colour 3, while x_a has one of colour 1. The transition t_a accepts a token of any colour u from the set *Unit* and moves a token of the same colour to place x_a. Without the read arc this coloured Petri net would unfold into three separate units. The read arc presents tokens to the transition via the variable n. The *is_neighbour* function on the transition guards it from firing unless the colours of the tokens in variables u and n meet some criteria.

Walton and De Roure used a function that enables the transition only if the token in n and the token in u are from neighbouring cells [93]. Figure 3.4 shows this model unfolded onto the linear graph $n_1 \longleftrightarrow n_2 \longleftrightarrow n_3$. This would be done by specifying $is_neighbour(u, n)$ as $(u = 1 \& n = 2)|(u = 2 \& (n = 1|n = 3))|(u = 3 \& n = 2)$.

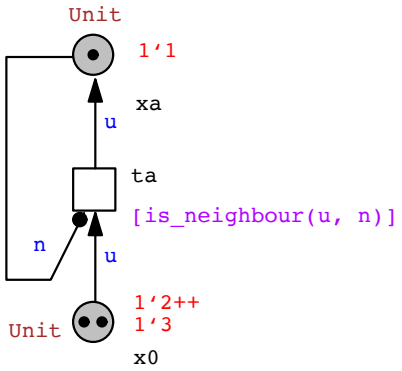


Figure 3.3: A coloured Petri net as displayed in the software package Snoopy. Figure 3.4 shows the same structure.

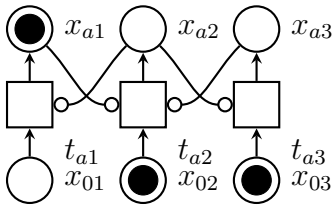


Figure 3.4: The coloured Petri net of figure 3.3 unfolded onto a graph. The graph is $n_1 \longleftrightarrow n_2 \longleftrightarrow n_3$

3.4 Whitehead's process metaphysics

Alfred Whitehead's process metaphysics was the launch pad for our approach to this work. Whitehead's writing is often hard to follow and so we introduce his ideas here in a way that helps with the rest of the thesis; however nothing here is required to understand what follows. We hope that this section might give the reader more of a first person view into the models developed in this work—and an appreciation for what informed our modelling process and where future work might take these models. For those interested Donald Sherburne offers a guide [76] to Whitehead's seminal book 'Process and Reality' [98].

We found Whitehead's process metaphysics helpful because: i) it includes a formal system capable of explaining how events of sense-awareness form and relate thus suggesting a way to handle this in our formalism, and ii) it questions our intuitive idea that to understand a behaviour we must identify an object with some permanence exhibiting that behaviour. We believe that taken together these ideas provide a useful way to understand social ecosystems in which defining solid material-like boundaries between parts is difficult.

We start with a sketch of what we're trying to represent with Whitehead's formalism.

Then we describe the essence of his earlier attempt to build a process rather than material theory of natural knowledge; and then move on to his later work in which he constructed arguably the most complete process-centric metaphysics, his *Philosophy of Organism* which should fill in the details of our sketch.

3.4.1 Introduction: Moments of experience as real things

Let's examine a social situation between people to start. Recall a conversation from a party or the workplace. Imagine where the other participants in this group were situated and how they were oriented. Then imagine that below you are a series of snapshots falling away into the blur of the past, with each capturing a shared experience. Perhaps someone was welcomed into the group, perhaps someone said something and everyone laughed or perhaps a thread of debate arose between two of the people with the rest looking on; or perhaps this thread of experiences broke into separate discussions. Process philosophers argue that those shared experiences are instances of the fundamental building blocks of reality—that reality is built up from things that are happening rather than from material objects.

Whitehead goes so far as to call these moments of experience, or *occasions of experience*, the living *creatures* of his *Philosophy of Organism*. After an experience has formed future experiences may draw from the perceptions that went into it. While a new occasion takes shape different combinations of these perceptions mix together to form the subjective experience that will be left as the occasion ends. For example, where a selection of perceptions rises up in the awareness of all it may be met with common agreement or laughter—and the process of other experiences on the way toward the same level of awareness may be forgotten. Rather than social entities calling to each other it is as if past experiences try to call forward to future experiences; while at the same time as the group seeks rich experience it listens backwards seeking past experience, or parts of past experience, to bring forward. In this way experiences that arose outside this group may also be pulled in and shared communally.

Figure 3.5 is a sketch showing a method for showing a trace of occasions of experience.

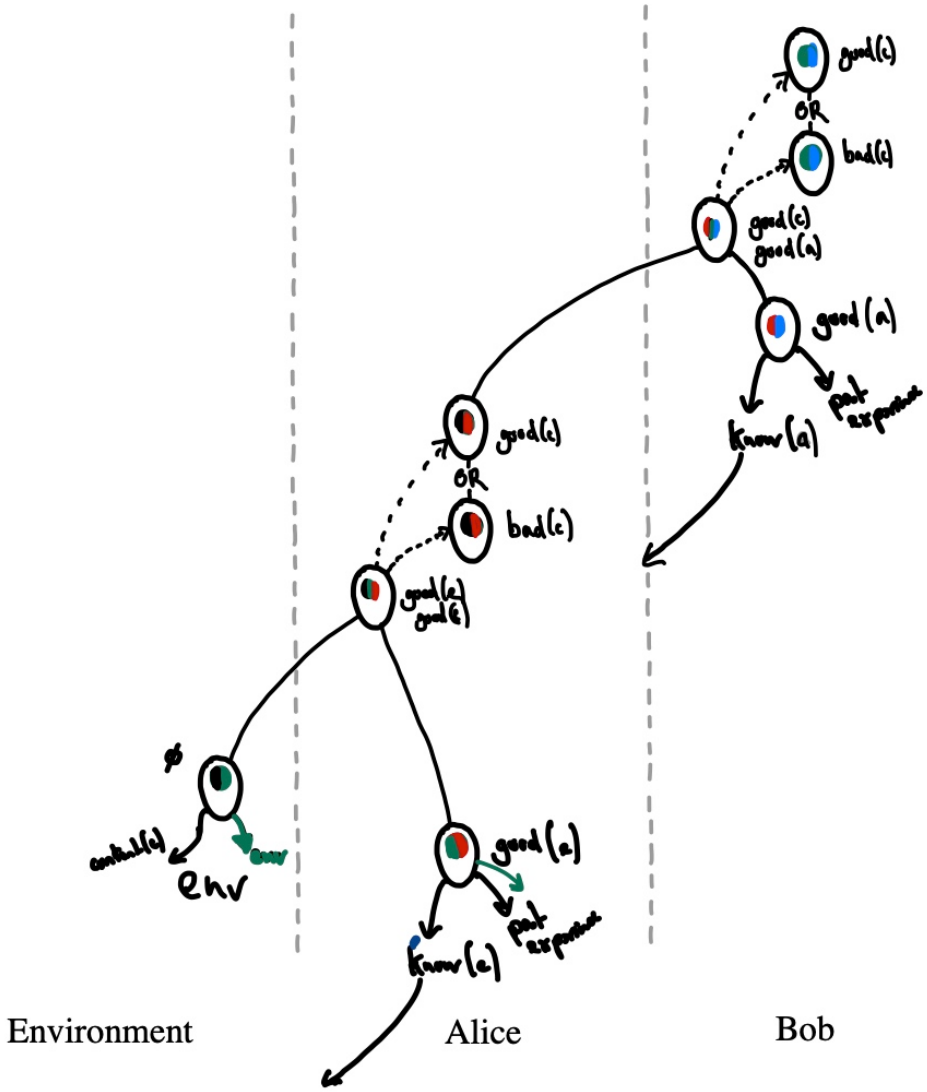


Figure 3.5: Sketch showing a method for representing a trace of occasions of experience. Alice, having come to know the environment, experiences a piece of content from it. If Alice feels this to be good then Bob, having come to know Alice, may experience this content. He first experiences it mixed up with his good feeling about Alice, before then re-experiencing it without Alice in mind.

These occasions arise in regions separated by the dashed lines—an environment to Alice, Alice and then Bob. Time is read upwards with lines incident to an occasion indicating its causes. The content of an occasion is indicated by the colours within and may be carried forward into new occasions. This coloured content might also be associated with statements such as $good(e)$. These might be viewed as logical formula or facts; the past occasions influencing the formation of a new occasion forming a knowledge base and the subjective experience formed in the occasion the result of selecting those of these that mix to form a rich experience.

3.4.2 Process philosophy

G. T. Kneebone, in his introductory survey on "Mathematical logic and the foundations of mathematics" identifies Whitehead's metaphysics as a well thought out attempt to describe how mathematics can be applied to the natural world [45, chapter 13]. He gives an introduction to *The concept of Nature* [97] in which Whitehead pursues a metaphysics compatible with the theories of quantum mechanics and relativity which were new at the time.

Whitehead starts by arguing that the ultimate fact we must start reasoning from is what he calls 'the passage of nature'. He describes the coming and going of acts of sense-awareness as evidence of this. Each such act he considers unique. He also considers them ultimately private. For example, we may not quite know what went into building our internal experience of a mug on our desk, but when we look away and that act ends we see that it was clearly unique, and that the subjective experience constituent to it is private, not only to us as an individual, but to that act itself. Whitehead labels these moments of sense awareness *occasions*. An occasion of new experience, Whitehead argues, comes to us at some level preformed like James' droplet of [38, Ch. X]. Only after we experience this may we analyse it to see what went into its formation. Whitehead describes the mechanisms behind the formation of these droplets, or occasions, of experience.

Each moment of sense awareness contains elements which are *discerned* directly from the senses and elements which are *discernible* being those derived from these. Other philosophers such as David Hume make the same distinction, but Whitehead considers the process of perception somewhat differently. Whitehead's Philosophy of Organism is a process rather than material metaphysics.

In developing his system Whitehead conceived no pre-existing 'screen' onto which these drops of sense-awareness or occasions are projected. Instead these are continually and actively pulled as much as pushed into a model of the present which is always based on prior experience. As such this model is of a short-lived slice through reality rather than an instant of zero length. During such a slice of time things will be changing; and so our perception of reality is necessarily of things happening.

Whitehead, calls this the *specious present*. The specious present then, being a duration through time, contains things going on rather than a snapshot of reality. Whitehead says that 'what we perceive as present is the vivid fringe of memory tinged with anticipation'. Further, he considers this limitation of perception also to apply to observation or measurement in general and thus to physical as well as perceptual reality.

Having started his metaphysics from this passage of acts of sense-awareness, Whitehead's early project was to build a system capable of supporting the practice of science. He proposes that the specious presents of contemporaneous observers overlap in a way that they form the neighbourhoods of some topological space; and that the resulting structure, described using his *method of extensive abstraction*, may be divided indefinitely to recover the zero-length snapshots of the state of nature at any *moment* that we are more used to dealing with. In his earlier work he labels each act of sense awareness, or a broader cut through reality as a whole, an *event*.

3.4.3 Method of extensive abstraction

The main purpose of this section is to build intuition about what Whitehead means by an *event*. It is important to remember that these events are not regions through space-time: they are things unto themselves, and later in Whitehead's work, they will become his 'atomic' building blocks of reality.

Whitehead's *method of extensive abstraction*, introduced above, is based on what he calls the *law of convergence to simplicity by diminution of extent*. The latter simply says that the relationship between parts of two events will be simpler than the relationship between the wholes of those two events. For example the complexity of the relationship between two people at one lunch is simpler than that between those two people throughout their whole lives; or taking it step further, the complexity of relationship between two people who never meet is simpler than two people that do. He builds his topology based on the relation aKb meaning that event a extends over event b, or that b is a proper part of a. The relation aKb has properties [96, p. 101]:

(i) $aKb \implies a \neq b$;

(ii) $\forall x(\exists y(xKy) \wedge \exists z(zKx))$;

(iii) $\forall x(bKx \implies aKx) \implies (a = b \vee aKb)$

(iv) $aKb \wedge bKc \implies aKc$;

(v) $aKc \implies \exists x(aKx \wedge xKc)$;

(vi) $\forall a \forall b \exists z(zKa \wedge zKy)$.

These can be interpreted as (where the set of events which an event e extends over is called the set of *parts* of e):

(i) Events don't extend over themselves;

(ii) Every event has parts and is a part of other events;

(iii) If the parts of b are also parts of a then aKb (or a=b);

(iv) The relation K is transitive;

(v) If aKc, there are events such as b where aKb and bKc;

(vi) Any pair of events are extended over by other events.

Whitehead defines an *abstractive set* as an infinite nest of events which would, by applying these rules, converge to a unique limit of a zero-length moment. For this to be the case Whitehead says that for any pair of events of the set, one must extend over the other, and that there must no event extended over by every other in the set. Two abstractive sets are *K-equal* if every member of one extends over some member of the other. K-equal abstractive sets are equivalent in that they converge to the same limit.

A geometry such as this one, in which there are regions but no points, is called a *point free geometry* and may be a partially ordered set, the study of which has moved on a long way from Whitehead's time. Gerla and Miranda describe how Whitehead later moved to a theory based on connection rather than extension [26]. They also demonstrate that a useful model can be built from the extension based theory of section 3.4.3 (or at least an inclusion based geometry they consider equivalent) by cutting a real Euclidean space into

regions consistent with a nest of events.

3.4.4 Philosophy of Organism

In his newer work Whitehead expanded greatly on the nature of these events or occasions. The only physical entities in this system of process metaphysics are these occasions. These are the atomic building blocks of Whitehead's view of both physical and perceptual reality. They form through a process Whitehead labels *concrescence*. They are the transient *creatures* of his *Philosophy of Organism*. An occasion is rooted in the branching causal networks of the upstream occasions which was its cause. In the case for mental rather than physical systems, pieces of a-priori knowledge and propositional statements will also have been folded in to these causal roots.

Where a procession of similar but evolving occasions form, for example, those formed while watching a leaf drop, or as a family regroups every year to recreate Christmas, Whitehead identifies a longer lived *nexus* of occasions. Once formed and experienced each occasion in this nexus dies. The leaf falling though still attracts attention and Christmas comes again. Such occasions as these then are destined to form the ancestors to the next generation of 'themselves'. There is a definite causal relationship in which such a subsequent occasion occurs after the former. Whitehead describes across eight categories the rules governing which occasions may form and how they do so. Key to understanding the scenario presented here is that two occasions will not form from exactly the same cause.

In this way, each occasion is what Whitehead calls a *superject*, both subject and object. While living, or forming, an occasion builds a subjective view of the objective data formed by its predecessors. Once it has formed then it becomes just another object to subsequent occasions. The predominantly held material view, that one system forms a subjective view of another, or that one system measures or observes another, is no longer primary. This is a big change from the way are used to understanding our reality.

Now the operation of a process of related occasions coming and going is actually more interesting than this. While an occasion is forming from what Whitehead calls its

actual world (the collection of dead occasions which are its potential cause) a number of neighbours will be forming contemporaneously (that is with no way of influencing each other but with overlapping actual worlds). These neighbours will be 'picking out' different combinations of occasions from the actual world for inclusion. Each such combination results in not one but a set of occasions, for while forming an occasion not only 'picks' which upstream occasions to include, but from the elements that went into forming these. Whitehead says that each such element is *prehended*. It is to these acts of prehension themselves which a prehension reaches back toward (possible through the *medium* of a series of occasions) bringing their data forward into the occasion of which it is a part.

Whitehead's approach results in structures of such processes of occasions, or more broadly a nestable partial hierarchy of *societies* of occasions. Whitehead identifies a collection of causally related occasions as a *duration*. Every duration contains other durations and every pair of durations is part a broader duration [45, chapter 13]. For example, we might pick out processes of causally related occasions associated with work and home life and section these into separate durations. Zooming 'outwards' we might find these contained by the duration of our entire lives and then that of life on Earth. Zooming back 'inwards' from here we find processes such as staying alive as a society, staying alive as an individual or maintaining good hygiene and relationships, which are all prerequisite to at least the duration containing work events. Thus we see that processes of occasions may be situated 'next to each other' and may depend on each other. This interdependency results in what Whitehead considers a society of occasions, or of overlapping societies of occasions. The resulting structure is suitable for modelling processes at the scale of the web while at the same time is key in Whitehead's understanding of individual and communal perception.

3.5 Conclusion

Driven by our initial assumptions in section 1.3 we have explored concepts that will support our modelling methods. We are led toward a model in which individuals are driven to form shared experiences and to share states of mind; and in which these experiences form

as we mix droplets of perceptions into existing states of mind. These behaviours may result in communities that can combine their abilities while working together.

The structure of communities on the web were found to be self-similar in some cases and we asked how such structure arises. Two simple mechanisms, for example, tend to cause parts to become more connected and for large communities with a certain purpose to dominate smaller ones with the same purpose. Balancing against these Axelrod describes a simple mechanism that prevents the drive to share our beliefs from resulting in complete homogeneity.

We described Petri nets in sufficient detail to understand their use in later chapters; and we introduced Whitehead' process metaphysics which was the launch pad for our approach to this work—however following this is not necessary to understand this thesis.

Chapter 4

Strands of experience

4.1 Ontology 63
4.2 Experience as content integration into states 67
4.3 Simple experience model(G^1) 68
4.4 Composite experience model (G^c) 69
4.5 Extended models 71
4.5.1 Prerequisite experience model 71
4.5.2 Trust or slow experience model 72
4.5.3 Labelled composite content and emergent strands 73
4.6 Discussion 74
4.7 Conclusion 74

Recall from section 1.6 that our goal at this stage is to develop a self-similar modelling part that we are calling a strand. We asserted the existence of such strands that, when interacting with other strands, act as if they support no internal interaction. We equated these strands with processes or with stories about a thing with constant identity. We begin our formalism by better defining what we mean by a strand of experience, and then by developing a model describing what these strands might experience.

We develop a model determining how strands' experiences build on each other; with past experiences filtering which objectively shareable experiences raise into awareness and how they are subjectively experienced. In this model experience occurs during the process of integrating shareable content into a strand's *state of mind* to result in a possibly new

state of mind. Some experiences are composite in that to form they require strands to work together. The model defines all possible paths of simple and composite experience. It describes which possible experience a strand may undergo at any moment—while dynamics discussed in the next few chapters will govern which occur and when.

We start building our ontology of parts in section 4.1. In section 4.2 we form a basis from which to develop generators that constrain how strands of experience unfold. Then we develop two models that are important to the argument in this thesis: in section 4.3 we describe experiences that strands may experience by themselves; and in section 4.4 composite experiences in which strands contribute different components. Two more models developed in section 4.5 are derivative approximations used in implementations tailored to particular scenarios presented later in the thesis.

4.1 Ontology

We start by stepping away from social machines. Over the bulk of the thesis we develop a theoretical model which supports systems with the nature of social machines captured by the assumptions in (B) and (C). Our assumptions about the nature of this theoretical model are captured in (A) and we build our model from these to start. These assumptions inform our choice of the categories of parts to model and in an iterative way, our subject domain.

Throughout the discussion, we will use the following definitions:

Interactant: A person or system capable of interacting socially, including those capable of forming and acting on shared intentions together.

Participant: When interacting in a way that contributes to a social machine's activities, people or automota are participants in that machine.

We will represent interactants and participants as collections of separate but overlapping strands of activity and talk at the modelling layer of interacting or participating strands rather than interactants or participants.

Our assumption (A1) is that the fullest levels of collaboration require shared intentionality. Recall from section 3.1.2 that shared intentionality requires: a way for people to shift their attention or perceptions to be about the same things; and a way to maintain coupled states of mind including those necessary to exhibit shared intentionality. We support the first by defining shareable *content*:

Shareable *content* are things that interacting strands of activity may notice or share with each other and that direct attention to something that may, in turn, affect *state of activity*. Interactants may direct each other's gaze onto parts of the physical environment or more broadly their attention to the more abstract things our states of mind may also be about. In the case of intentional agents, for example, these more abstract things include mutual goals or intentions to reach them.

To support the second, we define *states of activity* that can be held mutually by different participants:

***State* of activity.** An ongoing state of activity in which a mental representation is formed about something during cognition, some action is undertaken, or some emotional affect is held, or propositions made. A state of activity is live and carries with it a collection of often recurrent experiences that arise again and again. A state of activity may be shared by a collection of stable strands of activity in a way in which, we will argue in the upcoming chapter of forums, they act to share some states-of-activity between them in unison.

These are states where something is happening. This choice is in keeping with assumption (A2), chosen to embody (C1) that we should focus on the capacity for certain activities rather than sectioning out the parts of the world that are involved in their occurrence.

These definitions suggest a simple model in which the state changes when content is attended to and experiences formed—as captured in assumption (A3). There is support for just this kind of simple model. As introduced in section 3.4.2 James describes experiences as forming like droplets or buds of perception. In Whitehead's model, past occasions of

experiences inherit from each other and are mixed together to form new such experiences. We define an experience to fit this model:

Occasions of *Experience*. *Experiences* form as *content* is integrated into *states-of-activity* that result from prior such *experiences*. Experiences exist always as part of a *strand* or *strands* of experiences (see below).

The model developed in this chapter has the characteristic, like that of a state machine, that experiences that alter states are instantaneous and that the time spent in states is not relevant to the model. There is a mismatch when we try to apply this kind of modelling to strands of activities that vary over time. If the reaction of a strand's state-of-activities to integrating content, as just described, converges quickly to steady state then this could pragmatically be ignored during modelling. But we need to also capture the case where exposure to new content has an effect that takes longer to settle, or that does not settle—as strands carrying memories offer the ability to form composite experiences with access to their content, for example; or as a fact bumps up against propositions. It is not wrong in this system to say that content integration is quick, but we need to see it as the introduction of content into an ongoing state-of-activities, and of the new paths of experience that come with it.

We have not asked what an experience happens to and what this state is in. We don't want structure to thoroughly shape experience—that is, we don't want there to have be some object-like feature always existing prior to an experience forming—and in fact we will propose in chapter 10 that the birth of something akin to such a structure may result from an experience. In Whitehead's conception of these experiences, the actual occasion, there is no preexisting subject for the experience to happen to. Instead, a new subject forms with each experience, and is gone with that experience—the subject/object relationship, so familiar to us, remains but it is now between an unfolding experience in which the subject arises, and those, now completed, experiences that proceeded it and that it draws upon, that form the object.

So, we do not want to say that an experience happens to an interactant, participant, or a social machine, for example. Instead, we define experiences as arising as part of a *strand*

of unfolding activity that we choose to attribute to a thing that we label an interactant, participant, or social machine:

***Strand* of activity.** A relatively self-sustaining and self-causing sequence of experiences, that in our model carries state and is the thing that an experience occurs to. We will model scenarios using strands, or collections of strands, of activity instead of permanent objects or systems like people. The presence of objects or systems may be inferred from the presence of a strand of activities. We will see that a strand might be driven, for example, by some desired outcome, to maintain a model of some part of reality or some other collection of facts or beliefs, or just by the presence of some rising emotion in a person or group.

Content that will not fit may be considered by the strand, but, unless integrated, the strand will appear dormant in some cases, or remain in the same state-of-ongoing activity in others. For example, when exposed to a new fact while integrating content as part of a process of sense making [43][44] a strand of activity that carries our understanding about some part of reality will 'wake up' if that fact changes the model carried as part of its state.

Online communities host strands of activity that depend on defined interaction frameworks supported by services. We identify a framework-supported strand:

Framework-supported strand. A strand of activity facilitated by a social machine framework. For example, each subreddit, Twitter handle, and Facebook group supports a separate framework-supported strand of activity.

We propose that interactants or, more precisely, strands of activity associated with interactants may act together to form two classes of experience:

Composite experience. An experience to which multiple strands contribute their own experiences in a complementary way—for example, to ask and answer a question. As a more encompassing example, an experience may form in reaction to certain external content and this may have a cognitive element as well as an affective response and possibly an action. These three aspects of the resulting experience may be

embodied in reality in very different ways, and it makes sense to consider that each is contributed by a different strand. For example, our current emotional state may affect our cognitive response to given content; and this may sustain through many such otherwise unrelated experiences.

Aggregate experience. An experience in which multiple strands work together to form an experience that each might form alone but with the increased intensity that comes of sharing this with each other. In the limit where each strand's experience is identical, we might call this a common experience.

An example of a composite experience includes complementary actions toward a common goal, such as holding a tower of blocks steady while someone adds a block, or taking complementary roles during a conversation. An example of an aggregate experience would be to laugh together when the tower falls. In the context of social interaction facilitated by network technology, a substantial portion of online participants may follow the same conversation, and react to it to form aggregate experiences within groups that have formed around common past experiences.

4.2 Experience as content integration into states

People and communities carry systems of knowledge, belief systems and cognitive models as well as affective emotional states and shorter lived states of mind—or more formally, states of intentionality. We must allow strands to form different experiences in response to the same content based on their past histories of experience. And we assume that we can capture these past histories in the form of these states; and that the content is objectively identifiable but may be subjectively experienced differently. The process of integrating content may include the exploration of possible effects on our previous belief model. This may be a complex task with many leads tried and dropped; but we assume it occurs as an atomic event.

At this point we assume that stable beliefs and possible states of mind are eternal along with the possible experiences formed while moving between them. In reality the

state of holding a particular belief for example is more likely something that must be built and maintained. We are working toward a goal in which stable forums emerge within an ecosystem. While doing this we hope to see the emergence of stable belief systems composed of multiple smaller states—and we will show that we can in section 7.2.

4.3 Simple experience model(G^1)

We define a class of experience generators G describing the possible paths of experience of strands. In this section we start by looking at a simple generator supporting experiences that strands may form by themselves. The cognitive models or belief systems that a strand may host form vertices of a directed graph with edges between these vertices representing possible experiences. This section has been peer reviewed and published in [94].

A first order experience generator, G^1 describes the possible paths of experience of a single strand. As we'll see in chapter 5 such strands may still traverse these paths together to form communal experiences; however strands represented by these generators may also form these experiences by themselves.

We can implement a first order experience generator with a *labelled transition system* to take the form of a graph with labelled edges:

$$G^1_{\text{ts}} = (X, \Lambda, E) \tag{4.1}$$

where Λ is an alphabet of shareable experiences, used to label edges between states and X is a set of internal states resulting from integrating these shared experiences with previous ones—these states capture some portion of the shared experiences so far and affect the way future shared experiences are integrated. Finally, E is a set of transitions taken from $X \times \Lambda \times X$ where each element (p, α, q), or $p \xrightarrow{\alpha} q$, indicates that the shareable experience α, when experienced while in state p will be experienced in a way that moves the strand to state q.

Figure 4.1 shows an example experience graph. This has two separate components and can be described by a total of six (p, α, q) tuples. A strand is constrained to experiences

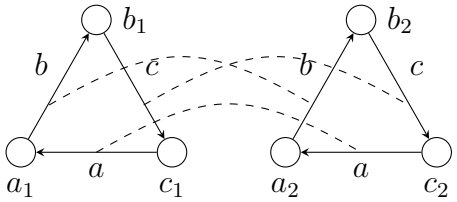


Figure 4.1: An experience graph showing two components coupled through overlapping content. The components $a_1 - b_1 - c_1$ and $a_2 - b_2 - c_2$ are coupled by a common alphabet of shareable content a, b, c. The content might form a short cyclic story. The left component might describe the experience of this content by someone with left-leaning political predispositions and the right component someone with right-leaning predispositions. The content might for example be the cycle: a="Shooting in US school" followed by b="canned response from our political side" and then c="anger at other political side"

on one or the other component. However, given that the shareable content $\{a, b, c\}$ is the same for both it is still possible that strands on different components may affect each other if the dynamics that we will add later support this. We try out this example generator in chapter 5.

It may be useful in the future to specify these generators as state machines rather than as labelled transition systems. To do this we would add x_0, an initial state, and F, a (possibly empty) set of final states indicating that the strand has nothing left to experience. This approach might be good for considering the life cycle of a strand or if using a generator to describe tasks that may be completed. Another alternative would be an action language [25] used to describe scenarios where actions effect an external world which in turn effects future options.

4.4 Composite experience model (G^c)

In chapter 5 we will introduce the idea that a collection of strands may form a *forum* of strands to which they cede their excitement about the content they have integrated and about the state of mind this has brought them to. For now we identify simply the more general concept of a *collective*. A collective might represent a single person comprised of multiple strands; or it might be a community comprised of many people with each person comprised of multiple strands. Strands within these collectives can work together to form

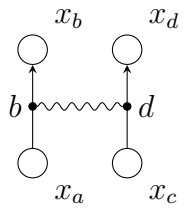


Figure 4.2: A simple example of a second order experience generator. The solid line between the edges b and d indicates the presence of the pair of edges $(x_a \xrightarrow{b} x_b, x_c \xrightarrow{d} x_d)$ in E^2.

a *composite experience*. While this occurs those strands might be considered to bind to form a single emergent strand.

To represent composite experience within a collective we define the class of c-order experience generators G^c as representing experiences that form during the coordinated movement of between 1 and c strands with each moving along their own parts of G^1. There is a choice of how to proceed when $c > 2$. We start with G^2 and return to this choice later.

We can implement a second order experience generator

$$G^2 = (G^1_{\text{ts}}, E^2) \tag{4.2}$$

where $G^1_{\text{ts}} = (X, \Lambda, E)$ as defined in equation (4.1), and E^2 are pairs of edges taken from E. The presence of a pair $\{e_a, e_b\}$ in E^2 adds a constraint on the first order generator indicating that the edges in G^1_{ts} can only be experienced together by strands in a situation where they may interact together to form a composite experience. For example, figure 4.2 shows an example in which a strand in state x_a and a strand in x_c are constrained such that they must experience b and d respectively together. We try out this example in chapter 5.

Because the edges in E^2 are unlabelled we cannot associate this composite experience with the integration of a single piece of content. Instead the experience is built 'upwards': where the content b and d is available to two strands they can move together. For example b and d may represent a question and its answer; or two facts or propositions, that come together to produce a compound proposition; or two experiences that can be merged

together in a more creative way.

We can of course imagine higher order generators and we will need them nearer the end of this work. A simple approach is to define

$$G^c = (G^1, E^2, ..., E^c)$$

where E^n is a tuple of n edges taken from E defined in G^1. This is a simple composition: each order adds another option to the movement of strands represented in the first order by G^1_{ts}. We continue this discussion in section 4.5.3 but this is not needed to follow the argument in this thesis.

4.5 Extended models

We develop two more models here. They are derivative approximations used in implementations tailored to particular scenarios presented later in the thesis.

4.5.1 Prerequisite experience model

There are many situations in which a collective of strands with the appropriate past experiences may open themselves up to new experience. For example figure 4.3(a) shows an experience that can be formed by strands in the state x_3 only with the help of strands in states x_1 and x_2. This arises so frequently that having a special shorthand is useful. Figure 4.3 (b) shows this alternative representation. The diagram indicates that the traversal of edge e_3 requires that strands be in states x_1 and x_2. This approach more directly captures some situations and results in smaller models. Expanding on the nomenclature of Petri nets we say that the edge e_3 is socially guarded, in that strands traversing it need the help of others. There is also a functional difference. The strands in x_1 and x_2 don't traverse an edge. In the models of dynamics of chapter 5 this means that there will be no recovery period where strands are stuck in these experiences, and no sharing of the content that labels e_1 and e_2.

There are a few ways socially guarded experience may be implemented. Take an edge



(a) **(b)**

Figure 4.3: A third order experience graph and an approximate form made using social guards. The edge e_3 requires two precursor experiences e_1 and e_2. Represented in (a) as a composite experience, and in (b) as a socially guarded experience

$p \xrightarrow{\alpha} q$. We might guard the edge itself, the content α, the source state p or the target state q. The most explicit solution is to guard each edge as indicated above. We can craft a socially guarded generator

$$G^{\text{guarded}} = (G^1_{\text{ts}}, G) \tag{4.3}$$

where $G^1_{\text{ts}} = (X, \Lambda, E)$ is the first order generator of equation (4.1), and G is a set of edges (e, g) indicating that edge e is guarded by state g.

4.5.2 Trust or slow experience model

It will be useful to represent some experiences that require longer exposure to content or more intense exposure to content to form; or that requires the trust of a source to follow. For example to change ones mind about certain topics of interest during communal decision making, sense making or culture formation. We can expand the generator described above to include experiences that require trust. We represent these on a generator graph as double edges. These indicate to any models of dynamics using the generator that an edge may need special handling. See figure 7.2 for an example. To this end we expand equation (4.1) to complement the set of edges E with a second set that require trust to traverse E_{trust}

$$G^1_{\text{trust}} = (X, \Lambda, E, E_{trust}). \tag{4.4}$$

An edge must be in only one of these two sets.

4.5.3 Labelled composite content and emergent strands

Although not necessary for the rest of this thesis we briefly consider the effect of explicitly labelling composite experiences with content. Recall that in section 4.4 we defined a composite experience in a way that it acted to constrain simpler experiences. The constrained experiences can only arise when each contribute to form the composite experience. If this form of composite experience occurs it is because it has been driven by these contributions. Now, if the composite experience itself is labelled by content then the strands contributing to it may be drawn into contributing not of their own accord but because of a drive to form the entire composite experience.

In a sense this turns a collection of strands forming a composite experience into an emergent strand. For example, suppose that a collective has recently experienced the second order content 'big dog'. The collective is excited to recreate this composite experience in a way that it could not have in the model of section 4.4 as it would have had no label with which to be identified. To form 'big dog' one strand in the collective must contribute the experience of the concept 'big' and another 'dog'. This would cause the content 'big' and 'dog' to be raised into the awareness of each strand and possibly then into other members of the collective; while at the same time the emergent strand might re-share the content 'big dog'. We can start to see that the collective having focus on the term 'dog', for example, might then through some form of association raise other content up into awareness; although at this point the dynamics of our models in chapter 5 would not support this.

Content that drives a composite experience in this way can be viewed as *composite content* in that the experience might be driven in a particular collective either from 'above' by this single piece of content or from 'below' by the content that drives each individual contributor; and that when experienced from 'above' releases its component parts that drive each individual contributor. It is then natural to open up the generator of section 4.4 to allow a composite experience to comprise other composite experiences rather than just first order experiences. At some point labelling the content becomes a convenience for our modelling approach; content might be viewed as packets of experience that would as a

first step be integrated into every collective of some class in the same way, but that would then result in a secondary more personal subjective experience that depends on the state of the collective.

4.6 Discussion

Obviously these fully predefined generative models will not scale well. At the conceptual level we may replace these predefined generative models with a process in which options continuously unfold. The system model here would be unchanged by doing this. Whether or not we believe in the existence of such a priori structures is a matter ultimately for philosophers, but this model forms a solid basis that we find in the next chapters to be a good place to reason from.

4.7 Conclusion

The ontology started here forms the basis for the modelling in the thesis, and as a way to connect the resulting models to real-world phenomena. The models presented here describe possible paths of experience that strands may undergo. These models determine how strands' experiences build on each other; with past experiences filtering which objectively shareable experiences raise into awareness and how this is subjectively experienced. Experience occurs during the process of integrating shareable content into a strands' states of mind to result in a possibly new state of mind. Some experiences are composite in that to form they require strands to work together. The models describe just what possible experience a strand may undergo at any moment—while the dynamics discussed in upcoming chapters will determine which actually arise and when.

These models are important as we will use them first to define the behaviours of pre-existing strands and then later as an approximation for the behaviour of emergent groups of such strands that move together coherently. We will gain more confidence in these models as we use them throughout the thesis.

Chapter 5

Collective experience in forums of strands and the emergence of strands

5.1 Ontology 76
5.1.1 Overlapping forums of strands 77
5.1.2 Dynamics of strand interaction in forums 79
5.1.3 Dynamics of aggregate strand formation 83
5.2 Examples 84
5.2.1 Shared intentionality 84
5.2.2 Discovery of peas galaxies 85
5.2.3 Social machine frameworks 88
5.3 Forum model 91
5.3.1 Simple experience 92
5.3.2 Composite experience 94
5.4 Experiments and results 95
5.4.1 Flock formation 96
5.4.2 Flock life-cycle and synchronisation 96
5.5 Discussion of results 99
5.5.1 Flock formation 99
5.5.2 Flock life-cycle and synchronisation 101
5.6 Conclusion 103

Having stated that there are parts of our ecosystem that act as strands governed by models describing paths of possible experience, we now need to model a collection of

these strands. This collection of pre-existing strands needs to support the movement of short-lived virtual strands associated with composite experiences described in section 4.4; and it must support the emergence of single coherent threads of shared experience—that might act as strands in other collections. Let's call our model of this collective a *forum*. A forum might represent a single person comprised of multiple *strands*; or it might represent a community comprised of strands cut from a collection of people.

Following from the discussion in the last chapter, first let's add a new requirement for our theory—and let this drive our model making:

Social machine characteristic: A social machine depends on and supports its participants' nature to form shared experience.

We further develop our ontology in section 5.1, before then showing that this can be used to describe a number of offline and online social situation in section 5.2. We develop two models in section 5.3. The first model, developed in section 5.3.1 supports only the individual experiences captured by the model of section 4.3. The second, developed in section 5.3.2 builds on this to support the composite experiences of section 4.4. We verify that these models support these required behaviours through experiment in section 5.4. The implementations are coded using the stochastic coloured Petri nets introduced in section 3.3.1. This chapter is expanded from [94].

5.1 Ontology

In chapter 4 we proposed that strands of activity integrate content to form sequences of experiences, and that these strands carry the sustained effect of such past experiences in the form of a state of activity. We now need a way to allow these strands to interact. Particularly, given assumption (A4), that a collection of strands of activity may act coherently within another collection as if it were a single strand, we seek behaviours that support this emergence.

We propose two different kinds of emergent strand. Take two interactants who, while reacting to a social media post, find themselves in a conversation. If we attribute a strand

of activity to each party's private view of these interactions, then these may follow a path of experiences that they could not access alone. This is not fully captured simply by each strand unfurling alternatively while one speaks and the other listens. It is reasonable to say that the pair's activity constitutes a strand in its own right—at least from the perspective of other strands that might interact with it. This *composite strand* may come into being as its components begin to interact in this way and fall out of being when these stop interacting in this way. At a possibly much larger scale where interactants may follow the same path of experience together, they form an *aggregate strand* in which they act from the outside like an emergent strand.

The first kind of emergent strand comprises a sequence of *composite experiences* as defined in the last chapter, and the second a sequence of *aggregate experiences*. The ultimate expression of assumption (A4) forms a defining characteristic of the theory: emergent strands of either type may themselves interact. So, for example, we may see the core groups of communities with opposing political views each speak with one voice and discuss issues in a way that two individual people might.

5.1.1 Overlapping forums of strands

Taking a step back, we need a way to model scenarios in which there is some structure relating which strands may interact directly to form experiences together from those that cannot. We envisage a continuum formed by overlapping strands similar to Whitehead's extended continuum. Figure 3.5 shows strands of activity associated with participants in which the participants must have previously met each other to interact. This interaction is recorded in each strand's past interactions, which we already capture in the form of their states. However, it will be useful to have an approach tailored to representing strands' potential for interaction.

One approach would be to situate the strands on a network. However, we notice that there are pools of activity where strands interact. For example, our home and work lives may result in separate collections of strands of activity, with one collection tending to overlap more with those of our family and friends and the other our work colleagues. Some

of the strands within these separate collections will interact within us—for example those related to personal logistics, scheduling activities, and managing energy and head space—and some will interact also outside us—for example, where we work with friends. We will say that where strands can interact, that they *overlap*. To support shared intentionality, recall that interactants must be mutually responsive to each other. Where any number of strands are able to mutually interact, rather than capturing a network of all the possible strand interactions, we will instead identify the fact of this overlap by saying that the strands are tied or bound to a *forum*:

***Forum* of overlapping strands.** We infer a forum where a collection of strands' past experiences allows them to form experiences together. This interaction may occur during a composite experience in which multiple strands work together to form an experience with each bringing something separate—or strands may form the same experience together en masse. Just as the state of a strand captures some measure of that strand's past experiences, the membership of a strand in a forum captures some measure of that strand's past interactions with other strands.

To support what we have called self-similarity, it is important that in some forums a core collection of strands may act together in a way that other strands see that core collection acting as one strand. This one strand may comprise either many very similar strands acting in unison, or multiple specialist strands working together to form what appears to be a single strand of activity carrying coherent state and with which other strands may interact coherently.

This approach more accurately captures scenarios where the fact of this mutual overlap is more relevant than the myriad of individual overlaps that make it up. It is also pragmatic in that the number of ties between pairs of strands increases exponentially with the number of mutually overlapping strands. When we do need to keep track of these individual connections, a forum can be created for each pair of connected strands. In fact, we'll show in chapter 10 that forums need to support a dual nature, as both a fully connected clique in which interaction between strands is anonymous; and potentially any smaller network of strands therein.

We will see in chapter 8 that a forum may be represented as a hyperedge connecting vertices representing relatively stable strands of activity to form a hypergraph; where a hyperedge is an extension to an edge that connects more than two vertices. More than one hypergraph will be needed for many scenarios; as what to strands in one forum may be a relatively stable strand of activity, may be to another far too short-lived, fast-changing, or slowly changing to interact with.

Strands in a scenario may have access to the same underlying machinery of interaction that supports the exchange of certain states such as emotion, agreement, or disagreement. On top of this, the mutual interaction of strands in a forum may depend on common abilities, such as the ability to form the recurrent experiences associated with a certain system of beliefs or of speaking a common language; including experiences of using common words or idioms, or those that inform grammar or syntax. Where these abilities are common to all forums in a scenario, and remain approximately fixed during its duration, we may attribute them to the underlying medium capable of forming strands—or to classes of strands such as those supported by humans, AI chat systems, or the environment. Some abilities, such as those afforded by shared cultural traits, may vary across a scenario and, in these cases, may be attributed to individual forums; these can be modelled with the *social guards* of equation (4.4).

Finally, like with the framework-supported strand above we identify also a framework-supported forum:

A *framework supported forum.* A forum supported by services running a social machine software framework.

5.1.2 Dynamics of strand interaction in forums

We have said that a forum of strands may form experiences together. We now ask what behaviours are necessary for this interaction; a discussion that builds on the nature of strands described in chapter 4 and driven by assumption (B1), that a theory of social machines must capture collaborative working and, (A1), that to do so fully the underlying theory must allow shared intentionality to emerge. The model of the last chapter captures

causality in the way that states of activity may change from one to another, but does not describe how long it takes for this to happen or how competing options are selected. We answer these questions while adding characteristics necessary to support shared intentionality: response to content from the environment, joint attention and coupled states of activity.

Responding to the environment. In the following discussion it will be useful to consider the strands in a forum responding to content made available to them by their environment. Following from assumption (A5), whether a strand of activity forming an interactant interacts with strands of another interactant, or with strands from a non-social part of the environment, the strand is just interacting with other strands. We can model scenarios in which we don't have a full model of the environment as long as we stick to the same methods of interaction between both types of strands.

Take a single strand forming part of an interactant; for example, one that builds and maintains a model of activities in part of its environment. To gain knowledge about that environment in our model, a strand must integrate content shared from strands in the environment; and for the strands to interact they must form part of the same forum. We'll see though, that we cannot consider an environmental strand as fully part of the forum unless there is a chance for mutual interaction.

Suppose now that some content in the environment is relatively static for a period. For example, a cat and a mouse may be running in the environment. With knowledge accessible to a strand that "cat's chase mice", the strand might form the experience "this cat is chasing this mouse"—perhaps there is another, normally dormant, strand, further from the senses, that recognises the first strand's experiences and forms the experience alongside it. The person hosting this strand might report first noticing the cat, and then the mouse, and then the proposition that the mouse is chasing the cat. In the following discussion it is natural to consider that the strands interacting in a forum may each be associated with different people, but they may also form part of one person.

Strands may have joint attention. To support communal gaze or attention so that the attention of strands in a forum tends to follow the same content, we need the effect that while integrating particular content, a strand also shares that content with other strands in any forums it is part of. To achieve this effect we add two characteristics:

Strands share content into forums that they overlap with. This may occur when a strand forms an experience. We use the term share even where this sharing is clearly not deliberate.

A forum makes active content available equally to all its strands. Where content is observable in a forum, it is excited or active in the forum. Contributions from multiple strands increase the value of excitement for the content. Strands in the forum are likely to integrate content that they can integrate if it is more highly excited.

The term excited is used in the sense that a physicist might say that a system is in an excited state: something that could be happening but is not always now is. In cases where the content tells a news storey, interactants may be excited by it in the more literal sense.

If all possible content is laid out in a long array then elements of active content will be lit up according to their excitement or activation—that is, there is an array of weights associated with this array of content. Alternatively, using the geometric analogy favoured by machine learning practitioners, each element of content adds a dimension to Cartesian space, and the weights define a vector within this space.

With the dynamics chosen later in this chapter, the excitement of each element of this vector is proportional to the rate at which it has been recently integrated by strands.

As an example of the first behaviour, while integrating particular content from the environment, such as someone taking a new block from their pocket, a strand will add to the activation of that content as seen by other strands in the forum; that is, the person whose strand notices the new content draws attention to it.

As an example impact of the second behaviour, if the environment offers two pieces of content and many strands are in the position to integrate these, this tendency causes a positive feedback loop and a tendency for the strands to focus on one item of content over the other—a communal decision will have appeared to have been made to focus on the chosen content.

Strands may have coupled or joint states of activity. We now need a mechanism for interactants to form and maintain joint states of activity and for aggregate collections of strands to follow proceedings together. Strands may respond to the same content differently—or not at all—depending on their current state of activity. Where strands need to maintain a joint state, such as the shared goal, coordinated action plan, and common commitment to these, that form an intention, there needs to be some way for strands to maintain and evolve this coupled state. Equivalently, we need collections of strands with compatible past experiences to potentially flock and act as a single strand, carrying for example a changing cultural belief system. Another example of the sharing of states is of course shared emotion: a parent might initially take their child's feeling of dread and worry about getting behind while off school and share this together, even if the parent knows that the actual concern is not true. A single mechanism meets these three cases.

Just as we allowed a strand to increase a forum's excitement about content while integrating that content, we allow a strand's excitement about a newly entered state to increase the whole forum of strands' excitement about that state. However, this differs from the first case in that this excitement will only affect strands that are also in that state. We add the characteristic that:

Strands in a forum stick together in a new state after entry. A newly entered state of activity initially includes the act of integrating and reacting to the content that brought about the state. This is indicated by allowing the strand's state to be temporarily excited or activated on entry. In a forum this excitement is shared and is indistinguishable from that of the other strands. While the state the strand is in is excited in the forum, the strand will tend not to integrate more content.

Equivalently we may view this dynamic as shorthand for a situation in which: upon integrating the same content in the same forum, strands find themselves in a new implicit forum associated with that state and mutually excited about implicit content associated with that state itself. In the implementation of this theoretical model given later in the chapter, we stick with a model in which the state is excited. We are forced however, in chapter 10, to add the creation of this new implicit forum while considering the life-cycle of forums (and social machines).

Now if a collection of strands in the same state start integrating the same content they will end up in the same state. In the presence of content that may then be followed, this new dynamic holds the collection together in the same state until the rate of new arrivals reduces enough that this new content may be followed. During a physical interaction, a mutual interest in this shared state might be seen by nodding or other forms of reaction that grows together, such as surprise or annoyance. In the forum, a new pool of strands that have found themselves in the same state has formed. Once the forums' focus on its experience of the integration and reaction to the content wears off, and if nothing comes along that causes them to move together again, the strands settle into a steady state of activity—in many cases this is passive and amounts to sustaining knowledge of past experiences that is sufficient to influence future choices.

5.1.3 Dynamics of aggregate strand formation

We will show that these dynamics support assumption item (A4): that a collection of strands of activity may act coherently within another collection as if it were a single strand: a serially ordered sequence of experiences that from the 'outside' appear to carry consistent internal *state* and where there is a choice of content to integrate will appear to agree on their choice. 'Internally' many strands may work together to produce this external consistency. For example, a pair of people walking may be aware of internal deliberation while choosing a path, but from the perspective of a person walking toward them they may seem to act in unison; or, a team may appear from the outside to be of one mind even if internally there is ongoing internal deliberation. A participant in a social

machine for example may appear to host a strand that carries a goal, knowledge, skills, a model of current reality and intention and where an intention forms this may be seen to change coherently—at this level, if queried, the participant would speak with one voice about these things, even if internally they cannot hold all these things together in mind concurrently.

This chapter includes experimental simulations that demonstrate this formation of aggregate strands. Further, section 6.2.2 demonstrates that we can choose dynamics that govern the strands in a network of overlapping forums that, in the right circumstances, are identical to those of a single forum.

5.2 Examples

The model we have developed, of strands interacting where they overlap to form potentially overlapping forums, is a more abstract view than that we usually take of communities. Here, we give some examples to show how it can be applied, and in turn highlight examples of real-world phenomena matching those of our developing ontology. The examples will include some level of shared intentionality, so we start with a generic example of this, before moving onto web-dependent communities that can be considered social machines.

5.2.1 Shared intentionality

Described in section 3.1.2, Tomasello et al. identify three levels of social engagement that infants progress through. These provide good examples to test our models against.

At the first level, *dyadic engagement*, interactants take turn behaving and may respond to each other's expression of emotions—but don't yet form shared perceptions, or coupled states with any feedback holding the states together or mutual knowledge of this. We could model this kind of example with two separate forums, with each containing strands internal to just one participant as well as strands sharing some outward representation of the other interactants' emotional state and state of outward activity. These forums overlap through sharing strands that connect to each others' outward appearance—albeit

from different sides—as well as parts of the environment these outward activities involve..

At the next level, *triadic engagement*, the infant and an adult may monitor each others' perceptions (for example, by attending to the same parts of external reality) as goals are pursued. To model this example, we could take the above model and include the outward representation of each other's gaze in addition to emotional state.

Finally, *collaborative engagement* occurs when the participants share joint intentions (internal goal, decided plans of action and commitment) rather than just goals; and joint attention in the form of mutual knowledge rather than just monitoring each other's perceptions, and are aware of this. Action plans at this level may be complementary with, for example, one interactant holding a tower steady while another adds a block. We could model this kind of example with a single forum. Strands of activity sharing a common state-of-activity within the same forum are able to do just this. For example they might hold a state of activity "we are building this tower together", or "we are engaged in this online community built around finding, sharing, and reacting to particular types of content".

5.2.2 Discovery of peas galaxies

Citizen scientists, classifying images on the Galaxy Zoo platform, discovered a class of galaxies unknown to astronomers at the time. Tarte et al, in their paper on archetypal narratives in social machines [82], trace the life of the collection of images that these participants created and that culminated in the recognition of the "pea galaxy". Working from this paper, we start with a model of the social machine, at what the authors identify as *Moment (0)*, before the "green peas" collection is created, and for a few moments afterwards. Table 5.1 identifies the strands inside and outside the social machine. We situate each strand of activity according to the crowdsourcing dimensions described by Malone et al.[49]. Strands beginning with a P are participants in a social machine, those with a D are discussions, and those with a C are collections. A collection of images is considered the state carried by a strand that maintains that collection—and may be implicit within a discussion. Notice that what is identified as input and output, depends

on which activities are chosen to be part of the social machine. For example, from the perspective of the space science community, output may be an addition or change to their corpus of knowledge—which does not occur until *Moment (6)*. Alternatively, we might have narrowed the activities to just those that classify green peas galaxies, loosing the context of this social machine within the broader ecosystem.

Following *Moment (0)* the green peas collection is born:

1. **Classifiers share green-looking galaxies.** $P_{\text{classifying-weird}}$ regularly found green galaxies and shared them in D_{weird} (implicitly adding them to C_{weird}). Before *Moment (2)*.

2. **A classifier starts the green peas collection.** A classifier creates the new thread on the forum "Give peas a chance", $D_{\text{give-peas-a-chance}}$, and implicitly creates the collection $C_{\text{green-peas}}$. Browsers of this thread may form a new participating strand $P_{\text{classifying-green-peas}}$. *Moment (2a)*.

3. **A moderator gives weight to the green peas collection.** A $P_{\text{moderating}}$ strand is activated in Alice (who has a formal role as a moderator), who invites participants who shared images of green-looking galaxies to reshare them in $D_{\text{give-peas-a-chance}}$. The collection now starts to grow very fast as the P_{browsing} strands notice the activity, with some adding $P_{\text{classifying-green-peas}}$ strands to the social machine. *Moment (2b)*.

4. **Participants with $P_{\textbf{classifying-weird}}$ strands form a tight community** calling themselves "the Peas Corps". *Moment (2c)*.

Where these strands interact, we infer a forum. At *Moment (0)*:

- **Input forum, $S_{\textbf{input}}$.** The SDSS website supports interactions in which P_{source} strands hosted by scientists in the machine's environment collect images into $C_{\text{all-images}}$ hosted by the SDSS.

- **Participants' forums, $S_{\textbf{participant}}$.** Each participating strand P depends on the abilities and situation of the person that hosts it. We represent each person by

Strand	Description	Crowdsourcing dimensions
Inputting from ecosystem:		
P_{source}	Scientists, possibly unaware of Galaxy Zoo, manage images. Feeds into $C_{\text{all-images}}$ hosted as part of the SDSS website operations—and considered part of *the* social machine by the authors.	(Not part of the social machine boundary chosen by the Tarte et al.) What: create a centralised collection for their science community, Who: experts, Why: professional reputation and contribution to science, How: collection or sharing own images.
Discussing:		
D_{root}	The root thread of the actual forums (called "The Objects"). The D indicates a discussion strand that maintains an archive of the discussion.	
D_{weird}	A thread under D_{root} where $P_{\text{classifiers}}$ post and discuss 'beautiful' images.	
Participating:		
$P_{\text{classifying-all}}$	Classifies galaxies found in images. Performs "the on-going repetitive tasks of classification, punctuated by remarks and exchange on images, which can range from short aesthetic interjections to long exchanges that are rich in scientific content."	What: create classified collection of galaxies, Who: crowd, Why: enjoyment & perceived contribution to science, How: collection (platform seeks consensus behind the scenes).
$P_{\text{classifying-weird}}$	Looks out for stunning or otherwise noteworthy images and adds them to C_{weird} by sharing them into D_{weird}. Doing so creates a focused forum for discussion.	As for $P_{\text{classifying-all}}$ with Why: social interaction and reputation in local community.
P_{browsing}	Watches activities in the discussion forum and partakes. Makes contributors to, for example, $P_{\text{classifying-weird}}$, feel good and may cause participants to pick up new collaboration strands.	What: support classifiers & source of classifiers to machine
$P_{\text{moderating}}$	Moderates the discussion forums and organises classifiers.	What: organise and focus the machine, Who: experts with a view of machine operations Why: reputation in local community and possibly professional interest, How: direct communication to classifiers with an air of authority.
Maintaining collections:		
$C_{\text{all-images}}$	Database activities carrying state that includes all the images.	
C_{weird}	Implicit strand of activity discernible within D_{weird}.	
Outputting to ecosystem:		
$P_{\text{consuming}}$	Strands of activity that use the output of the social machine. The output in this case are found by following the *collecting* strands.	

Table 5.1: Active strands in the green peas scenario at *Moment (0)*.

means of a forum of strands. This allows a person to maintain more than one strand of experiences and for these to interact.

- **Classification session forums, $S_{\textbf{classifying}}$.** While a participant is classifying their $P_{\text{classifying-all}}$ strand interacts with $C_{\text{all-images}}$. A single forum containing $C_{\text{all-images}}$ would allow multiple participants to interact here, which does not happen in reality. These are recurrent forums supported by the social machine's framework.

  $P_{\text{classifying-all}}$ is part of a $S_{\text{participant}}$ forum. If the participant also hosts more focused classification strands, such as $P_{\text{classifying-weird}}$ then this strand will see the content experienced by $P_{\text{classifying-all}}$ and may activate.

- **Discussion forums, $S_{\textbf{root-discussion}}$, $S_{\textbf{stunning-discussion}}$, $S_{\textbf{wonderful-discussion}}$.** Stepping back we see just $S_{\text{root-discussion}}$. However, participants may be interested in only some parts of this: where their strands form part of a smaller group of people with narrowed interests, there are smaller forums.

  - **Stunning images discussion forum, $S_{\textbf{stunning-discussion}}$.** Participating strands in here include $P_{\text{classifying-weird}}$ and $P_{\text{moderating}}$. $P_{\text{classifying-weird}}$ has its 'home' here—Participants hosting one of these strands are driven in part by the community.

Strands, and thus the forums in which they interact, comprise certain recurrent experiences—some of these may form part of a grammar of possible experiences—where an activity is part of an experience. Forums overlap where they share strands, resulting in a structure. These recurrent experiences differ across this structure. A more complete analysis would include a study of these recurrent experiences; or of classes of recurrent experience that perform a similar function.

5.2.3 Social machine frameworks

Without repeating the relatively in-depth analysis of the green peas galaxies scenario, it is useful to identify some strands of activity and the forums these interact in, for a few social

machine frameworks. These are chosen to sample a range of the crowdsourcing dimensions described by Malone et al.[49].

Reddit's typical classification and reaction Subreddit. Reddit comprises a collection of communities, or subreddits, with each focused on particular topics. Many subreddits act to find and classify content such as pictures or news articles.

Each subreddit comprises:

1. **Participating strands.** Participants may integrate content in the form of posts and comments. They may also form experiences in which they vote for this content or click to go into a forum to react to a comment further. Strands can also share content with the community.

2. **A framework-maintained top-level discussion strand for participants' posts.** The experiences in this strand form together with that of participants when they create posts. The state carried by this strand is the trace of all posts.

3. **A framework-maintained forum hosting this strand.** Participants interact with the top-level discussion strand through a forum common to all participants.

   (a) This provides a pool of participants that when integrating this content all end up in the same state of reacting to it.

   (b) The framework draws attention to certain active content. The framework presents content in a way that some is more likely to be followed than others; in a way that is compatible with the theoretical model's content excitement or activation level. Although the algorithms used to judge this are propriety, when sorting posts in a Subreddit by "Hot" rank is determined "using votes and post age"[1].

   The framework provides a way to couple states of mind. For those wishing to compare their reaction to content with others, the framework provides a

[1]From Reddit's official blog: https://www.reddit.com/r/blog/comments/o5tjcn/evolving_the_best_sort_for_reddits_home_feed/

way to interact with a strand of activity stemming from this post in a forum dynamically created for this purpose (see next item).

4. **Recursive discussion strands and forums around parts of this.** The experiences of these strands form together with those of the participants when they create posts. The state carried by these strands is a trace of comments and other such strands.

Wikipedia. Most people who visit Wikipedia do not participate in creating the encyclopaedia's entries, and there is no forum in which these users interact with each other or otherwise affect each others' visits. However, the content is, of course, created by a community of volunteers who interact by creating and editing entries. There is a framework-supported strand of interactions that carries the text of each entry. Here, the content of each interaction, rather than a post or comment, is an 'edit' submitted by a participating strand. The forum in which the participating strands interact with this strand does not support the large-scale aggregate interactions of a subreddit. To those viewing the trace of this strand these edits are seen in time order making newer edits more visible than old—and very loosely we can say that newer posts or those that indicate many changed lines are more exciting to a participant. There is the capacity for one-on-one interaction in this forum both through private 'chat' messages to editors and by 'replying' to an edit in the history in a way that effects the change. In terms of crowdsourcing dimensions, Wikipedia differs from most subreddits in that although there may be a crowd of editors working across subject areas, most edits are made by experts in a field.

Now, if we take a big step back, Wikipedia could be viewed as a single social machine working on a single interconnected corpus of knowledge that ties into existing knowledge sources. At this level, there are many independent strands of activity that shape the knowledge base. The social machine framework includes rules that govern engagements driven by disagreement. These could be captured as constraints on paths of experience that form a grammar; and this might be enforced by either a technical component of the framework or by convention.

Pokemon Go. De Roure et al. write that "Pokemon Go conceals its purpose as an exercise app, intended to co-opt smartphones to drive exercise and social interaction, in opposition to their reputation as an isolating and alienating technology"[19]. This stands out from other examples, as experiences include physical actions that are constrained by participants' location and the opportunity for participants to interact face to face. This example is a game, and as such participants maintain a number of strands of experience driven by different purpose. For example, to find and catch Pokemon; to find and catch a particular Pokemon; to acquire a Pokemon by walking while holding a Pokemon egg; or to fight other participants in a gym.

Strands of activity can be identified in human participants. Other strands of activity maintain game rules and govern which strands may interact, how, where, when, and with whom. Some of this structure is imposed by game rules shaping framework-supported forums, and some by the possible experiences supported as participating strands form composite experiences with framework-supported strands.

In the first game mechanic given above, players hunt Pokemon by walking in the environment until they find one. A player hosts a strand with this purpose, and it appears to the player that a strand of activity governs the behaviour of each Pokemon out in the virtual environment. As with other scenarios, the presence of a forum indicates the fact that strands may interact—in this case this occurs when the history of each strand of activity brings them into 'physical' proximity. Perhaps more interesting is that certain kinds of fight require participants to be physically in the same place and thus able to interact outside many of the constraints imposed by the game's rules.

5.3 Forum model

A forum comprises a collection of indistinguishable strands. These form a fully-connected social network in that they all hear each other and are heard by each other equally. Each strand may be excited to be in any state of mind and separately excited by any shareable content that has been integrated. strands share excitement about content even when this moves them to incompatible states of mind. These levels of excitement are ceded to the

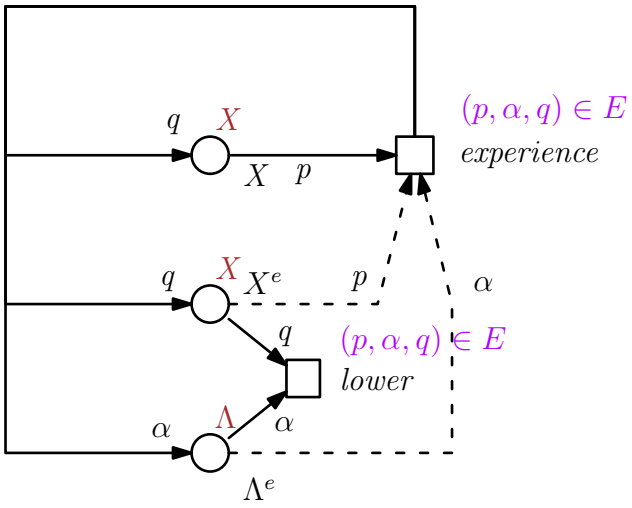


Figure 5.1: A coloured Petri net describing a fully-connected community, or forum, model. Places X and X^e contain tokens from the colour set X, and place Λ^e from the colour set Λ.

whole such that strands cannot tell their own excitement apart from their neighbours'. We will encode the generators from chapter 4 into Petri nets and add these dynamics.

The models of [93] used the colour of the tokens to differentiate strands in a network. Here we model a network that is so simple that the strands are indistinguishable. Token colour will instead be used to represent the state of experience taken from a graph of possible experiences. We will see that this makes it simpler to encode complex dynamics that govern movement through the graph.

5.3.1 Simple experience

A single coloured Petri-net shown in figure 5.1 was developed to model this scenario. Each strand is represented by a token in place X. This is drawn from a colour set X (comprising $\{a_1, b_1, c_1, a_2, b_2$ and $c_2\}$ in the example of figure 4.1) with the colour indicating that strand's initial state of experience. The firing of the transition *experience* will take an strand's token from the place X and replace it with one coloured to represent a new state. The guard equations $(p, \alpha, q) \in E$ ensures that any such transition will move the strand from state p along an edge labelled α to state p on the experience graph: it encodes the experience graph G in the form of a list of edges. This firing will also add a token to X^e indicating that an strand is excited to be in state q, and add a token to Λ^e indicating that

an strand is excited to have experienced the content α. The second transition *lower* will act to lower these excitement levels again.

Not shown on figure 5.1 are the rate equations for the two transitions. The rate equation for *lower* was:

$$\lambda^{\text{lower}} = X^{\text{e}}[q] \cdot \delta \tag{5.1}$$

where $X^{\text{e}}[q]$ is the number of tokens of colour q in X^{e} and δ is rate constant indicating how quickly excitement is lost. The guard function on this transition ensures that each strand loses excitement about a state they have entered and the content that led them as a single atomic event. This is tidy, but this implementation assumes that every state is reached by only one piece of content and every piece of content moves an strand to only one state. This approximation follows through the remaining models in the thesis; we will be careful to ensure that the experience graphs we wish to execute meet this constraint.

Of more interest is the rate equation for *experience*. The dotted lines, indicating modifier arcs, allow the rate to vary as a function of X^{e} and Λ^{e} without effecting the tokens in these places. The rate equation was:

$$\lambda^{\text{experience}}(p, \alpha) = X[p] \cdot \left(\lambda^{\text{self}} + \lambda^{\text{pull}}(\alpha) - \lambda^{\text{stick}}(p)\right) \tag{5.2}$$

where

$$\lambda^{\text{self}} = w^{\text{self}} \tag{5.3}$$

$$\lambda^{\text{pull}}(\alpha) = \frac{1}{N} w^{\text{pull}} \cdot \Lambda^{\text{e}}[\alpha] \tag{5.4}$$

$$\lambda^{\text{stick}}(p) = \frac{1}{N} w^{\text{stick}} \cdot X^{\text{e}}[p]. \tag{5.5}$$

Here $\lambda^{\text{experience}}(p, \alpha)$ is the total rate of all strands leaving state p along an edge labelled α. The λ terms that contribute to this are for each single strand. N is the number of strands in the community and w^{self}, w^{pull} and w^{stick} are weights capturing three forces. The first term w^{self} determines the rate at which any strand will spontaneously move along an edge. The term $w^{\text{pull}} \cdot A^{\text{e}}[\alpha]$ creates a pull along edges that strands have recently experienced.

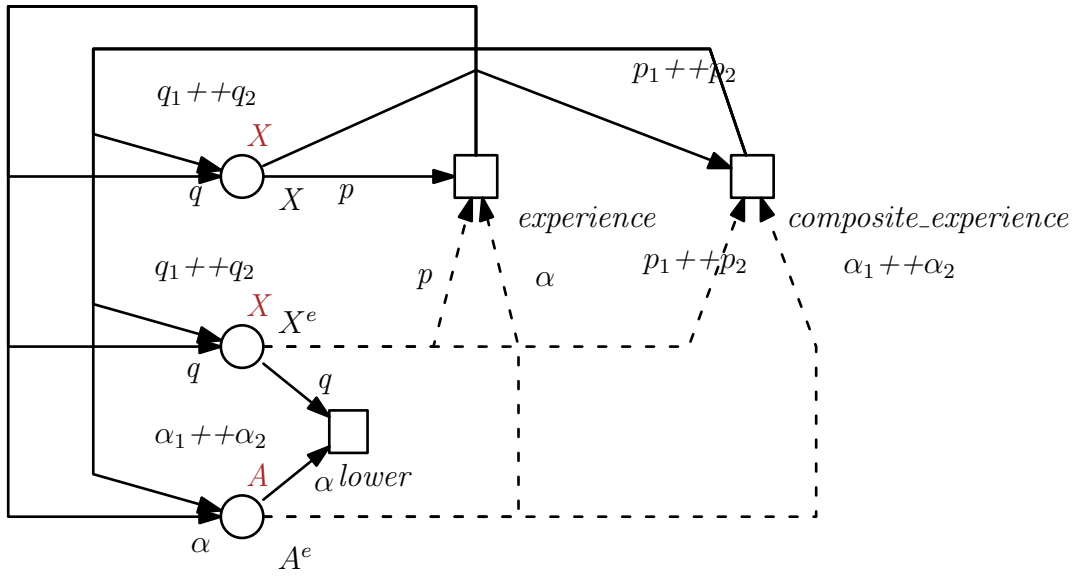
The last term $w^{\text{stick}} \cdot X^{\text{e}}[p]$ creates a pull to stick in states of experience that have been recently entered.

As we will see, the inclusion of the normalisation term N ensures that as the number of strands is increased, the evolution of the system state approaches dynamics described by differential equations. Combined with the dynamics of equation (5.1) the second and third terms of this equation describe self-exciting Hawkes processes with exponential memory kernels [31, 69].

5.3.2 Composite experience

It will be useful to model collectives that have a composite state formed by strands that must work together to form an experience. These collectives might represent a single person, a pair of people or multiple people connected in some way. The model of figure 5.1 supports only the first-order experiences held by single strands and represented by equation (4.1). Figure 5.2 shows a model that extends this to support the second order experiences captured by equation (4.2) ($G^2 = (G^1_{\text{ts}}, E^2)$). Notice the addition of the *composite experience* transition. This runs concurrently with the *experience* transition. To reduce the size of the models we've introduced Λ^2 as an approximation to E^2. Where E^2 contains pairs of edges from $X \times \Lambda \times X$, Λ^2 contains pairs of the content labelling these edges taken from Λ. This approximation is equivalent when each α in Λ is used on only one edge. The guard function for the *experience* transition is constrained to ignore edges that form part of a composite experience; and the new *composite_experience* transition is guarded appropriately. Notice that when *composite_experience* fires the forum is excited by the content labelling the two component edges in the same way as if each was fired separately.

The rate equations for the transitions *lower* and *experience* can stay unchanged from equation (5.1) and equation (5.2). We do need to update the rate equation for the *composite_experience* transition. A simple choice is to ensure strands move together at the



$g^{composite_experience} : [(p_1, \alpha_1, q_1) \in E] \& [(p_2, \alpha_2, q_2) \in E] \& [\{\alpha_1, \alpha_2\} \in \Lambda^2]$

$g^{experience} : [(p_1, \alpha_1, q_1) \in E] \& [\forall c \in \Lambda^2 (\alpha \notin c)]$

$g^{lower} : (p_1, \alpha_1, q_1) \in E$

Figure 5.2: A Petri net model encoding a forum supporting composite experience.

rate of the slowest,

$$\lambda^{\text{experience}^2}(p_1, \alpha_1, p_2, \alpha_2) = \min\left(\lambda^{\text{experience}}(p_1, \alpha_1), \lambda^{\text{experience}}(p_2, \alpha_2)\right). \tag{5.6}$$

where $\lambda^{\text{experience}}$ is equation (5.2). Another option would be to use something more like the chemist's *law of mass action* where two reactants change form during a substitution reaction.

5.4 Experiments and results

We used our software described in appendix A to explore some dynamics supported by the model of a single forum described in section 5.3.

5.4.1 Flock formation

We began with a collection of strands starting in evenly distributed states of experience on the left component of the experience graph of figure 4.1. We tuned the parameters of equations (5.1) and (5.2) until simulations converged toward synchronised communal experiences. Figure 5.3 shows a few such runs with an increasing number of strands N. Each coloured line represents the number of strands in a certain state of experience over time. The results are consistent with the strands moving around the experience graph together: the tokens are seen to cycle between states x_a, x_b and x_c. Notice that as N is increased the lines converge to the same shapes and follow the same *a*-*b*-*c* cycle—however, because the system starts off in a completely symmetrical state, the first strongly held communal experience to arise remains random with the time until the symmetry breaks tending to increase with N.

5.4.2 Flock life-cycle and synchronisation

We then added a third component to the experience graph. We split the strands equally across the three, with inhabitants of the first component all starting in state x_{1a}, the second in x_{2b} and the third in x_{3c}. Due to the common alphabet of shareable experience $\{a, b, c\}$ these three groups interact. Figure 5.4 shows the time series for the resulting simulation runs. Notice that what would have been three flocks first tear each other apart before slowly reforming into a single flock; and that they do this despite the force that holds strands in a particular state of mind working only within the flocks.

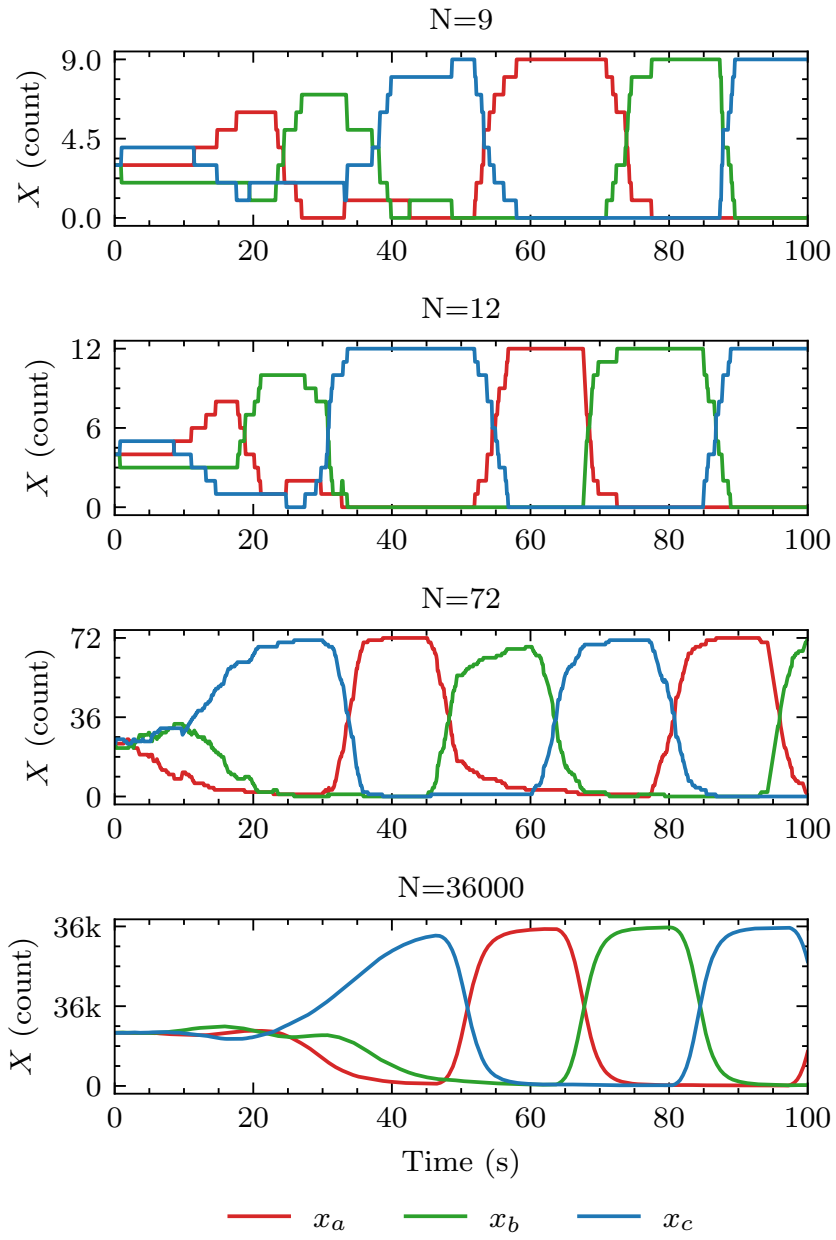


Figure 5.3: Traces showing the emergence of a flock of strands in a forum. Simulations showing convergence from initially evenly distributed states of experience. Each line shows the number of strands in a give state over time. Weights were $w^{\text{self}} = 0.1$, $w^{\text{pull}} = 1$, $w^{\text{stick}} = 1$, $\delta = 0.2$, with N varied by figure.

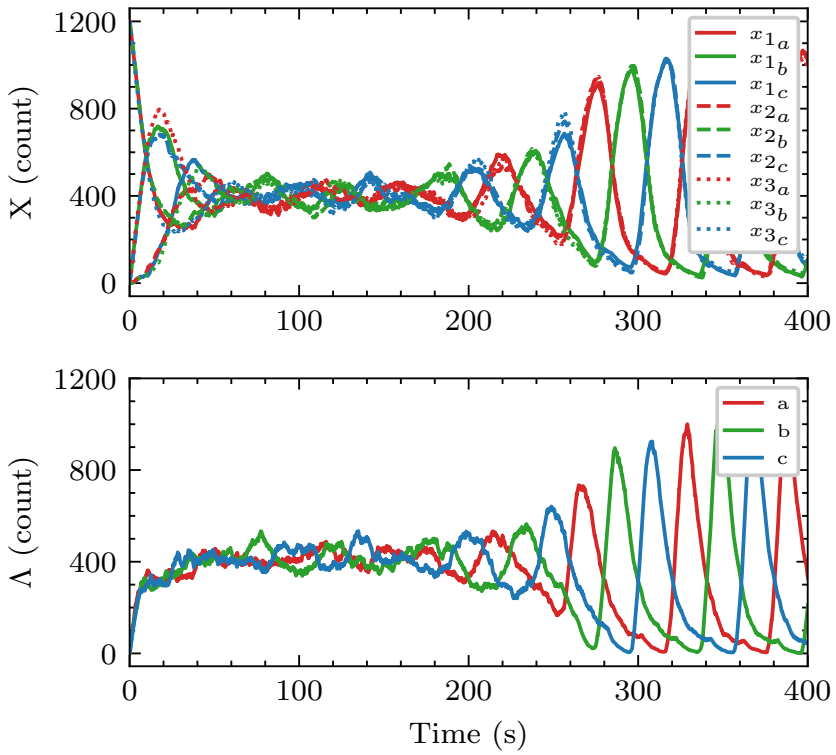


Figure 5.4: Simulations showing the breaking of three flocks and the formation of a single flock. Weights were $w^{\text{self}} = 0.05$, $w^{\text{pull}} = 0.5$, $w^{\text{stick}} = 1$, $\delta = 0.2$, with $N = 3600$. The top plot shows the number of strands in each state of experience; the bottom plot shows the degree of excitement for each piece of content.

5.5 Discussion of results

We discuss the results of the two experiments presented in the previous section. We will interpret the models as representing two scenarios. These scenarios may be viewed as containing social machines and we apply the HTP model of Tinati et al. [84] to account for the results.

The experiments of section 5.4.1 and section 5.4.1 demonstrate some important features of the dynamics of strand interaction. The experience graph of figure 4.1 supports a cycle of three experiences that a strand may undergo. In the figure's caption we gave the example of a cycle of experiences in which strands learn of a news event, then have a reaction to it, and then react to others' reaction to it. This interpretation however does not perfectly fit the simulation which includes no external event to start the cycle. Instead, lets consider the phenomena of the re-telling of a story with a beginning where we meet the story's characters, a middle where they are exposed to some peril, and an end where this is resolved, by a collection of people who know the story well. Relevant only to the second experiment, these people may have differing histories which effect their reaction to each experience in the story.

A single strand by itself will recall the story at their own pace and start again once complete. Each strand would result in a separate graph like that of figure 5.3 with $N = 1$: as the strand moves between the three states of activity x_a, x_b and x_c we would see a cycle of: i) a count in x_a and none in the other states, ii) a count in x_b and none in the other states, and iii) a count in x_c and none in the other states.

5.5.1 Flock formation

In the first experiment, we start with a collection of separate people each hosting a strand of experience in which they retell the story to themselves and that all have the same past histories and will react in the same way to each experience. We then bring these people together such that their strands may interact together to form a forum of strands.

Now, if those strands where all by coincidence in state x_a when brought together, for example at 60s for the cases where N was 9 (or 12 or 36000), then we would see that

their story telling is synchronised and will in all likelihood remain synchronised: a stable phase cycling between i) many counts in x_a and few in the other states, ii) many counts in x_b and few in the other states, and iii) many counts in x_c and few in the other states. What is important is not so much that they form a cycle—although many experiences or activities are recurrent—but that they form a coherent flock of strands. From the perspective of others in the forum it appears as one strand that carries a lot of weight: in that it activates content more strongly than a single strand would and that those that share its state are held to it more strongly than they would that of a single strand.

In the experiment, however, the strands did not all start in the same stage of the story telling. Instead when the individual strands are all brought together, a third of the strands will be recalling the start of the story, a third the middle and a third the end. The forum is initially muddled by the sharing of content associated with each part and only the strands that are recalling the same part will have a chance to experience shared states of mind. It is like there are three overlapping conversations at a dinner table all talking across each other. Those in the same stage of the story will work this out and will focus their attention onto each other and ignore those in the other stages—to form three cohorts. As time passes however some strands may be pulled ahead of their cohort as others lag behind; and very quickly we see a single cohort, or flock, form as the strands become synchronised in their story telling.

HTP account of results. The HTP (Heterogeneous networks—Translation—Phases) model builds on actor network theory (ANT). It conceptualises the process in which networks of heterogeneous parts, that include people and technologies, form around common agendas and influence each other and their agendas: i) an initially unorganised network of **heterogeneous** parts starts with the same agenda; ii) through what ANT calls **translation**, this forms a mobilised network of organised activities advancing that agenda; which in turn iii) effects the networks around them which in turn effects the initial network and its agenda—resulting in **phases** of activity.

In the simulation just described the heterogeneous parts share an agenda to tell the same story together. They are primed to do this but are disorganised and there is confusion

in the resulting social machine. A larger network of parts moving together forms within the initial network during translation. The resulting phase of activity is stable—and even in the unlikely event that random chance breaks the cohesion it will quickly reform.

Another way to account for the results is to zoom into the above interpretation. As above, participants share an agenda to tell the story again and again. We now identify a sub-agenda for each of the three stages of the story. Three networks within this social machine form initially around the agenda to tell the stages (x_a, x_b or x_c) of the story they find themselves focused on. As the rate or arrival of newcomers to the stage they are telling slows a pull to the next stage asserts itself and they move onto this.

Looking at the traces on figure 5.3, for $N = 9$ for example, we see three networks segmented from the larger network with each holding a different agenda. Left separate these, according to the HRT account, would shift between phases as their agendas advance and complete allowing a move to the next agenda. When brought together however at time zero, again a larger network of parts moving together forms within the initial network during translation. With this account however the resulting network cycles between three phases as the agenda cycles. The next scenario is more interesting in that the three networks sustain with some independence.

5.5.2 Flock life-cycle and synchronisation

In the second experiment we take three separate social machines. Each comprises participants that share common paths of past experience; but these histories are different between the machines. These social machines share the common agenda to tell the story and they may work together to do so but each will have its own reaction to the story. They may form experiences in which they appreciate the content associated with each stage of the story-telling together—for example the introduction of characters—in objectively understandable terms. However, the more subjective states of experience that form in reaction to this content will differ across the machines—perhaps the story is about a family of migrants for example.

The strands in each of the three social machines, like that of the simulation just

described, if left separate would act as large single strand or flock. When brought together however this could potentially remain true; but the initial conditions of the simulation had each social machine starting in a different stage of the story telling. There is confusion as the content shared by each social machine during its story-telling clashes in the forum. Looking at figure 5.4, 60 seconds into the simulation all memory of the past order has been lost: the bottom plot shows that the content associated with each phase is excited equally, and the top that participants in each machine are spread equally between the stages of the story telling. However as time advances we see the familiar cycles from the previous example reemerging: the three social machines are advancing a common purpose by telling the story together, while each forms its own experiences as informed by their differing histories. These differing histories: might carry some residual mood of the participants' into the experiences formed during the story telling; or, perhaps they might instill differing beliefs that their interpretation of the story must remain compatible with; or, cause them to favour interpretations that advance an argument they favour;or perhaps their differing histories might have instilled differing agendas—such as to bring about a humanising or dehumanising of immigrants.

HTP account of results. Again we identify an agenda to tell the story that is common to a single social machine comprising all participating strands of experience; and again we identify a sub-agenda to tell each stage of the story. Now we add an agenda that differs between the three initial social machines but that is prevalent within them—this could be to maintain a certain mood, belief, or to advance a certain state of affairs, for example. At time zero we see that the three networks each have agendas to tell different stages of the story. Translation in each network is hindered by the differing agendas of the other networks and the phase of order each held before they came together is lost. As these separate networks seek translation however so does the whole driven by an agenda to tell the story. By 200 seconds into the simulation it is clear that a new network has formed whose parts have managed to align their story-telling agendas such that they are able to tell the story together. This is interesting as this new network is cut equally from the networks of the three initial social machines. As time progresses the remaining

participating strands respond to the translation of this core network by aligning themselves with it. Throughout this evolution the networks associated with the three initial social machines maintain their differing agendas and in turn their incompatible subjective states of experience. The newly formed network consisting of the original actors translates by keeping the agenda of the actors of each of the original networks aligned—at least as the cycling phases of agenda come and go while the story is told.

5.6 Conclusion

The work described here builds on the ontology and generative model proposed in chapter 4 in which experiences build on each other; with past experiences filtering which content raises into awareness and how it is experienced. We defined a forum as a very simple community comprising a number of strands, each described by an instance of this generative experience model. These strands cede their level of excitement about both entering a state of experience and having experienced particular content to the whole. These models support an indirect feedback that shared content or state has been followed in that the communal excitement level goes up. This forum supports both the small scale interaction of a few strands and very large scale communal experiences. Our simulations show that strands comprising large forums may act like a single coherent flock. This supports our larger argument that the pre-existing strands we start with may always in reality be emergent.

As well as this convergence the model also supports divergence. strands in a forum may undergo incompatible experience causing them act as multiple forums. At this stage they would continue to share the same single channel for sharing excitement about content—but with separate channels for sharing excitement about each state of mind implicit in the model. Upcoming models will support the complete separation of forums as well as their joining.

The example we chose was contrived in that components of the experience graph formed closed loops allowing strands that lag behind the bulk of the flock another chance to latch on during the next cycle. Here and later in the thesis we hard code the dynamics

in a way that required tuning the model parameters to achieve flocking. In reality the dynamics are likely more complex and self regulating—with people's drive toward social engagement coming and going based on how recently they've engaged for example.

Chapter 6

Forums as static networks of forums

6.1 Single layer networked forums model 106
6.2 Analysis and application 108
6.2.1 Simplified form 108
6.2.2 Proof of self-similar dynamics 109
6.2.3 Scenario: Two forums connected by a strand 111
6.3 Experiments and results 111
6.3.1 Fully connected network of forums 112
6.3.2 Synchronisation on a grid 112
6.4 Conclusion 114

We've developed a model of a single forum that governs the dynamics of its component pre-existing strands along paths of possible simple and composite experience; and we've shown that this forum supports the emergence of both short-lived emergent strands formed during the coordinated experience of multiple pre-existing strands, and of longer lived flocks that themselves act like strands. We now expand the model to support a network of interacting forums. This must be done in a way that closely connected forums similarly situated in a broader network may act like a single forum in this broader network.

In this chapter, and the one following, we temporarily step away from the overlapping forums of strands model introduced in section 5.1.1. In section 6.1 we will add a following network to the model allowing forums to follow one another directly—rather than strands

interacting with strands through mutual overlap in the same forums. For forums containing multiple strands, this is a useful approximation of the overlapping forums model, in which the strength of the tie between two forums is proportional to the number of strands common to both. Of more importance to the main narrative of the thesis is that, we can choose the dynamics of strand evolution within a forum such that, with just one strand per forum, a collection of forums that all follow each other equally have *exactly* the same dynamics as if those strands were included in just a single forum. We demonstrate this through closed-form analysis in section 6.2. Then in section 6.3 we show the same result through simulation. This is an important result, as it demonstrates the self-similarity we consider important to developing a theory of social machines as captured in assumptions (A4) and (C2). We will return to the overlapping forums model in chapter 8 and for the remainder of the thesis. We defer a formal addition of the term 'binding' to our ontology until chapter 9.

6.1 Single layer networked forums model

When studying social networks like Twitter or Reddit it is useful to define a few networks. The *following network*, N^f, and the *sharing network*, N^s, are weighted directed graphs cut from an undirected *connection network* N^c. In our model a forum will share with any forum following it; and so we encode only a following network. We define the connectivity between forums S in the form of a directed graph (S, F), where F is a set of edges taken from $S \times S$; and a function F^c over F acting as edge weights. This network describes the potential rate of shared experience from one forum to another. The actual rate of shared experience will depend on the compatibility of the state of mind of each forums' strands. In reality these of course relate—and in the next chapter we will begin addressing this.

We extended the model of a single forum given in figure 5.1 to support multiple forums resulting in the model shown in figure 6.1. We modified places X, X^e, and Λ^e to represent the state of all forums, by expanding them hold tokens coloured by $S \times X$ and $S \times \Lambda$. The variable s represents a forum of interest and r a potential neighbour that this may follow. The guard function on the *experience* transition has been expanded so that a forum is

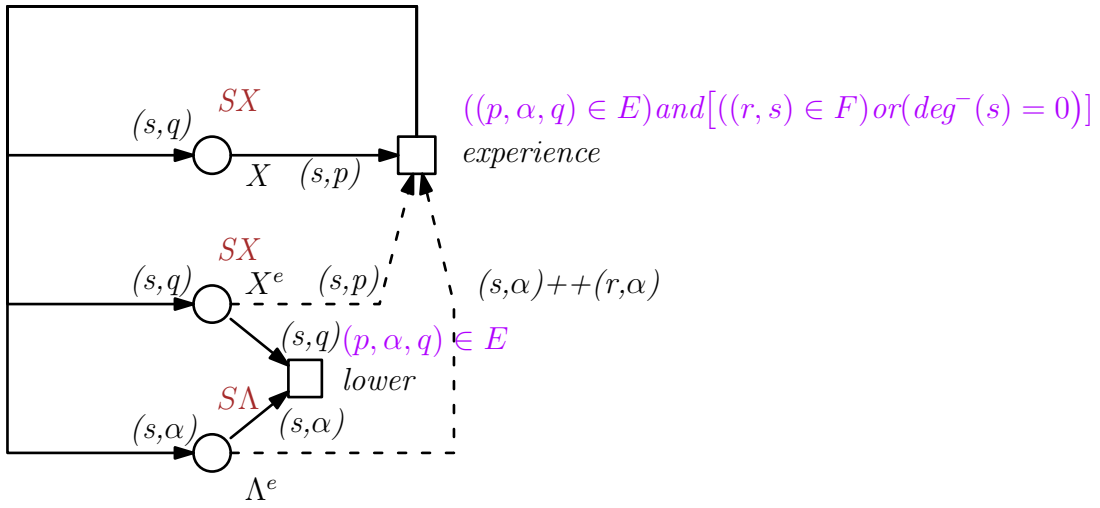


Figure 6.1: A Petri net model encoding a network of forums.

only effected by those it follows. There are three cases for the *experience* transition rate function. Where a forum s follows no other forum the rate is just that of λ^{forum} from equation (5.2) and if it does have followers then we add a following term to this:

$$\lambda^{\text{experience}} = X[s,p]\cdot\frac{X[s,g]}{N_s}\cdot\begin{cases}\lambda^{\text{forum}}(s,\alpha,p), & \text{if } (p,\alpha,q)\in E\wedge \deg^-(s)=0\\ \lambda^{\text{forum}}(s,\alpha,p)+\lambda^{\text{follow-pull}}(s,r,\alpha), & \text{if } (p,\alpha,q)\in E\wedge(r,s)\in F\end{cases} \tag{6.1}$$

where

$$\begin{aligned}\lambda^{\text{forum}}(s,\alpha,p) &= \lambda^{\text{self}}+\lambda^{\text{pull}}(s,\alpha)-\lambda^{\text{stick}}(s,p)\\ \lambda^{\text{self}} &= w^{\text{self}}\\ \lambda^{\text{pull}}(s,\alpha) &= \frac{1}{N_s}w^{\text{pull}}\cdot\Lambda^{\text{e}}[s,\alpha]\\ \lambda^{\text{stick}}(s,p) &= \frac{1}{N_s}w^{\text{stick}}\cdot X^{\text{e}}[s,p]\\ \lambda^{\text{follow-pull}}(s,r,\alpha) &= \frac{1}{N_r\deg^-(s)}w^{\text{follow-pull}}\cdot\Lambda^{\text{e}}[r,\alpha]\end{aligned} \tag{6.2}$$

and $w^{\text{inter-pull}}$ is a weight determining how hard neighbouring forums' excitement over content is felt, N_s is the number of strands in forum s, and $\deg^-(s)$ is the indegree of forum s indicating the number of forums it follows. We have not added a shared $w^{\text{stick-inter}}$ term; that is there is no attraction between flocks formed in different forums—although

this could be added easily.

The rate function on the transition *lower* must be expanded out from that of equation (5.1) so that each forum's excitement decays independently:

$$\lambda_{lower} = X^e[s, q] \cdot \delta.$$

6.2 Analysis and application

The implementation of a forum captured in this chapter and in the one that follows, has an inherent limitation: strands cannot explicitly form part of more than one forum. To represent this we equate the concept of a strand with a forum by showing that fully connected networks of forums may act like a single forum.

6.2.1 Simplified form

We start by relating the single forums model of section 5.3 with the networked forums just described in equations (6.1) and (6.2). We can simplify this representation by adding for each forum a self loop, representing internal connections, to the directed graph of inter forum connections and writing the rate equation:

$$\lambda(s, p, \alpha) = X[s, p] \cdot \left(w_{self} + \sum_{(r,s) \in F} w(r, s) \cdot (w_{pull} \cdot \Lambda^e[r, \alpha] - w_{stick} \cdot X^e[r, p]) \right) \quad (6.3)$$

where $w(r, s)$ is the weight of the edge from neighbour r to the system of interest s. Let's see how this works by recreating some models from the previous chapters.

Single forum. To recreate the dynamics of the single forum introduced in section 5.3 we let $S = \{s\}$ where s is our single forum, $F = \{(s, s)\}$ and $w((s, s)) = \frac{1}{N_s}$; where N_s is the number of strands in S ($N_s = \sum_{p \in X} X[s, p]$). The single edge in in F requires that $r = s$. Using this reduces equation (6.3) down to

$$\lambda(s, p, \alpha) = X[s, p] \cdot (w_{self} + \frac{1}{N_s} \cdot (w_{pull} . \Lambda^e[s, \alpha] - w_{stick} \cdot X^e[s, p]))$$

which, if we drop the redundant s as there is only one forum, is equivalent to equation (5.2) as desired.

Networked forums. To recreate the dynamic of the networked forums above we create a set S of forums and a network F of edges between them as before; however this time we also add a self edge (s, s) for every forum s of S. As above we then define the weights: first adding the weight $w((s, s)) = \frac{1}{N_s}$ for all s; and then $w((r, s)) = \frac{1}{N_r \cdot \deg^-(s)}$. Notice though that this results in a contribution from neighbouring forums to the force that sticks strands in a state. Otherwise the resulting dynamics match those of equation (6.1) if the weights between neighbouring forums are further scaled by $w_{follow-pull}/w_{pull}$.

6.2.2 Proof of self-similar dynamics

Equation (6.3) supports configurations of forums that act like a single forum that can in turn be described by this equation. This is an example of the self-similar property we claim for forums and social machines. To demonstrate let's take N forums each comprising a single strand and fully connect these in a way that they are indistinguishable. This is how we defined the forum in section 5.3. The forum model is self-similar if the dynamics are unchanged by redrawing the forum boundaries around these strands to create a single forum with N strands.

We first create a set S of N forums each with one strand such that $N_s = 1$ for all s in S. For each pair of different forums (r, s) we add a link to F and set the weight $w(r, s)$ to $w((r, s)) = \frac{1}{N_r \cdot \deg^-(s)}$ or $1/N$. As above we add a self link to F for each and in order that each forum pay its single strand no more or less regard than their neighbours we also the weights $w((s, s)) = 1/N$ for all forums s.

We seek the rate functional summed across all forums in S,

$$\lambda(p, \alpha) = \sum_{s \in S} \lambda(s, p, \alpha)$$

where $\lambda(s, p, \alpha)$ was just proposed in equation (6.3). Let's write this out, noticing $w(r, s)$

is always $1/N$ and expanding out the *pull* and *stick* terms:

$$\lambda(p,\alpha) = \sum_{s\in S} X[s,p] \cdot \left(w_{self} + \frac{w_{pull}}{N} \sum_{(r,s)\in F} \Lambda^e[r,\alpha] - \frac{w_{stick}}{N} \sum_{(r,s)\in F} X^e[r,p] \right)$$

As F is a complete graph with self loops, then for a given s, $\sum_{(r,s)\in F}(\cdot)$ is just $\sum_{r\in S}(\cdot)$. This substitution results in

$$\lambda(p,\alpha) = \sum_{s\in S} X[s,p] \cdot \left(w_{self} + \frac{w_{pull}}{N} \sum_{r\in S} \Lambda^e[r,\alpha] - \frac{w_{stick}}{N} \sum_{r\in S} X^e[s,p] \right)$$

Now we define $\sum_{s\in S} \Lambda^e[s,\alpha]$ to be simply $\Lambda^e[\alpha]$ (and likewise for the sums over $X[s,p]$ and $X^e[s,p]$) reducing the rate equation to

$$\lambda(p,\alpha) = \sum_{s\in S} X[s,p] \cdot \left(w_{self} + \frac{w_{pull}}{N} \Lambda^e[\alpha] - \frac{w_{stick}}{N} X^e[p] \right)$$

and finally

$$\lambda(p,\alpha) = X[p] \cdot \left(w_{self} + \frac{w_{pull}}{N} \Lambda^e[\alpha] - \frac{w_{stick}}{N} X^e[p] \right)$$

as we sought to demonstrate.

This is an important result as it shows that an aggregate of strands across forums may act as if they are in a single forum. It also affords greatly simplified models. For example the models of Figure 6.2 are equivalent.



(a) Networked strands

(b) Collapsed into networked forums

Figure 6.2: Equivalence between strands on a network and strands in forums. What might have been one forum is held apart by differing relationships to the rest of the network indicated by dotted lines.

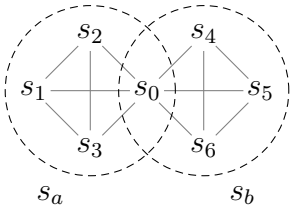


Figure 6.3: A scenario where one person has strands overlapping with two forums. Seven forums $s_0 - s_6$ connected to form two fully-connected virtual, or emergent, forums s_a and s_b, linked by a common component s_0.

6.2.3 Scenario: Two forums connected by a strand

When a person is a member of two groups that person may act to connect those groups. The seven forums $s_0 - s_7$ shown in figure 6.3 might each represent a person. If it were not for s_0 these would form two fully connected forums s_a and s_b. In this case we might represent this equivalently as either two independent forums or as two fully connected but separate components of four forums each.

Now if the strands comprising the forums in the two groups exist on parts of the experience graph that don't interact, then in principle we may still approximate the dynamics as two separate forums, basically splitting s_0 into two. This would be approximate as s_0's connection weight to each group would be halved splitting its attention between the two groups. Otherwise if the groups exist on the same part of the experience graph then s_0 will cause them to act essentially as one forum but one with reduced cohesion. The community formation dynamics reviewed in section 3.2.3 would tend to cause strands from each side to come to know each other through s_0—we add a behaviour supporting this in chapter 9.

6.3 Experiments and results

We tested the dynamics of this model. We first confirmed that a fully connected network of forums can converge in cases equivalent to models in which a single forum would converge. We then looked for the same convergence on a model with spatial structure.

6.3.1 Fully connected network of forums

We created models with N forums all following each other resulting in $N(N-1)$ directional edges. We used a small experience graph with just two states x_a and x_b connected in a way that forms a cycle. Each forum was started with strands shared equally between these two states, making everything symmetrical. We varied the number of strands per forum but kept this the same across the network.

We saw that each forum initially settled in one state randomly, and then, as the excitement in that state wore off, was pulled by their neighbours away from this to the other state. Figure 6.4 shows the number of strands in state x_a for each forum as a function of time for the case where there were 3600 strands per forum. We see that the two flocks converge after 8 or 9 cycles. With this number of strands per forum the period of oscillation was very uniform between runs. As the number of forums was increased the period of oscillation was seen to remain around 24 time units; but there was a tendency for larger networks to take longer to synchronise.

6.3.2 Synchronisation on a grid

We then performed exactly the same experiment but with the forums spread across a three row by four column grid. The traces in figure 6.5 show results from an example run; where the traces are separated by row. In this case we see that the bottom row happened to stay synchronised throughout and that the rest of the system converged after 4 or 5 cycles.

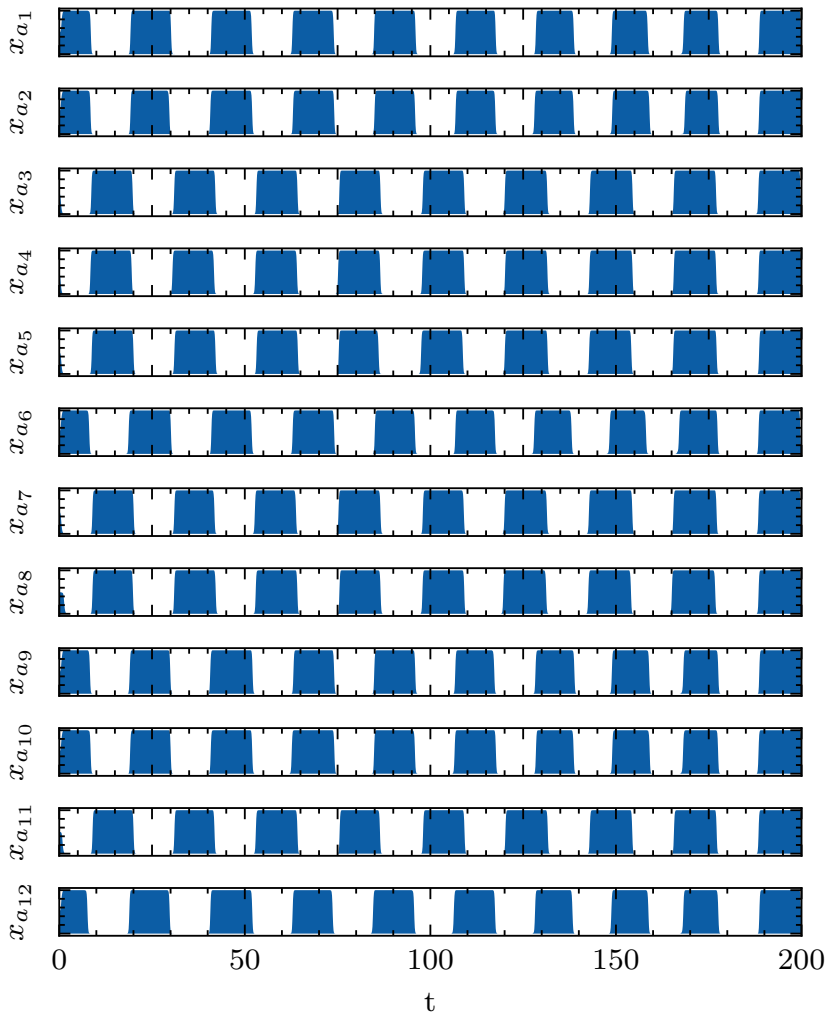


Figure 6.4: Traces showing synchronisation across a network of fully connected forums. Each trace shows the number of strands in state x_a over time for each of the 12 forums. The seed to the simulator was 1.

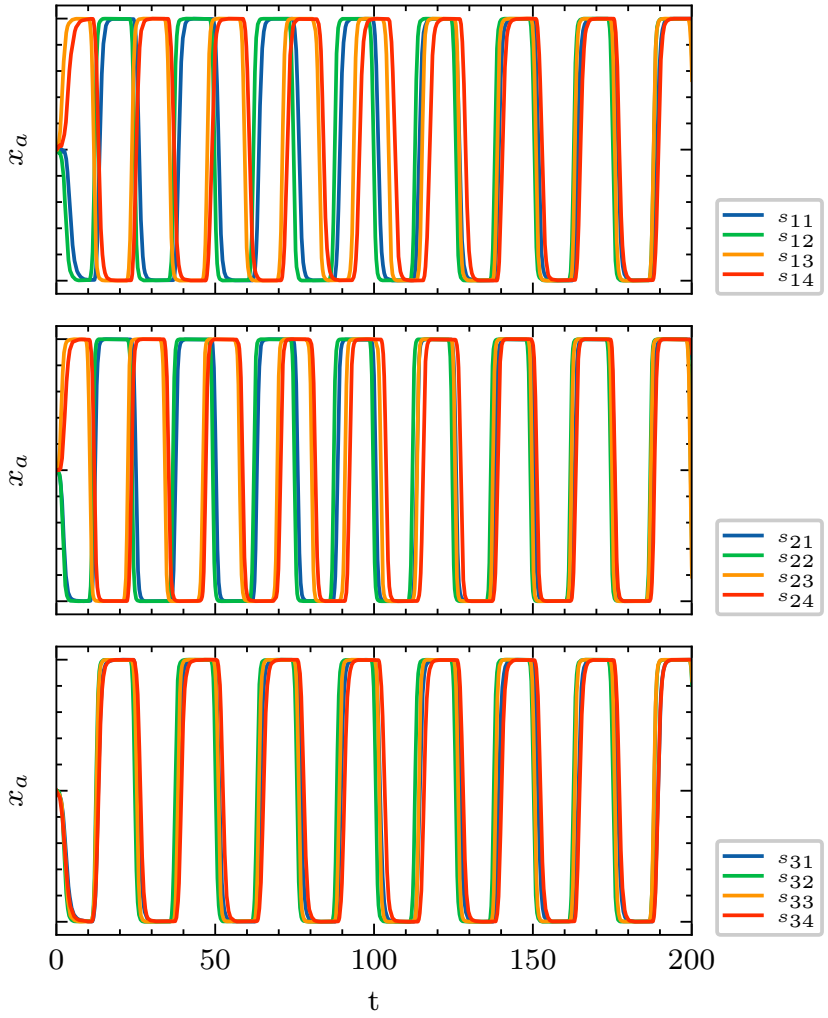


Figure 6.5: Traces showing synchronisation across a 3 by 4 grid of mutually connected forums. Each trace shows the number of strands in state x_a over time for each of the 12 forums. The seed to the simulator was 2.

6.4 Conclusion

We added a following network to the model allowing forums to follow one another with a certain weight. Just as strands in a forum cede their excitement about their state of mind and of content to the whole they also share a common connection to other forums. The dynamics governing a number of forums of single strands on a fully-connected following network match those of a single forum containing those same strands. We proved this and showed it experimentally. We demonstrated that the same flocking behaviours that arise in a single forum can occur across a network of forums.

Chapter 7

Dynamic networks of forums

7.1 Modelling dynamic binding between forums 116
7.2 Axelrod's model of dissemination of culture 119
7.2.1 Representation with our model 119
7.3 Experiments and results 122
7.4 Discussion and conclusion 125

In this chapter, as with the preceding one, we temporarily step away from the overlapping forums of strands model introduced in section 5.1.1. We further develop the model in which forums follow one another directly—rather than strands interacting with strands through mutual overlap in the same forums. With this model, we have shown that we can contrive scenarios in which a forum of pre-existing strands may act as one strand and that a network of forums following each other may act as one forum. We now need to add a mechanism that evolves the network in a way that these contrived structures form naturally.

In section 7.1 we propose a simple mechanism for doing this. Starting with a pre-existing network of connections between forums we add a mechanism for a forum to become excited about an edge in this network when following content from that edge. We then verify this binding mechanism and exercise the entire model. We need to show that forums emerge across regions of a network of forums to form a new network with fewer edges. We

therefore looked for a scenario that encouraged the growth of few larger forums—one in which differing paths of past experience did not prevent convergence while at the same time did not always converge to a single forum without differences. To this end, recall that we planned to build a scenario that could be represented by Axelrod's model and show through simulation that it tends to produce larger emergent forums. We develop this scenario this in section 7.2 and show that it works in section 7.3.

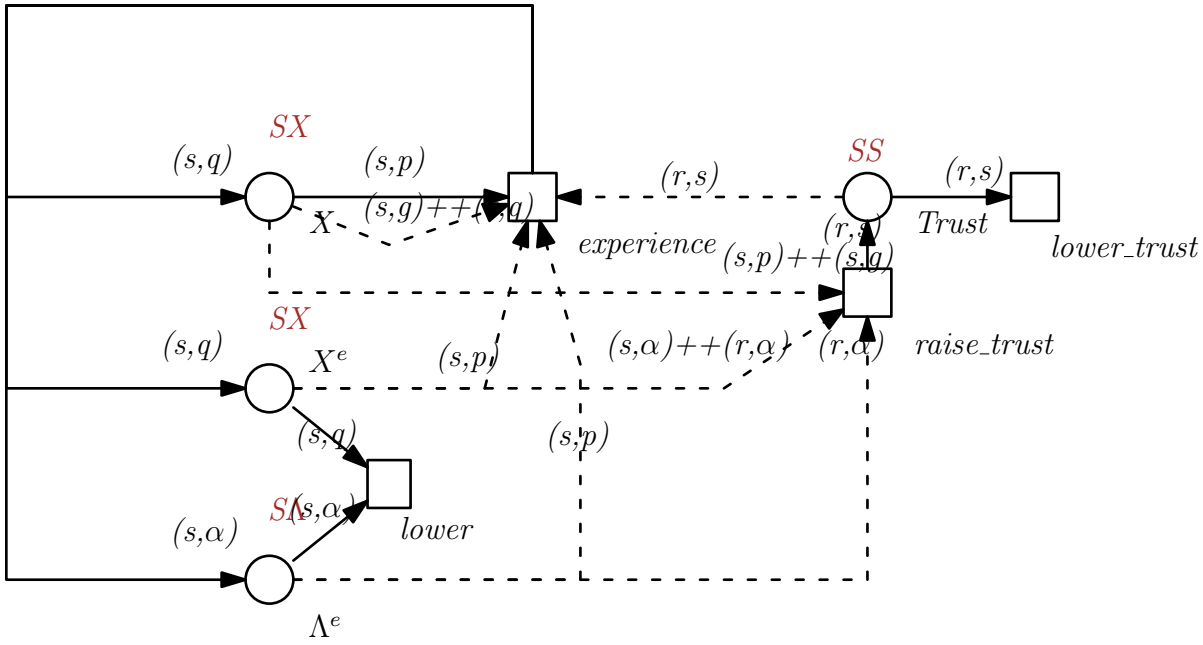
We will return to the overlapping forums model in the next chapter, and for the remainder of the thesis. In chapter 9 we remove the need for the static underlying network, introduced in this chapter, and formally add a 'binding' mechanism to our ontology.

7.1 Modelling dynamic binding between forums

Starting with a pre-existing network of connections between forums we need a mechanism for a forum to become excited about an edge in this network when following content from that edge. In this way forums that interact more will bind more tightly than those that interact less. We might achieve this by making the weights between forums from our last model described in section 6.1 dynamic. We tried this but found it hard to prevent positive feedback from irreparably breaking connections—and at this stage we have no way to reform them. Instead we used the generator of section 4.5.2 that supports special experiences that require longer exposure to content or more intense exposure to content to form. Here we interpret these experiences as requiring trust to follow. Where a forum becomes excited by another this excitement effects just these special experiences; thus leaving the capacity for forums to interact through sharing standard experiences across the static edges even if they are not excited by each other.

To this end we expand the model of figure 6.1 to create figure 7.1 by adding two new features: a shared feeling of trust or excitement between a forum and each other forum; and edges in a generator graph that require this trust to follow as described in section 4.5.2. This model also supports composite experience based on the social guards described in equation (4.4).

The guard function on the *experience* transition, $g^{experience}$, supports both standard



$g^{experience} : [(p, \alpha, q) \in E \vee (p, \alpha, q) \in E^{trust}] \wedge [(r, s) \in F \vee deg^{-}(s) = 0] \wedge (p, g) \in G$

$g^{raise_trust} : [(p, \alpha, q) \in E] \wedge [(r, s) \in F] \wedge [(p, g) \in G]$

$g^{lower} : [(p, \alpha, q) \in E] \vee [(p, \alpha, q) \in E^{trust}]$

Figure 7.1: Petri net model for forum supporting trust or excitement toward neighbouring forums.

and special experiences but there is a separate rate function for each case. For special experiences the rate function will depend on the level of trust or excitement between a forum of interest s and each neighbour r. We will return to this rate function shortly. In the same way that each forum s has decaying levels of excitement about being in a state, as captured in X^e, and decaying excitement about each piece of shareable content, as captured in Λ^e, we add a decaying level of trust or excitement as captured by the new place T. This place ranges over $S \times S$ with a token in (r, s) indicating that forum s trusts forum r. In keeping with the decay rate represented by the transition *lower* trust is assumed to decay at the same rate

$$\lambda_{lower_trust} = T[r, s] \cdot \delta.$$

Returning to the rate equation for the transition *experience* we add another term to equation (6.1) that is used when an edge requires trust. We choose that strands within a forum implicitly trust each other and so include the contribution $\lambda^{\text{forum}}(s, \alpha, p)$. However we choose that strands should not spontaneously undergo experience that requires trust so remove the λ^{self} term. We add a term that determines the rate at which a strand in forum s will follow a strand in forum r into state q:

$$\lambda^{\text{trusted}}(s, r, q) = w^{\text{trust}} \cdot \frac{1}{\deg^-(s)} \cdot \frac{T[r, s]}{N_s} \cdot \frac{X[r, q]}{N_r}$$

The rate is proportional to a number of factors: w^{trust} can be tuned to determine the influence of trusted neighbours generally; $1/\deg^-(s)$ prevents the forum's internal dynamics from being overwhelmed as it follow more neighbours, essentially splitting its attention across them all; $T[r, s]/N_s$ is the average trust that the N_s strands in forum s share about forum r; and finally $X[r, q]/N_r$ is the proportion of a followed forum r that is in the target state q.

Making these changes to equation (6.1) results in the rate equation

$$\lambda^{\text{experience}} = X[s,p] \cdot \frac{X[s,g]}{N_s} \cdot \begin{cases} \lambda^{\text{forum}}(s,\alpha,p) & \text{if } (p,\alpha,q) \in E \wedge \deg^-(s) = 0 \\ \lambda^{\text{forum}}(s,\alpha,p) + \lambda^{\text{follow_pull}}(s,r,\alpha) & \text{if } (p,\alpha,q) \in E \wedge (r,s) \in R \\ \lambda^{\text{forum}}(s,\alpha,p) - w^{\text{self}} + \lambda^{\text{trusted}}(s,r,q) & \text{if} (p,\alpha,q) \in E^{trust} \end{cases} \cdot \tag{7.1}$$

7.2 Axelrod's model of dissemination of culture

In this section we will recreate Axelrod's work presented in section 3.2.2 using our techniques in order to help verify these techniques and complete our first major goal.

Formalising Axelrod's model we define a set of N^{f} features $R = \{r_1, r_2, \ldots, r_{N^{\text{f}}}\}$. For each feature f_i there is a set of $N^{\text{t}}_{f_i}$ traits $T_{f_i} = \{t_{f_i 1}, t_{f_i 2}, \ldots, t_{f_i N^{\text{t}}_{f_i}}\}$. A person's cultural belief c can be represented as a list of these traits ordered by feature; for example $c = (t_{f_{12}}, t_{f_{23}}, t_{f_{31}})$ represents a the person holding the second trait for feature 1, the third trait for the feature 2 and the first trait for feature 3. An $n_{\text{rows}} \times n_{\text{columns}}$ grid of sites can then be defined.

7.2.1 Representation with our model

We now need to represent this model and something approximating the dynamics just described. We'll start by defining a generator graph and then create a grid of networked forums to represent the grid of Axelrod's sites. For reasons described below the forums must support the trust and social guard mechanisms supported by the model just introduced.

To match Axelrod's model each person will be represented by a forum containing at least one strand per feature; with the movement between traits caused by the influence of neighbouring forums. To build the experience generator we will need a state for every possible trait x_{ft} where f is the index of the feature and t the index of a possible trait for that feature, such that

$$\{x_{ft} \mid f \in R \wedge t \in T_f\} \in X.$$

We will need a way for strands to change their mind about which trait they hold. Firstly, for each trait t in each feature f we add shareable content α_{ft} that will take strands to state x_{ft}, such that

$$\{\alpha_{ft} \mid f \in R \wedge t \in T_f\} \in \Lambda.$$

We need strands to follow neighbouring forums at a rate proportional to how similar their cultural traits are. Let's assume for now that we have a way to grow trust between forums based on this similarity. For each feature f we visit each target trait t and add trusted links from each other trait s in that feature to that trait t; these we label with the shareable content that will lead to the target trait, such that

$$\{(x_{fs}, \alpha_{ft}, q_{ft}) \mid f \in R \wedge t \in T_f \wedge s \in T_f \wedge s \neq t\} \in E^{\text{trust}}.$$

We now need to add this way for a forum to trust another to a degree proportional to the number of cultural features they agree on. A natural way to achieve this is to provide a way for forums to discuss traits that they agree on and for this discussion to result in trust. One solution is to create a loop of experience for each trait—perhaps representing small talk or some other form of trait reinforcement—that is socially guarded such that this loop can only be experienced when a forum contains strands holding the trait it depends on. Figure 7.2 shows an example experience graph capturing this. Finally we add a pair of states for each trait such that

$$\{x_{a_{ft}}, x_{b_{ft}} \mid f \in R \wedge t \in T_f\} \in X,$$

and edges between these such that

$$\{(x_{a_{ft}}, \alpha_{b_{ft}}, x_{b_{ft}}), (x_{b_{ft}}, \alpha_{a_{ft}}, x_{a_{ft}}) \mid f \in R \wedge t \in T_f\} \in E,$$

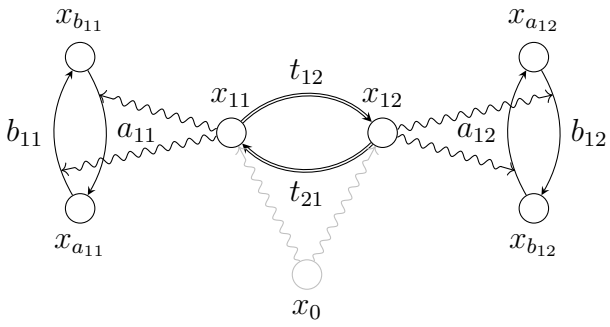


Figure 7.2: The experience graph for one feature of our representation of the Axelrod model. There will be a strand in x_0; one strand in either x_{11} or x_{12} representing the first or second trait of feature 1 respectively; and one strand in each of the two conversational loops at the top. The dynamics will enable traversal around these conversational loops only if the prerequisite trait is held in a forum. Conversation experienced with strands in neighbouring forums will provide the trust required for strands to follow their neighbours between the states representing the feature's traits.

and the shareable content to cycle between these states must be included such that

$$\{\alpha_{b_{ft}}, \alpha_{a_{ft}} \mid f \in R \wedge t \in T_f\} \in \Lambda.$$

We must socially guard these states so edges between them represent composite experience requiring a strand to believe in the trait they depend on. That is

$$\{(x_{a_{ft}}, x_{ft}), (x_{b_{ft}}, x_{ft}) \mid f \in R \wedge t \in T_f\} \in G.$$

Note that although the states representing our traits don't require composite experience to leave, for completeness they must depend on some other state. We add $x_0 \in X$and make the traits for all features socially dependent on this, such that

$$\{(x_{ft}, x_0) \mid f \in R \wedge t \in T_f\} \in G.$$

Figure 7.2 represents the experience graph we've just described for one feature f_1.

7.3 Experiments and results

We modelled a number of specific scenarios to explore this method for representing Axelrod's model. We were limited by the software stack to quite small grids of around 4 by 3 with two features both with two traits each and something to chat about. As described in appendix A this limitation arises because every model was expanded out to a non-coloured Petri net resulting in very fast growing models. This is a downside to keeping the mathematics tractable. More customised approaches could be optimised to run this form of simulation more quickly and at larger scale.

Verification of underlying model

Before running these simulations we tested the behaviour of the underlying model captured in figure 7.1 and our implementation of it. To confirm that the guarding mechanism worked: we tested very simple cases building up to a scenario in which strands remain synchronised across a network while moving on an experience graph that can be interpreted as the cycling of a two-bit counter. We then verified that as neighbouring forums moved together that their mutual connection was reinforced; and that where they could not interact due to their incompatible positions on the experience graph that their connection was not reinforced. The software includes an option to recreate these model runs.

Verification of steady state configurations

We then verified that on a three by four non-periodic grid of forums that configurations which should be stable were indeed stable. With all forums agreeing on their traits there should be no opportunity to change traits. We saw just this, with the forums synchronising on their small talk loops discussing the traits they held, and these traits never changing. We also designed a number of initial configurations in which there were regions with completely non-overlapping traits. The forums within a region held small talk but the two regions were unable to interact and the traits never changed. Again the software includes an option to recreate these model runs.

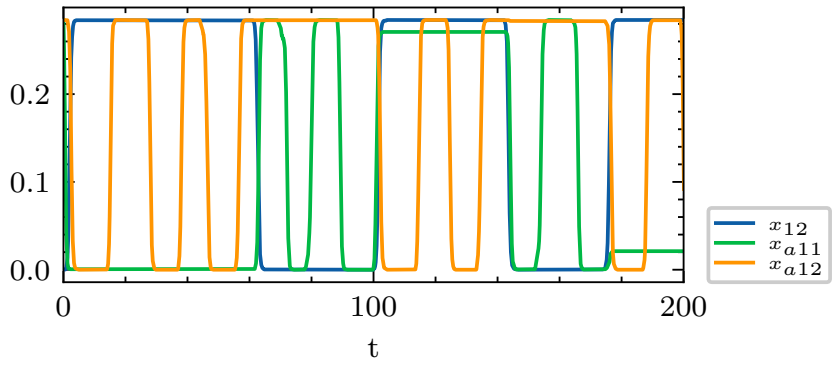


Figure 7.3: Traces showing one forum's changing culture and corresponding chatter. Here x_{12} is the portion of forum s_{23} in the second trait of feature 1. Notice that x_{a11} stops changing when the forum is in the second trait while x_{a12} continues to change indicating ongoing chatter dependent on the trait. The random simulation seed was 1 and the initial trait seed was 7.

Verification of Axelrod implementation

On the same three by four non-periodic grid of forums we then started with randomly assigned traits and positions on the small talk cycles. Figure 7.3 shows illustrative traces for system s_{23}. The slower changing trace x_{12} shows the number of strands believing the second trait of the first cultural dimension. While this is high strands go around this trait's chat loop as evidenced by the changing trace x_{a12}. When it goes low this indicates that the system believes in the second trait and that the strands go around the other chat loop as evidenced by the changing trace x_{a11}.

The majority of runs converged to a homogeneous culture within a few hundred time units. Some runs settled into multiple regions. The traces in figure 7.4 show the changing fraction of each forum in each trait during such a run. By looking carefully at these traces we infer the state of culture across the grid. Figure 7.5 shows the dominant cultural traits in each forum at the start of the run, during the stable period around 20 time units into the run and at steady state. The links between the systems indicate the cultural overlap between neighbours—with a thick line indicating two shared traits, a thin line indicating one shared trait, and the absence of a line indicating no shared traits. In this run there is a large group with the traits 22 and a little group comprising s_{33} and s_{34} with the traits 11. The model also allows access to the directional measure of trust between each system and its neighbours. Although not showed here, this was seen to be consistent.

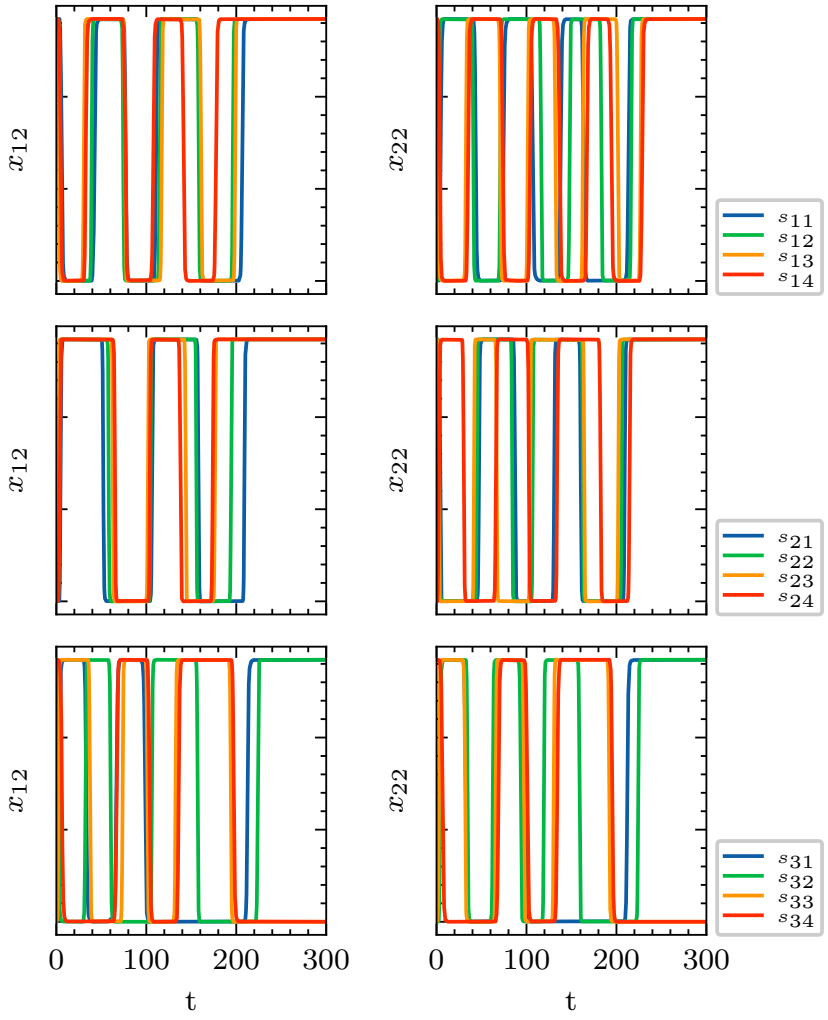
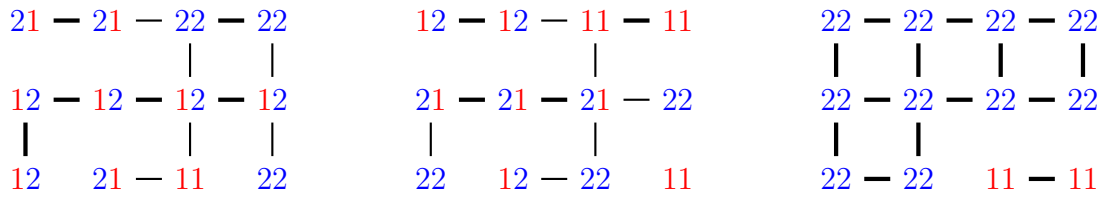


Figure 7.4: Traces showing the changing culture across a grid of forums during an Axlerod model run. Same run as above



(a) $t = 0$ **(b)** $t = 20$ **(c)** $t > 250$

Figure 7.5: Cultures at three times during an Axelrod model run. Same run as above.

7.4 Discussion and conclusion

At the start of this chapter we had shown that we can contrive scenarios in which a forum of pre-existing strands may act as one strand and that a network of forums following each other may act as one forum. We further expanded the model such that when a strand integrates content from an external forum its host forum becomes more strongly bound to that external forum. This model assumes that the host forum already has a connection to the external forum—we address this later.

We tested this new binding mechanism and indeed the bulk of the rest of the modelling approach developed in the thesis to this point. We did this by modelling a scenario described by Axelrod's model of the dissemination of culture. We could not build models with as many people or cultural traits as Axelrod as the size of the resulting Petri nets grew too quickly to execute. Using Petri nets is rigorous but models do get very large when expanded out onto uncoloured Petri nets for execution. Custom code could handle larger models and run a lot more quickly—we suggest this as future work in section 13.2.

In this way we have reached our first major goal of the thesis. We have shown that after having to assume forums with strand nature exist, others can arise naturally. That is we can replace a strand in this model with another forum. If a forum contains a group of n very coherent strands then it might be abstracted as a *n-token* and our dynamics might be reduced by treating it as single token; albeit one with more heavily weighted effect on the excitement on the graph.

We're not quite ready to discuss the stable forums that we will suggest represent social machines though. In the next chapter we'll see that in order to progress with this next big goal we must tweak the model to support partially overlapping forums.

Chapter 8

Higher order overlapping forums model

8.1 Model 127
8.2 Visualising overlapping forums of strands 129
8.3 Conclusion 130

In the last two chapters we've been working with a useful approximation to the overlapping forums of strands model introduced in section 5.1.1. To prepare for the argument in the following chapters, and to better support the modelling of real world systems, we now return to it. In the model advanced in the previous two chapters, strands are members of just one base forum and these forums are linked to form a network. Emergent forums may arise on this network. As it stands these emergent forums are something that might be inferred during analysis—but they are not something we can explicitly add into a model to constrain its behaviour. Such synthesis is ultimately useful not just for exploring ideas and generating data for testing analysis methods but for the analysis of real world data itself: many techniques for fitting models to data require an iterative approach in which models are tuned to minimise the difference between the data they produce and real data. The 'overlapping forums of strands' model introduced in section 5.1.1 allows us to explicitly add such higher order forums to our models.

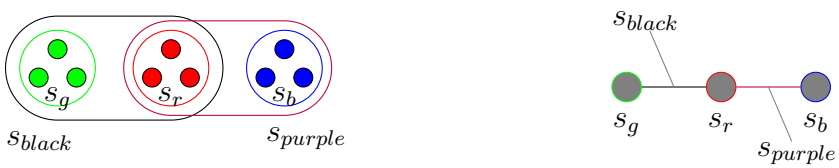
In section 8.1 we expand upon the 'overlapping forums of strands' model, and relate it to the 'network of forums' approximation to used in the previous two chapters. Then in section 8.2 we discuss an approach for visualising it and the more general class of forum dynamics that might mediate the interaction of strands on an experience graph.

8.1 Model

Some scenarios can be represented using both the 'overlapping forums of strands' model, introduced in section 5.1.1, and the 'network of forums' one we tested in the last chapter. Figure 8.1a shows three sets of strands in base forums s_{red}, s_{green} and s_{blue}, with s_{green} and s_{blue} connected to s_{red} through their membership in s_{purple} and s_{black}. We could model this using the networked forum techniques perfectly adequately, as shown in figure 8.1b. Notice that the first representation suggests that the forums of three strands are contained by the larger forums but that it could be redrawn to remove this suggestion; however the second representation explicitly indicates absolute containment—as do our networked forum models developed in chapters 6 and 7.

However, not all scenarios can be represented using the networked forums approach. Figure 8.2a shows an example that we can't represent using a network of forums unless breaking it down such that there is one strand per forum as shown in figure 8.2c.

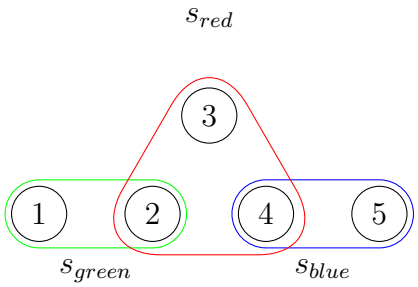
The representation in figure 8.2a can be viewed as a hypergraph. A hypergraph is a generalisation of graph where edges may join more than two vertices. We represent strands by vertices and forums by hyper edges. This form is useful as some methods used to work with standard graphs apply. Figure 8.2b shows this alternative representation in which



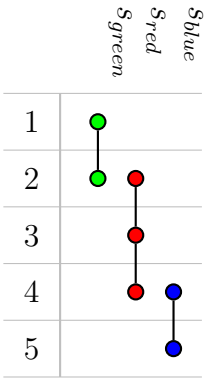
(a) Nested scenario

(b) Equivalent network

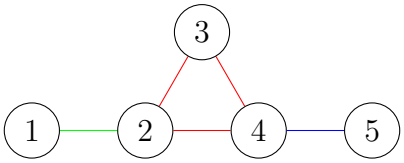
Figure 8.1: Scenario showing overlapping forums. Three sets of strands in base forums s_{red}, s_{green} and s_{blue}, with s_{green} and s_{blue} connected to s_{red} through their membership in s_{purple} and s_{black}.



(a) Partially overlapping forums representation as hypergraph



(b) Alternative hypergraph representation



(c) Expanded network representation

Figure 8.2: Three representations of a scenario with overlapping forums. Two forums are partially linked by a third that maintains potentially private state

forums are shown as vertical lines connecting strands. This is a useful view as it removes any hint of absolute containment suggested for example in figure 8.1a. The representation is explored in [91] and seems particularly good for representing hypergraphs that change over time.

To modify the networked forums model to support this new behaviour we first colour every strand with a unique value from a set of all strands I rather than with the colour of its single home forum taken from S. Now we must capture which forums that strand forms part of. We start by replacing the following network F with a hypergraph. For example, to encode the example of figure 8.2b we take our set of strands $I = \{i_1, i_2, i_3, i_4, i_5\}$ and identify an edge for each forum $S = \{s_{red}, s_{green}, s_{blue}\}$. Each edge represents a forum and is a collection of strands:

$$F^{hypergraph} = \{s_{red}, s_{green}, s_{blue}\} = \{\{i_1, i_2\}, \{i_2, i_3, i_4\}, \{i_4, i_5\}\}.$$

To complete the model we modify places that range over the set of forums S to range over

the set of strands I:

- X and X^e now range over $I \times X$ rather than $S \times X$,
- Λ^e now ranges over $I \times \Lambda$ rather than $S \times \Lambda$, and
- F now ranges over $I \times S$ rather than $S \times S$ and must be expanded out from $F^{hypergraph}$.
- T now ranges over $I \times S$ rather than $S \times S$,

In this model then the strands maintain their own level of excitement. The dynamics of the forum may still apply unchanged however—with the communally held levels of excitement obtained by summing across all strands in a forum for example. The explicit links between forums are now replaced by the strands that comprise both—but this is now not directional. The trust or excitement about links that was between forums is now between strands and forums. At any moment two forums will be more connected when a strand they share in common is exited about its connections to both simultaneously.

8.2 Visualising overlapping forums of strands

Recall that a coloured place can be expanded out to many uncoloured places, and that a place with a Cartesian product colour can be expanded out onto any of its components. In our models prior to this chapter a strand is represented by a token in the place X which was coloured with $S \times X$ where S was the set of forums and X the set of possible states on the experience graph. We therefore had a choice to visualise a fixed structure of forums frequented by strands whose colour changes as they move over the experience graph; or a fixed structure representing the experience graph roamed across by tokens coloured by forum.

In this overlapping forums model it is the strands that are uniquely colored. However when imaging them on the experience graph we may instead consider them as having the colour of each of the forums they comprise—and we may even imagine the intensity of their connection to those forums as brightness. (Although an expansion to Petri nets to support multi coloured tokens has been discussed in [22], we only need a trick to aid

visualisation.) The nodes of the experience graph can be visualised as being adorned by the decaying excitement of each forum to have recently entered them, and the labels on each edge by the excitement of each forum to have recently integrated the content they represent. Tokens sharing a colour with these adornments may be influenced by them: staying in a mutually excited state until the excitement of the content on an edge pulls them forward.

We now imagine a web the colour of each forum draped across the experience graph held at its vertices by the strands that are connected to that forum. In fact, with the idealised version of the forum we've defined throughout, these are not webs that might be viewed as supporting peer-to-peer connections between strands but single hypergraph edges indicating that all vertices are connected indistinguishably. Although we've described some particular dynamics for governing the movement of these tokens we can get far by imagining just a tension to these webs pulling the tokens together against the constraints of the experience graph.

8.3 Conclusion

We've returned to overlapping forum model by allowing strands to again connect with more than one forum. It is now applicable for models in which forums overlap through the participation of common strands. We have identified the places required to capture the state in a Petri net implementation but have not fully specified this net. Unfortunately Petri net representations of our models from this point onward are too large to execute using our software. We instead compare the model with reality in section 11.3, note that it is based the ontology for which we have identified corresponding example phenomena, and will use it to represent and reason about scenarios in the coming chapters.

Chapter 9

Persistence through ongoing binding

9.1 Ontology and example application 133
9.1.1 Peas galaxy example 134
9.2 Stable solutions in a dispersive environment 134
9.3 Model: Explicitly purposeful forums 135
9.3.1 Experiencing purpose 136
9.3.2 Purpose filtered forums 137
9.4 Ecosystem boundary scenario 138
9.5 Purpose as recurrent component of composite experience . . 140
9.6 Model: binding to source of motivation 142
9.7 Example: normative commitment as binding 143
9.8 Conclusion 144

In work so far we have developed a model of a self-similar forum. As described in chapter 8 a strand can participate in more than one forum. There is no explicitly defined network of forums; instead forums are joined by virtue of sharing common strands.

Recall the assumptions about the nature of the social machine formalism we seek, given in chapter 1. Our forum model now largely supports these, and we'll shift focus back onto them in chapter 11 when the model is complete. While developing the forum model we needed forums to emerge naturally and chose a solution in which a strand's connection to a forum will decay if not used. If we are to use this forum model to describe a social machine then this implies a defining characteristic of the social machine:

Social machine characteristic: A social machine exists over the period in which it maintains cohesion through an ongoing process of binding—at least in a dispersive environment such as an ecosystem competing for attention.

We must ensure that our forum model can meet this requirement.

We ask what causes a forum to persist. We start in section 9.1 by adding the categories 'binding' and 'ongoing binding' to our ontology and by applying these to the ongoing peas galaxy example started in section 5.2.2. In section 9.2 we reiterate why an ongoing binding process is necessary and very loosely defining what we mean by a dispersive environment; before then describing a few different 'strategies' supported by forums in our model to maintain this binding. Then in section 9.3, picking up from the social machine literature, we focus on advancing a shared purpose as a binding mechanism. We create a new *filtered forum* that supports the sharing of only experience that advances a certain purpose. We use this in section 9.4 to show that differing purpose across forums may result in a porous but stable boundary between them—something necessary in an ecosystem of forums. In section 9.5 we show that a filtered forum can be expanded out to produce a model where the filtering occurs in the members of an unfiltered forum. We do this by mixing the purpose of the forum into each of the experiences shared in it. In this way, while specialists in a forum may not be able to share each other's full experiences they can still share each other's excitement about having advanced a shared purpose. In section 9.6 we add—or in a sense, uncover—a mechanism to our forum model allowing all strands that form part of a composite experience to bind to the forums that motivated that experience; even if they previously were unaware of those forums. Finally in section 9.7 we show that the model may represent the phenomena of normative commitment.

Simplification of model. Recall that, although a strand's level of interest in a forum may vary based on interaction, the model lacks the ability for strands to connect or disconnect from forums. Instead we gave strands a permanent bond to certain forums and always allowed basic experience across these bonds while requiring excitement in an edge to share more privileged experience. In the remainder of this work we use a forum model

in which there are no permanent connections and no privileged experience and in which every edge must be excited if it is to share influence. We add a mechanism for strands to discover and connect to forums as well as to disconnect from them.

9.1 Ontology and example application

In keeping with the initial definition of a *forum* introduced in section 5.1.1, that "we infer a forum where a collection of strands' past experiences allows them to form experiences together", we introduce the category of binding and give some examples. We take the active sounding term binding and define it in terms of the definition of a forum to result in something more passive:

The *binding* of a strand to a forum occurs when an experience allows a strand to start forming experiences with a collection of other strands that together constitute that forum.

We will shortly introduce the idea of a *dispersive environment* as one in which forums of strands tend to disperse without some ongoing binding process. Where a strand has paths of experience that might cause it to stop forming experiences with the other strands in the forum, we notice that some scenarios exhibit ongoing binding:

The *ongoing binding* of a strand to a forum occurs when something about that strand, or the medium hosting it, maintains its ability to form experiences with the other strands in the forum: for example, a common purpose that is better advanced together, or experiences in a niche that can't be found elsewhere in strands' neighbourhoods.

Although we will hone in on purpose as a mechanism of ongoing binding there are more physical sources. The bind of a participant's strand to a forum may exist in something as simple as their habit of typing in a url to a browser at a certain time of day; or it may exist in part as a link to that url in their browser's frequently opened sites tab.

9.1.1 Peas galaxy example

The green peas scenario outlined in section 5.2.2, offers some examples of binding. After a classifier creates a new conversation strand $D_{\text{give-peas-a-chance}}$, some P_{browsing} strands visit this strand and are able to form experiences by interacting with it. At a trivial level these strands are bound to the technical strand itself to form a private forum. The browsing strands, and of course the people that host these, however may feel bound at a social level to the classifiers that commented into this thread and are able to interact with each other through it. These strands, in forming a connection to this forum, have become bound to it.

The green peas scenario also offers some examples of ongoing binding. After P_{browsing} strands initially bind to the forum that forms around the conversation strand $D_{\text{give-peas-a-chance}}$, their hosts may form a new $P_{\text{classifying-green-peas}}$ strand that remains part of this forum after their P_{browsing} strand looses interest. If their $P_{\text{classifying-all}}$ strand comes across something that could be classified as a green peas galaxy then this $P_{\text{classifying-green-peas}}$ strand may activate and share this into this forum. Now the people that host these $P_{\text{classifying-green-peas}}$ strands of activity may forget about the forum or become more interested in other activities. We discuss the normative binding that might hold them together in section 9.7.

9.2 Stable solutions in a dispersive environment

Let's look at a real-world behaviour that might contribute to the decay of the connections in our model. Assume that there are forces that cause inactive forums to decay. For example, in an ecosystem competing for attention strands' interests might be swayed toward more engaging forums and the sense of shared history and shared state of mind carried by other strands in the forum may be lost. Where this is the case a forum needs some ongoing process of binding to hold it together against these forces of decay and dispersal. The shape of the experience graph frequented by a forum's strands and those of its neighbours is key to this process—as is the structure of the ecosystem supporting a forum.

Such a binding mechanism requires an ongoing state of shared excitement about something. This might occur as a forum uncovers a single unfolding story, or keeps revisiting the same collection or ensemble of retold stories, or keeps reusing the same story fragments or other simpler components of experience. A single unfolding story might be that of the world's unfolding events observed objectively or as told by a particular authority and then—either way—shaped by the forum itself. A revisited story might be the Nativity or a fragment of an unfolding story that is often told. Simpler components of experience might include the building of known concepts or of emotional affect.

Within a forum a collection of strands may operate across all these parts of the experience graph together. In fact many of the simpler shorter lived story fragments or components of experience will be used across forums maintained by more specialised stories or experiences or across more unique combinations of connections to external sources.

In the next sections we'll jump into the idea that the sense of advancing a shared purpose is a strong binding process. Web scientists picked up on this process because it is clearly very evident and important. Perhaps we will want to define our social machine abstraction around purpose, but we need our forum model to be as generic as possible. Features other than purpose may carry through a series of experiences. For example we might share a common emotional response to some content; and, as it can feel good to keep sharing this, it may act to guide our future experiences together thus binding us. Sometimes just being in motion with others is enough to hold us together even if we don't perceive the cause of the motion. We can also share appreciation for something we mutually feel to be aesthetically pleasing.

9.3 Model: Explicitly purposeful forums

Longer term binding might be associated with a region of the experience graph that once entered will be hard to leave. One way to do this is to pick up a sense of shared purpose.

9.3.1 Experiencing purpose

We can reason about purpose by finding a place for it on an experience generator graph G. We start with the first order generator of equation (4.1) $G^1_{ts} = (X, \Lambda, E)$. Recall that E is a set of edges between states in X each labelled with an element of shareable content from Λ. Now if each edge in E advanced an independent purpose then elements of E could be used to define purpose. However we'd like multiple edges to advance the same purpose, and for an edge to potentially support multiple purpose. To represent this we create an alphabet of purposes Φ and use these to further label the edges of the labelled transition system.

There is a structure to purposes in that advancing some advances others. For example by finding and labeling images of big dogs a social machine participant is implicitly advancing the purposes "*label big dogs*" and "*label dogs*". We represent this structure as a directed graph Ω with nodes taken from Φ and edges from more specific purposes to more general. For example "*label big dogs*" might advance "*label dogs*" which might advance "*label domesticated animals*". We call this set of directed edges Δ, resulting in the directed graph

$$\Omega = (\Phi, \Delta)$$

There are many ways to describe the same purpose; we should think of each element of Φ as a set of these. Finally it will be useful to define a universal purpose $\top$ that is advanced by an edge labelled with any purpose or by no purpose.

We can now define a first order generator with purpose,

$$G^1_{purpose} = (G^1_{ts}, \Omega, T)$$

where T is a set of tuples (ϕ, e) in which $\phi \in \Phi$ and $e \in E$. Each tuple labels an edge of the experience graph with a purpose situated in the purpose structure Ω. Not all edges need be labelled but to keep the representation tight these are implicitly labelled with the universal purpose $\top$.

As an aside for now, we may consider sharing a state of mind as being itself a purpose—

and this whole work is about sharing states of mind. Some states of mind are conceptual. A concept may have many words, or synonyms, that communicate it [53]. We might encode these synonyms as experiences in G and the concept that connects them a purpose from Φ. We can go further though. We may configure G so that integrating content that causes the experience of one of a concept's synonyms also raises that concept into a collective's awareness. We could add a composite experience of both the synonym and the concept for each synonym; and ensure that there is a strand in the collective able to experience the concept at that time.

9.3.2 Purpose filtered forums

We introduce a *filtered forum* as a forum that allows the sharing of experiences that advance only a certain purpose. Recall from chapter 8 that a model includes: a set of strands I, a set of forums S and a hypergraph with an edge for each element of S connecting vertices from I. We simply add a function $\phi(s)$ that assigns a purpose taken from Φ for every forum s in S. Equivalently the filtered forum can be seen to function over a subset of the directed experience graph. A filtered forum s will only support the sharing of experiences that further a purpose in ϕ_s or a purpose that is more specific than this purpose.

In the models defined so far an experience is shared not by the edge traversed but by the content that causes it and the state of mind reached. Now two strands may be in different states and so while a certain piece of content might move one along an edge labelled with a given purpose, for another strand it might not. Therefore a strand can't know in our simple models how others will react to content. We define the filtered forum such that they can: i) a strand will share into a filtered forum excitement about the content that moved it only if that movement advances the purpose of the filtered forum; and ii) a strand will respond to excitement shared via this forum only if it advances that purpose for *itself*. This symmetry will make our later expansion simpler. It also implies that a strand is aware of the purpose of the forum it receives content over and that this is part of the content.

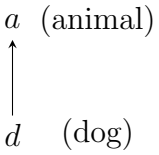


Figure 9.1: Graph of purpose for ecosystem boundary scenario. Purpose d advances purpose a. $\Omega = (\Phi, \Delta) = (\{a, d\}, \{(d, a)\})$

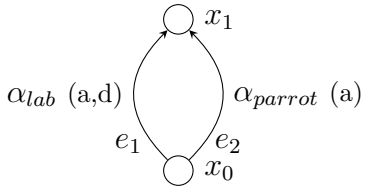


Figure 9.2: Experience graph for ecosystem boundary scenario. The graph is labelled with content and with purpose in brackets. $G_{purpose} = (G, \Omega, T)$, where $G = \{X, \Lambda, E\}$, with $X = \{x_o, x_1\}$, $\Lambda = \{\alpha_{parrot}, \alpha_{lab}\}$ and $E = \{e_p, e_l\}$, where $e_p = (x_0, \alpha_{parrot}, x_1)$ and $e_l = (x_0, \alpha_{lab}, x_1)$, and $T = \{(a, e_1), (d, e_1), (a, e_2)\}$.

9.4 Ecosystem boundary scenario

We apply this new kind of forum model to a simple scenario and then will show how regular unfiltered forums can be used to achieve the same effect. An important behaviour in social machine ecosystems is the sharing of experience between machines via participants that form parts of both. We explore this scenario.

We first define a scenario using an experience graph with purpose and filtered forums. Figure 9.1 shows a purpose graph where there is a purpose a to label, share or discuss content that includes animals and a purpose d to do the same for dogs. Advancing purpose d will also advance purpose a but not necessarily the other way around.

Figure 9.2 shows a simple experience graph where strands in state x_0 may experience either content capturing a parrot or Labrador dog. This moves them to the same state of mind x_1 which is simplistic but good enough for our purposes—perhaps the state of having just wasted a little time. The bracketed terms after the content indicates the purposes edges are labelled with.

Figure 9.3 shows the spatial aspect of the scenario. It includes four strands, $I = \{i_1, i_2, i_{2'} i_3\}$. There are two filtered forums m_{dog} that is interested in dogs and m_{animal} that is focused on animals. These are linked by the strands i_2 and $i_{2'}$. These two are

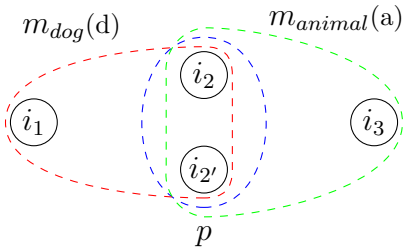


Figure 9.3: Overlapping structure for ecosystem boundary scenario. Membership of strands $F^{hypergraph} = \{m_{dog}, p, m_{animal}\}$, where $m_{dog} = \{i_1\}$, $p = \{i_2, i_{2'}\}$ and, $m_{animal} = \{i_3\}$, with $\phi(m_{dog}) = d$, $\phi(m_{animal}) = a$ and $\phi(p) = \top$.

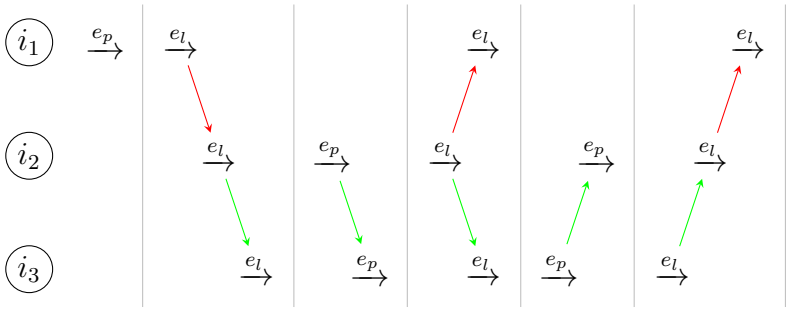


Figure 9.4: Execution traces for ecosystem boundary scenario. Six execution traces for each possible initial experience. Each row represents a strand and each column a trace. Entries are the name of the edge traversed and lines connecting these are coloured to indicate which forum the experience was shared through.

members of the unfiltered forum p that represents a person or participant. We included the strand $i_{2'}$ only to get across the idea that the person has other parts to their lives—we assume this is on parts of the experience graph that share no content with the parts of the graph frequented by our strands.

The three strands involved in the scenario start in state x_0 of the experience graph in figure 9.2. Suppose that one of these experiences content of either a parrot or a dog by traversing edge e_p or e_d. There are six possible such events. Figure 9.4 shows an execution trace for each. Notice for example that when i_1 experiences content relating to a parrot this is not shared through m_{dog}. The experience of dog related content will flow both ways between the filtered forums but non-dog animal related content will not.

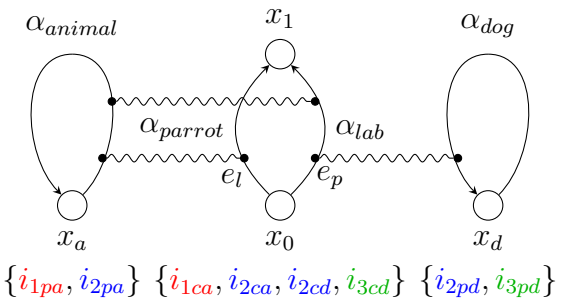


Figure 9.5: Experience graph for ecosystem boundary scenario. Including three composite experiences and two first order experiences.

9.5 Purpose as recurrent component of composite experience

We have just defined a new mechanism for representing purpose on the experience graph and introduced filtered forums. We now ask if scenarios that can be described using these techniques can also be described without this new kind of forum. Any approach must move the filtering into the participants. One solution is to make every experience shared into or from a forum part of a composite experience comprising both the experience of the content itself and that of advancing the forum's purpose. We present this approach using the scenario described above.

Figure 9.5 shows a possible experience graph for our running example. Strands in a focused forum must experience content mixed with the knowledge that it advances a purpose. Composing these experiences requires two strands. Strands in the state x_0 can integrate α_{parrot} or α_{lab} as before. However they need the help of a strand focused on advancing a purpose. Strands in the states x_a and x_d are in positions to do this. The need for this cooperation is indicated by the squiggly horizontal lines connecting edges that must be traversed together; as introduced in section 4.4. a strand in x_a with a focus on animals can help a strand in x_0 experience either dog or animal-related content. A strand in x_d however can only help such a strand integrate dog-related content.

Figure 9.6 shows one way to recreate the spatial aspect of this scenario. Again two forums m_{dog} and m_{animal} are linked by a common participant p_2. This common participant now contains four strands relevant to the scenario rather than one. The pair i_{2c_a} and i_{2p_a}

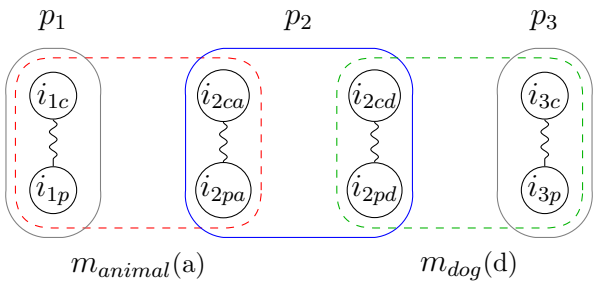


Figure 9.6: Spatial view of hypergraph for ecosystem boundary scenario.

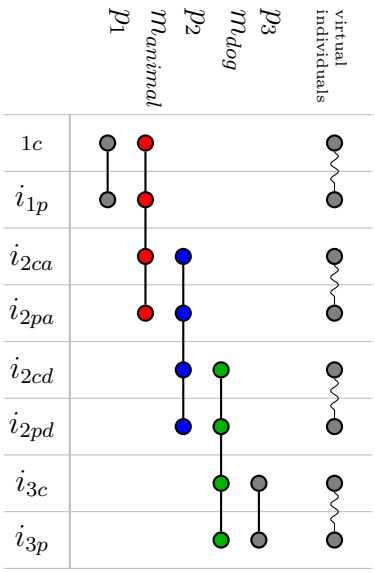


Figure 9.7: Connectivity of hypergraph for ecosystem boundary scenario.

together traverse the content and purpose of animal related experiences; while the pair i_{2c_d} and i_{2p_d} focus only on dogs. The participants p_1 and p_3 have been expanded in the same way.

Figure 9.7 shows an alternative representation of the connectivity of the strands. Rows represent strands, or vertices on the hypergraph, and each column a forum, or edge on the hypergraph. The rightmost column identifies what we have been calling virtual strands. These are collections of strands in the position to contribute to a composite experience. Virtual strands may come and go as strands move on the experience graph—in this example each will dissolve when a component leaves state x_0.

9.6 Model: binding to source of motivation

When a pair of strands from the example in the last section work together to form a composite experience they share their excitement about both the specific content of the experience and the forum's purpose. Participants, formed by collections of strands, in a forum should be bound while working on the forum's purpose even if they don't share the same experiences—for example participants with differing specialisms might form a team and be bound to that team by each other's excitement in their own progress even if the details of this work is not experienced directly. In the models prior to chapter 7 the forum was the primary building block and might be used to model a participant; and with all strands in a forum sharing the same connections to other forums this is exactly what would happen. In this chapter strands are primary and we don't want every strand in a forum to be bound more tightly to another forum when just one strand followed that forum. Instead we complicate the dynamics of chapter 8 to handle composite experiences slightly differently.

Suppose that a specialist was driven by a forum's ongoing excitement about some purpose to favour a composite experience that advanced that purpose. We want all the strands that participated in composing that experience to be bound more tightly to the forum that excited them into doing so. If they are not members of that forum then they should be added. Thus we add a new requirement to the dynamics described at the end of chapter 8:

> **Modelling mechanism: Binding to source of motivation.** A collection of strands within a forum may form a composite experience as described in section 4.4. Where a forum has shared content that excited some of the strands to contribute to this experience then *all* of the strands should bind more strongly to this forum. Where they have no connection to this forum, one should be made.

Figure 9.8 shows a sketch of this new mechanism. The blue lines indicate parts of the experience graph. Four strands are moving on this graph. Two virtual strands indicated by the curly lines form during separate experiences. Both virtual strands include a strand on

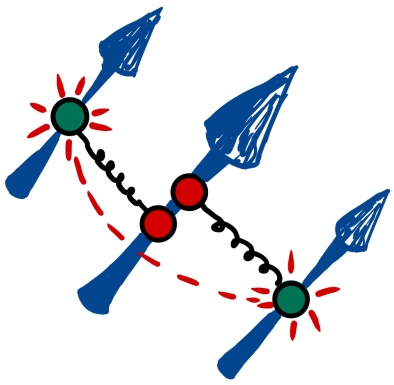


Figure 9.8: A sketch showing a method for representing the binding of strands to their source of motivation when this source is a forum unknown to them. Here the blue arrows indicate paths on the experience graph; disks represent strands coloured by forum membership; and the curly lines indicate virtual strands. The overlap of these virtual strands with the central part of the experience graph indicates common motivation causing the two green strands to join the red forum—indicated by the comic appearance lines. Following this experience the two green strands will have red stripes; as seen in figure 10.1.

the central part of the experience graph—which might represent purpose. The red strands form part of the same hypergraph edge. The forum this edge represents in the source of motivation in that it will support the mutual excitement of the two red components of the virtual strands. As such the other strands join the red forum—as indicated by the comic style appearance lines. The red dotted line indicates part of the red hypergraph edge and serves just to highlight this new connection. Another way to look at this is that the emergent strand that forms during the experience is a short lived forum—and that all the strands comprising this forum are bound to the forums that drove the experience.

This mechanism supports the well studied social network formation mechanism, described in section 3.2.3, in which people with a common friend will tend to become friends.

9.7 Example: normative commitment as binding

There is a community that applies science to develop guidelines and best practices for building successful online communities. Some of this community's work is summarised by Robert Kraut et al [46]. They uses the concept of a *design claim* as a basis growing this knowledge:

> We use the device of design claims to translate theory to design alternatives that achieve community goals. Design claims follow a positivist scientific paradigm, seeking to state general claims—that under certain observable conditions certain outcomes can be expected. In our case, the conditions that are of particular interest are design alternatives and the outcomes are desirable features of an online community, which we refer to as "design goals."

The authors describe a normative commitment as "a feeling that one has obligations to the community to be loyal and act on its behaviour". They identify three factors in creating normative commitment: a commitment to the purpose or cause, others' normative commitment, and reciprocity. Given a collection of people with strands of activity driven by compatible purpose, in section 9.5 we developed a model in which, when a person advances that purpose they share the experience that advanced it alongside the experience of advancing the purpose itself. This acts to highlight the community's purpose and a member's success at achieving the purpose. Kraut et al. claim that: "Highlighting a community's purpose and successes at achieving that purpose can translate members' commitment to the purpose into normative commitment to the community."[46, Design claim 27]. In our theoretical model, strands are bound to the forum when sharing in either of these experiences and we may therefore interpret this binding an representing normative commitment.

9.8 Conclusion

We started this chapter by observing that forums in our model require an ongoing process of binding to sustain. In turn we realise that if our model represents our theory closely then a social machine also needs an ongoing process of binding to sustain. We described a few different 'strategies' that forums might use to achieve this. Participants may have different experiences but if just some component of those experience is common then they may become excited by this and this may cause an ongoing process of new experiences and subsequent binding. Picking up from the social machine literature, we focused on advancing

a shared purpose as a binding mechanism. We extended the binding mechanism so that when a few strands contribute to a composite experience they will bind to the forums that motivated each other even if not previously aware of them. This is important as it provides a way for strands to join new forums and in turn for forums to start overlapping with new forums. Acting as an example scenario we showed that a stable boundary may form between two forums where one advances a more specific purpose than the other. As discussed in section 8.3 we could not create a concrete Petri net implementation of this model.

Chapter 10

Life cycle of persisting forums

10.1 Ontology 148
10.2 Spawning new forums 148
10.3 Coalescence of persisting forums 151
10.4 Future: long and short term binding 152
10.5 Conclusion 152

We've just uncovered and explored a new requirement of our forum model and are nearly able to say that this model supports an implementation of a social machine: that of a forum persisting through an ongoing state of binding. In this chapter we consider the life cycle of forums. We start by proposing three new characteristics of a social machine:

- Social machines must support the emergence of new social machines. We will say that a social machine must act as a *platform* for new social machines; the ongoing presence of binding is not absolute and allows a new process of binding to form around a different feature of the experience graph.
- A social machine, acting in its capacity as a platform, should support the coalescence of social machines that share compatible binding processes and connections to their environment.
- Social machines should support the emergence of ecosystems in which entirely separate

machines that have compatible binding processes will tend to have components in a common machine—and indeed we expect ecosystems where this is true to out compete those where it is not.

We will discuss these here and of course we need to ensure in turn that our forum modelling methods can support structures with these characteristics. To this end, in section 10.1, we add a new characteristic to our ontology's 'forum' category: that a forum has a dual nature as both a collection of strands that may interact according to the dynamics we've described, and as a crucible allowing its component strands to form, possibly overlapping, more localised forums. Then in section 10.2 we add behaviours to our model that result in this characteristic. Particularly, we ask how a new forum might arise within a forum that maintains an ongoing state of binding. One reason for this to occur is to allow specialist forums to work together alongside a more general host. We notice that our forum model must in turn have a dual nature: it is a simplistic forum of the nature we've been discussing with a single channel suitable for sharing experience, but must also support peer-to-peer connections needed to situate strands such that they may interact locally. We propose a natural way for this to arise and observe that this happens in many social machine frameworks: a composite experience creates a new forum that may sustain if a situation supports it.

Then in section 10.3 we ask if our model supports the coalescence of forums that share compatible purpose and connections to their environment. We identify a positive feedback loop that will encourage the growth of a single forum at least when strands share a connection to a common forum. We then reason that there is a network effect in which regions of a broader ecosystem that provide a structure that makes it easier for strands and forums of strands with a common interest to find each other through a coupled interest in some broader forum will out compete those regions that don't.

10.1 Ontology

To allow the creation of new forums we add a new characteristic to the ontology's forum category:

Dual nature of a forum. A forum has a dual nature, as both a collection of strands that may interact according to the dynamics we've characterised, and as a crucible in which these strands may also form more localised forums. These more localised forums are not absolutely contained by the forum that brought their component strands together. As such, they may overlap with each other, and also with the forum that initially brought them together, or with any other strands their forums overlap with.

In real-world application, these more localised forums may act as if they are contained by the forum that brought their strands together, but this will be a feature of both the resulting structure of overlapping forums and of the paths of experience constraining the strands these comprise. When we say there is no absolute containment, this indicates that the ontology itself, has no explicit place to capture the fact that one forum brought together the strands that comprise another, or that one forum is constrained to comprise strands that all comprise another.

10.2 Spawning new forums

Suppose that a collection of strands are working away in a forum to advance a common purpose. How might they form a new forum with a new purpose and with new channels for communication?

To answer this let's ask what might drive it to happen. For example, why might a forum in a state of binding around a certain purpose need to specialise? In the abstract, let's suppose that traversing two fairly distinct regions of the experience graph will both advance a host forum's purpose. Now suppose that both regions include composite experiences that include contributions from strands on a common region of the graph—other than

that of this shared purpose—and that they therefore interfere with each other's progress. There would be an advantage to having a separate forum for strands on each of these regions. For example a team or community of scientists might be advancing a common goal. Suppose that they have two specialisms and thus the shared experiences of each take part on different regions of the graph. However although their corpuses of knowledge overlap there are regions that don't and regions that would cause confusion if mixed; thus they find it hard to work together in the same forums.

Now, this thesis is based on the premise that we find it appealing to share experience. The dynamics of our forum models favour the movement of strands along the same path or along paths in which they work together to form more complex composite experiences. If two or more strands come to be moving together in a forum they will tend to bind themselves (and those that helped form any composite experiences) more to this forum. If this forum is large or busy however the chance encounter in which they found themselves travelling a path off the main route will remain just that. We are looking for a natural mechanism in which these fellow travellers might form a new forum in which to share experiences in a quieter and more focused way.

The first issue is that in our models the strands in these forums are anonymous and there is no way for them to know of each other by location. Recall though that in developing our self-similar forum models described before chapter 8 we started with base forums comprising anonymous strands with no way to interact other than through communal channels of excitement. We allowed these forums identity in the form of a location on a network and showed that if they bound to each other when sharing experience that a densely connected region can form and that this acts itself as a forum.

We have identified a new requirement of our generic forum model. The anonymity in this forum must not be absolute if new forums are to arise between strands:

Modelling mechanism: Dual nature of forum. A forum is actually two things: a simplistic forum of the nature we've been discussing with a single channel suitable for sharing experience and also a network of potential peer-to-peer connections.

A simple approach to allow peers in the forum to connect directly is just to allow every

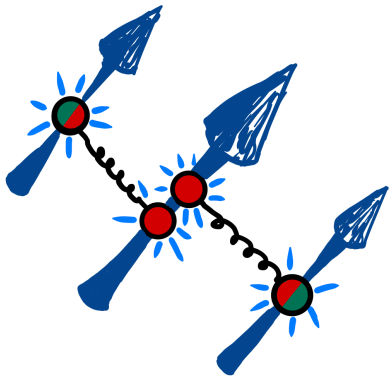


Figure 10.1: A sketch showing a method for representing the formation of a new forum in response to strands forming a composite experience. This example carries on from figure 9.8. The strands are all now in the red forum. They form a composite experience together this time; and in doing so create the new green forum as indicated by the comic appearance lines on all strands.

experience to create a new forum, and then to depend on the mechanisms we have so far to allow the chance for this to sustain. This requires another new behaviour of our model:

Modelling mechanism: Composite experiences create new forums.

Figure 10.1 shows a sketch of this new mechanism. The strands from the previous sketch form a composite experience together and in doing so create a new forum shown in green. This is behaviour consistent with the operation of the social machine frameworks used by all of the well known social machine frameworks. We'll look at frameworks that encourage more personal reaction to content that has been shared into large and otherwise anonymous forums in section 11.3.3.

These new forums may exist after the experience that they formed alongside has completed and they may then act as a forum for strands to continue to meet in. In this way once they've met to form one composite experience the strands may, if they have more chance to interact on their own private graphs, continue to carry this forum with them. When they have no way to interact further or if their attention is spread elsewhere then the forum will dissipate.

A big factor in whether the new forum may sustain is the nature of the forum that sparked it off. If the new forum's purpose is compatible or orthogonal to those within its ecosystem then it more likely to persist than if it opposes it. For example a conversation about pets might sustain in a left leaning echo chamber but a conversation that would

sustain in a right leaning echo chamber might not; or a new forum working on a more focused purpose than its host might sustain because its participants will be binding to both forums.

10.3 Coalescence of persisting forums

We've just described how a forum might come to spawn a forum that persists. To form the self-similar structure of the web described in section 3.2.1 forums that share compatible purpose and connections to their environment should tend to coalesce. Does the model we have support this?

Suppose that two forums are binding around the same purpose. We have a mechanism for strands to join a forum when they form a composite experience with one of its members. So there is a mechanism for strands to end up bound to both of these forums. Where a strand is then pulled by the excitement in both of these forums the binding mechanism will increase the strength of its tie to each forum at a rate proportional to how strongly each influenced it. Thus there is a positive feedback loop in which these strands will tend to be bound more to the larger forum and thus be influenced by it more.

However, to coalesce within a broader ecosystem, the strands comprising separate forums sustained by the same purpose must have a chance to meet. For this to occur requires that strands from both exist in the same forum and interact through this. For example two separate forums of people interested in dogs might come together through a shared interest in animals; in that the community of people interested in animals is larger. Thus our model supports a network effect (see section 3.2.3) in which regions of a broader ecosystem that provide a structure that makes it easier for strands and forums of strands with a common interest to find each other through a coupled interest in some broader forum will out compete those regions that don't.

10.4 Future: long and short term binding

Each time a strand experiences something from a forum it increases its interest in that forum. In our models this excitement decays in discrete steps over time. These decay events form a Poisson process, which, as described in section 3.3.2, is memoryless in that the chance of an event occurring is unrelated to the timing of past events.

We can imagine any number of more realistic models. For example an excitement model might capture both the chance of being influenced by a forum with which we have a long term ongoing relationship and one that we just passed through. We might notice for example that we maintain different average rates of connection with parts of our family and friend groups, but that once engaged with them the time between interactions will not change this significantly.

In the short term a strand may shift their attention towards and away from a forum. In the longer term, like a plate kept spinning, these short term interactions may maintain longer-term connections. When connecting to a new forum a strand need not necessarily form a long term connection. For example, if, while advancing the purpose of a forum it has a long term relationship with, a strand binds to the source of a peer's experience then this need not leave a long-term tie. Perhaps if this edge is visited a few times in succession then a longer term saturating connection might form.

Of course to stick with our theme of self-similarity we should expect many such time scales. I feel that providing a way to model just two should be sufficient for these to arise. The short term binding to an external forum of an ensemble of strands in a forum, for example, might result in, from the perspective in which that forum acts as a strand in another forum, a long term binding to that external forum.

10.5 Conclusion

Our theory needed to address the life cycle of social machines. We proposed three new theoretical properties: that a social machine may act as a platform to start new social machines and to aid the coalescence of social machines that share compatible binding

processes; and that successful social machines will form part of an ecosystem in which separate social machines that could coalesce will tend to have components in a common platform.

We updated our model with one final behaviour to support the first two of these properties: every composite experience creates a new forum. This may sustain if it gains an ongoing binding process. We observed that sustaining forums will coalesce naturally in the model due to a positive feedback loop that forms. We suggest that natural selection in which ecosystems that bring together forums with compatible binding processes will out compete those that don't. As discussed in section 8.3 we could not create a concrete Petri net implementation of this model. We will compare the resulting model with reality in section 11.3.

Chapter 11

Theory and model of social machines

11.1 Theory of social machines 155
11.2 Idealised social machine model 156
11.2.1 Model operation 156
11.2.2 Topology of social machines 158
11.2.3 Executable models 159
11.3 Model agreement 160
11.3.1 Channel for sharing excitement about content in a forum . . . 160
11.3.2 Channel for sharing each state of mind in a forum 161
11.3.3 Private forum for reacting to content 162
11.4 Model limitations 163
11.4.1 Social machines may sustain a flow of participants 164
11.5 Model implications 165
11.5.1 Social machine ecosystem abstraction from medium 165
11.5.2 A social machine needs a habitat 166
11.5.3 Archives and strand state are inseparable 166
11.6 Conclusion 166

We started this work by asking if the ecosystem of overlapping communities on the web can be understood as many instances of the same building block. We proposed that the resulting ecosystem should include social machines. We then started our theory of social machines by adding the characteristic that they should be self-similar. To explore the feasibility of this and to understand the repercussions we developed a simple model of this

ecosystem hoping that something with these characteristics would arise. Our conception of an idealised social machine, as an emergent structure with these characteristics grew while doing this. In some cases the behaviours required of these structures seemed inescapable and in these cases we added them to the characteristics required of a social machine in our theory.

We present the resulting theory in section 11.1 and model in section 11.2. We then ask in section 11.3 how reasonable our idealised model of a social machine is by comparing and contrasting it with examples from the real word. In section 11.4 we identify some limitations of the model. Finally, in section 11.5 we discussion some implications of the theory on how we understand phenomena to which it applies.

11.1 Theory of social machines

Our theory of social machines was developed for modelling ecosystems of overlapping communities. As such our conception of a social machine is not just something that depends on the social nature of its participants, but something that is social in itself. We propose that a social machine:

1. depends on and supports its participants' nature to form shared experience;
2. exists only while maintaining cohesion through an ongoing process of binding that occurs while participants form shared experience;
3. is self-similar: the social parts that comprise a social machine may themselves be social machines and a collection of social machines might comprise a social machine;
4. must, at least in principle, support a single overarching linear sequence of experiences, despite being inherently concurrent—it must exhibit some personal ordering of events;
5. must act as a platform supporting the emergence of new social machines;
6. should, acting in its capacity as a platform, support the coalescence of social machines that share compatible binding processes and connections to their environment; and

7. should support the emergence of ecosystems that encourage the coalescence of social machines.

Some repercussions of this theory uncovered in the sections below include:

1. Despite being self-similar a social machine cannot absolutely contain another social machine.

2. A social machine does not include the archives it might create during its operation.

A formal model of an abstract social machine:

1. should be written in part in some formal representation of experience;

2. might represent phenomena in any media and comprising any form of participant to which it applies.

11.2 Idealised social machine model

Here we summarise the model developed to represent this theory and discuss the topology of ecosystem the theory implies. We then reiterate the implications that the concrete implementation of this model is executable.

11.2.1 Model operation

It is useful to revisit the forum model with the details of its Petri net implementation abstracted away. The model must support the most fundamental aspect of our theory: that social machines are self-similar in that their participants are social machines and that they may participate in social machines.

We start by assuming that there are objects with some permanence that may undergo a linear thread of experiences all of which effect the objects' common state effecting future experience. We then model every path of experience that these strands may undergo. We assume that experiences form only in reaction to pieces of external content. Integrating this content changes a strand's state in a way that effects how future pieces of content

will be integrated, and even if a piece of content may be integrated at all. Some states might represent a model of reality, and the integration of content the folding of new facts of experience into this.

These strands exist always in a forum. A forum supports a single channel for communicating excitement about content. This channel shares the forum's current level of excitement about each piece of available content. This is simply the average of each strand's excitement about that content. Strands will try to integrate content at a rate proportional to how excited the forum is by it—but only if their current state of mind permits. Where strands' states of mind are similar enough that they can appreciate this they will hold themselves in this common state through communicating their excitement about being in it. To facilitate this a forum supports a separate channel for each state of mind a strand or strands find themselves in.

In this way, a forum supports multiple bundles of strands undergoing shared experiences by way of sharing their state of mind. A little like a dinner conversation among friends, a forum supports only a single channel for communicating content such as spoken words and phrases, but people who have entered synchronous states of mind are able to hold separate conversations over this channel.

While a collection of strands is following the same path of experience in a forum then that collection of strands may act as a strand in other forums. We propose that all strands—the ones we started this discussion with, and these new ones that emerge from those—are all the same class of object. That is all strands are actually collections of strands moving together. Stepping back, we did not model this with our Petri net approach. We used a crutch to get things going by creating eternally existing strands and demonstrating that these may form flocks that act like single strands that are slightly spread out on their path of experience.

Strands may—and we'll argue in section 11.2.2 that they must—comprise more than one forum. When a strand integrates content from a forum's content channel then it binds to that forum. Strands are influenced by all the forums they comprise. This influence is weighted according to the strength of a strand's tie to each of these forums. These

bonds decay with time and may break entirely if not used. Thus for a forum to persist the strands that comprise it must maintain an ongoing process of sharing experience with each other. Such a forum we propose is a social machine.

Strands comprising a social machine may follow different paths of experience yet sill be bound together. If just part of the content of these experiences is common then they may share excitement about this common part; and when they maintain an ongoing state of excitement about this part then this may result in the ongoing process of sharing experience required to bind them. For example strands may be bound by advancing a common purpose, maintaining a state of emotion, reusing a particular concept or story fragment, or reusing parts of a common corpus of knowledge or belief system.

It may take more than one strand working together to form an experience, with each forming their own simpler experience required to create the whole. In our Petri net models we called these composite experiences. When this occurs then all of the contributing strands involved will be bound to all of the forums that drove the simpler experiences. In this way strands may be 'introduced' to new forums that a partner in this experience was motivated by.

The process of forming such a composite experience results in a new forum comprising the strands that contributed to it. In many cases this will not be reused and when all binds to it have decayed completely then it may be dropped—and potentially archived outside the scope of our model but in a way that creates new possible paths of experience for strands to traverse. In some case though the new forum may sustain trough an ongoing process of binding and become a new social machine.

11.2.2 Topology of social machines

The self-similar definition of a forum that we began with is circular and raises immediate questions. Most immediately is that these parts can't keep getting smaller, and so what happens at the bottom; and that they can't keep getting more encompassing, and so what happens at the top? These questions are based on an assumption that there is an absolute hierarchy of containment—and as we can't answer those questions if this is the case then

we conclude just that the model must not support this assumption. This gives us a clue as to the nature required of these forums. The model will support no global containment if it supports no absolute local containment. Thus: a forum is comprised of other forums but it cannot be said to fully contain them—and the corollary, where a forum comprises other forums it cannot be said to be fully contained by them. For it to be true that a forum cannot be fully contained by another says that it must always comprise more than one forum—forums exist therefore at the overlap of forums.

This lack of containment is key to building a model which supports our theory that a social machine must support the emergence of new social machines. In this way a new social machine emerges in a way that it is not part of the social machine that it started in—but just another social machine. In some cases the new machine's environment will be such that it is destined to act as a component of what in this case acts as its host—for example as a specialist or worker—but this is not a function of the model itself.

Although one social machine cannot absolutely contain another, social machines may overlap and indeed we've just said that they must. In our idealised model they overlap by sharing some strands in common. From the perspective of a social machine a strand it shares with another social machine is simply part of itself. More particularly the strand is a part that expresses an aspect of itself that is compatible with an aspect of the machine it overlaps with. This mutual compatibility exists when the strand can maintain a single state of mind while interacting with other strands from both social machines in an ongoing way. Overlapping machines must therefor be different enough from each other that they will not coalesce, while being similar enough that they can share this part of their states of mind. A strand then represents a single process acting on a single state of mind that exists in two or more social machines concurrently.

11.2.3 Executable models

The approach we used to define scenarios and to then represent them as Petri nets is described in the text. The Petri net models are executable and this leads to something interesting.

If a Petri net implementation of the model in the text is executed some forums will persist; and if they do it will be because they support an ongoing process that binds their participants together through a tendency to bind to the source of shared experiences. Such a persisting forum would be a model of a social machine that might arise in the scenario. Given the broad scope of our theory of social machines we could in fact say that this persisting forum *is* a social machine.

In section 7.3 we showed that the idealised model supports the emergence of separate regions of culture. Each region is held together by an ongoing state of binding as its members chat about certain beliefs they share. However these structures were not technically social machines according to the theory as the most complete Petri net implementation of the model did not support features added in chapter 8 and later. This Petri net model lacks the behaviours of chapter 8 that allows forums to overlap rather than comprise an explicit network. This is not required by the theory for the emergent structure to be a social machine. However the behaviours added in chapters 9 and 10 are required by the theory in order that the resulting structures support an ecosystem of social machines: the Petri net did not bind all contributors to an experience to the sources of motivation, and nor did it create a new forum for each composite experience.

11.3 Model agreement

The idealised model shares fundamental behaviours with the interaction models used by the dominant social media companies identified in section 1.5 [57].

11.3.1 Channel for sharing excitement about content in a forum

Forums include a channel for sharing excitement about content. Recall that this presents the average level of excitement about content among strands in the forum. The stochastic decay functions we used result in a level of excitement about content that is proportional to the rate that it has been recently integrated; and a level of excitement about being in a state of mind that is proportional to the rate that strands have recently entered it.

This central channel forms a way of directing attention toward content that will in turn shift strands to certain states of mind. The behaviour this channel supports is evident in social situations in which people tend to shift each other's attention onto the same topic or part of their environment. It is a key element supported by the frameworks of dominant social media brands.

If we take it that Reddit's subreddits, Twitter's hashtags and Facebook's groups loosely equate with forums then it is true that they share this characteristic. Each presents a linear list of content; where this content may be posts or replies to posts. Our list would be ordered by the rate at which this content has been recently integrated—akin to the way many social media frameworks offer a view of 'trending' content. In fact, rather than a list our model might be seen to present this excitement about content in a way more akin to a tag cloud where more popular content is shown larger and more centrally in an image.

Regardless of presentation, it is clear that these exemplars all have methods to draw participants toward the content that is most popular in the instant—and in turn have some way for participants to indicate either explicitly or implicitly how interested they are in content. Of course the algorithms used by the companies managing these systems are more complex than those of our idealised model. Where a platform has an agenda perhaps we might model this as an idealised forum with purely democratic pulls; but one in which privileged participants have more control over shifting attention than the majority.

One fundamental difference is that, as just described, the forum dynamics currently capture only the current level of excitement about content and not a measure of how many have visited content. We discuss this in section 11.4.1.

11.3.2 Channel for sharing each state of mind in a forum

Forums include a channel for sharing excitement about each separate state of mind. This is consistent with conversation at a dinner table for example where separate conversations may form across each other, with each held together by its participants' shared state of mind. It feels good to be in agreement with each other; and the feeling that others' don't understand what you are trying to share feels bad. It seems reasonable for the dynamics of

this idealised model to hold strands together in a shared state of mind while new strands continue to join.

One issue with the model is that where different content leads to the same state of mind, strands whom are perhaps following different trails of content will find themselves together in the same state of mind. In reality although strands following different threads of conversation might enter the same state of mind they would not be aware of others outside this conversation whom are in the same state of mind. This is an approximation then. However, we might consider maintaining a shared state of mind to provide an ongoing binding process; and given that our theory requires social machines when acting in their capacity as a platform to support the coalescence of strands bound by compatible processes of binding, then the theory says that those in the same state of mind should indeed find each other.

11.3.3 Private forum for reacting to content

As just recalled strands who have recently entered a certain state of mind are held in that state by the shared sense of excitement of those that recently entered it. In reality this would occur through interactions that comprise reaction to this content. If these discussions were to occur in the single central content channel then forums would be very limited in their throughput before content and reaction became confused. We may therefore interpret the forum model's channels for sharing excitement about each separate state of mind instead as channels for reacting to a particular piece of content. These reaction channels allow the forum to split naturally into concurrently operating parts.

These behaviours are consistent with the operation of many social media platforms' interaction models. In interaction models such as Reddit's, participants solve this scaling problem by reacting to content in a more contained location than the channel in which it was shared. This is supported by a culture in which if one replied to a post in the same channel that it was shared in then your response would be ignored or voted down; or in which someone might reply explaining that this is "wrong". In these ways, participants are trained to use the Reddit platform in certain expected ways.

This is also supported by the software side of Reddit's interaction model. At every level when a participant shares content there is a mechanism to reply in a forum that is more contained than the forum that the content was shared into. Reddit's framework creates a fresh comment thread under each post, and each reply in this thread spawns a new thread; resulting in a tree of replies rooted in the original content. Other frameworks provide just one automatic level of nesting, often with the option for participants to quote another reply to provide some structure, but otherwise leaving the onus on the reader to work out which reply is in response to which—akin to the reader finding or creating the forum from which they are combining content by excluding some other parts. Between these culturally and technically administered interaction models most participants interact in defined ways that support scaling. This nesting of reaction is such a common approach to solving the scaling problem that it is very reassuring to find it as a core behaviour of our idealised social machine model.

It's interesting to notice that where there is no such easy mechanic for creating a new forum there are social rules that prevent confusing a central channel. For example in a formal situation it's considered rude to talk over someone, or to change the topic while others have contributions to make, or to interrupt an exchange between a smaller fragment of the group; while at the same time it's considered rude to take too much of the group's time. These rules form part of an interaction model for a social machine that forms off the web. In this kind of situation any amount of mixing of the communal content and communally approached states of mind may occur privately in the mind of a participant, but it would be unproductive to share *all* of this into the central channel. We might consider that in these situations separate forums for responding to each new piece of content do in fact form, but that these arise one after another rather than concurrently.

11.4 Model limitations

The idealised model is the simplest model we could conceive that is capable of representing the interaction of self-similar parts trough sharing experience. We discuss some limitations.

11.4.1 Social machines may sustain a flow of participants

Take a real situation such a collection of people on a Subreddit interested in finding and sharing a class of content. As the whole subreddit shifts its collective focus across content on the web or some other corpus of knowledge then we might view its strands as comprising an ensemble roaming across this content. This introduces us to a new view of the forum and the social machine requiring two behaviours not supported by the forum model. Firstly, a strand must be able to switch their attention between forums while maintaining longer term binds to forums that govern how likely their attention is to move to each. We discussed this requirement in section 10.4 in terms of the need for longer and shorter term memory kernels for binding strength.

Secondly, the forum model supports only the live interaction of participants. This might be quite suitable for physical meetings, or live discussion over a phone or video call, or even to describe the way people use instant messaging apps, if they were to pay full attention to them the whole time. However what seems to make social media work is its ability to make it feel like you are sharing experience with someone or some community even though this is not live. A forum instead may maintain the interest of a *flow of strands* that shift their short term attention onto and away from the forum. We might address this by creating longer and shorter term memory kernels governing excitement about content. Particular strands would need a way to escape any long term excitement about content perhaps through a dynamic that governs a growing disinterest in being stuck in the same state of mind for too long. Of course real social media frameworks shared to a forum's main channel and any reaction this. This requires an association between the archives left by a social machine and the machine itself. This is of course something that we see in many social media frameworks where the same way a participant might join a live conversation might provide easy access to old conversations.

11.5 Model implications

This theory of social machines has some differences from existing conceptions. We discuss these in chapter 12. Here we consider some implications of the theory on how we understand phenomena to which it applies.

11.5.1 Social machine ecosystem abstraction from medium

An ecosystem of overlapping communities must form through the agency of some underlying medium. This medium might comprise anything, although our focus has been on a medium comprising people and infrastructure. This medium supports the sharing of experience between parts of itself. Driven by its desire to form experiences, and its social nature to bind to parts of itself that it has shared experiences with, this medium interacts with itself to form this ecosystem of overlapping communities—or social machines.

These social machines and the ecosystem they form operate to some degree regardless of the hosting medium's structure. For example, suppose that part of this medium is held together by a common interest. In advancing this interest and binding together to form a sustained social machine the part will form a series of experiences. From the perspective of the whole it doesn't matter which parts contributed to form these experiences: we can just say that the machine had these experiences. The same of course can be said of the neurons in our brain—from our perspective it doesn't matter which are firing while reading this text for example. As another example, we enjoy the feeling of connection in meetings where we enter this state of sharing in which no one seeks ownership of a contribution. Although, while a healthy team might operate with no desire to own a contribution, the binding of the team is actually the sum of the binds between its participants as well as a bind to the idea of the team itself, and in maintaining these ties participants will bind to contributors and contributors to participants.

11.5.2 A social machine needs a habitat

There is a complication in that by considering a social machine to exist only while its cohesion is maintained by its ongoing operation: we can't say that a subreddit, Facebook group, or Twitter hashtag *is* a social machine—but yet they clearly relate. As part of hosting this ecosystem of social machines the medium maintains archives and services that, together with other machines in the ecosystem, form a social machine's *habitat*. These archives include for example Wikipedia pages, and social media posts and comments attributed to a subreddit, Facebook group, or Twitter hashtag, as well as much broader corpuses of knowledge. These services will include an implementation of an interaction model that supports a social machine; and they might include collections of physical workflow fragments such as moving a person between locations in a city or performing part of a scientific data collection. Subreddits, Facebook groups, and Twitter hashtags then, are *places* in this habitat where social machines tend to form.

11.5.3 Archives and strand state are inseparable

If a social machine has an obvious input and output then these are from and to this habitat. But it's not this clear cut. A social machine may also be primarily participatory in that it may have no obvious output and its purpose is to have an effect on its participants. If this effect is lasting, even if it is clearly embodied in actual people, we must say just that it is in the part of the medium that hosts the machine. We must conclude that the archives maintained by this medium include not just technical artifacts but human memories and left over states of emotion that we've identified as being held in the state of a collection of people. There is nothing in our model that distinguishes between these two.

11.6 Conclusion

We formed our theory by gathering up the requirements we used while developing our model and finessing these into defining characteristics of a social machine. We then summarised our model without reference to the mathematics we used to define scenarios

and to turn them into Petri net models. The self-similar definition of a forum raises some immediate questions. In addressing these we reasoned that a forum is comprised of other forums but that it cannot be said to fully contain them—and the corollary, that where a forum comprises other forums it cannot be said to be fully contained by them. As proposed in section 1.5 we compared the behaviours of our idealised model against the familiar behaviours of a few dominant social media platforms. We found that our model matched their main behaviours but also found some limitations: most notably the forum model as it stands supports only the live interaction of strands.

Finally we explored the implications of the theory. We noticed that the social machine ecosystem is abstracted from the medium it exists in. For example if a social machine solves a problem, then it is not important which parts contributed to that solution—other than that this may effect their binding. A key element of the theory is that a social machine exists only while maintaining cohesion through its ongoing operation. This drove us to conclude that a subreddit, Facebook group, or Twitter hashtag are not social machines in this theory. Instead we proposed that the medium that hosts an ecosystem of social machines maintains memories, archives and services that, together with other machines in the ecosystem, form a social machine's habitat. Subreddits, Facebook groups, and Twitter hashtags then, provide places in this habitat where social machines tend to form. This idea that a social machine may exist independently of its mediating infrastructure really opens up the way that we think about them—particularly in terms of their mobility and independence.

Chapter 12

Discussion

12.1 Comparison with other conceptions of a social machine . . . 169
12.2 Initial theory and adjustments made during research 172
12.3 Social interaction facilitated network technology 175
12.3.1 Representing a social interaction with the ontology 176
12.3.2 Personal difference in network-facilitated interaction 177
12.3.3 Effects on community structure. 179
12.4 Applications of theory 180

We reviewed the literature regarding social machines in chapter 2. The nature of the concepts of social machine described there all at first glance loosely agree with each other and in turn with the conception that we describe. However there are quite fundamental differences when we look at the details. In section 12.1 we compare the conception captured by our theory against some of these existing conceptions. Then in section 12.2 we take a step back and recall the nature of the theory we started with and how it developed into what is presented here. Along with the theory itself, this helps us address the question, in section 12.3, as to how social interaction facilitated by network technology differs from fact-to-face interaction. Finally, in section 12.4, we end with some example applications of the theoretical model presented here.

12.1 Comparison with other conceptions of a social machine

Scope is less constrained. In section 2.1.2 we discussed some academics' differing views on the definition of a social machine. Our work does not follow Berners-Lee's conception in which "people do creative work and the technology does the administration". This is not just because, as Smart and Shadbolt observe, the idea that technology cannot perform creative work is limiting. If we consider administration to include message routing and tracking the relationship between participants; then in our social machine models the connection to a social machine is as much held by the participants as it is by any state of the technical parts of the machine.

The theory presented leads to questions as to how small a part might be considered a social machine. We might ask if while two people are conversing they are flipping between two conversations with differing purpose then how many social machines are there? It seems natural that the situation in which they are conversing is a social machine, and clearly there are two threads of activity. In our proposed model these threads of activity meet the criteria of social machines in that either might spawn other such threads of activity.

As discussed in chapter 1, removing these constraints may make the concept less useful to web scientists. Given the self-similar nature of our conception however this was necessary to ease the representation of situations in which a social machine ecosystem includes parts that don't depend on the web for their existence.

Life cycle and boundary delineation by compatibility of binding process rather than purpose. De Roure et al. originated the idea that a social machine might be delineated by purpose, or by evolving purpose. This is a good way of describing the spatial boundaries of a social machine—although depending on what we mean by purpose it is possibly overly constraining. This idea provides a less than satisfying picture of a social machine's life cycle. There is a question of identity; as to when a social machine stops being itself.

In our theory "a social machine exists only while maintaining cohesion through an ongoing process of binding that occurs while participants form shared experience". This provides a more concrete idea about life cycle and, again, depending on how we understand the meaning of purpose, a broader scope for what holds a social machine together.

Advancing a shared sense of purpose is a good example of such a process of ongoing shared experience. Web scientists picked up on this process because it is clearly very evident and important. Other features may carry through a series of experiences though. For example we might share a common emotional response to some content; and as it can feel good to keep sharing this it may act to guide our future experiences together thus binding us. Sometimes just being in motion with others is enough to hold us together even if we don't perceive the cause of the motion—such as sharing a sense of camaraderie while letting events pass by, or helping to carry something without knowing where you are headed.

Self-similar and explicitly part of an ecosystem. The idea that social machines form part of an ecosystem was raised De Roure et al. Our work started by thoroughly embracing this. It's possible that by doing so we've lost some other meaning of a social machine; but we feel not. We discussed in chapter 2 the very thorough work of Shadbolt et al. to form a taxonomy of social machines. From the perspective of our theory here this reflects a reductionist view that was very complete for a machine operating in a vacuum, but that misses the idea that a social machine is social itself, not just comprised of social parts.

The platform is an aspect of the machine. As just introduced, a social machine might be delineated by possibly changing purpose. Shadbolt, O'Hara et al. wrote a book drawing together work performed during the SOCIAM project [74]. They add the idea that social machines with a purpose are hosted on a social platform that is, at least in theory, without purpose or agenda. This view seems quite reasonable if we think about web based social machines or work teams. However this is not how we think about computers or the Turing machine. These machines are general purpose and it is the software that runs on

them that adds purpose. In fact the social platform seems more akin to a general purpose computer.

Our requirement that a social machine theory be self-similar led to the idea that they may spawn social machines. In this way we have reconciled this difference by stating that the social machine both coalesces around a purpose and supports the coalescence of its participants around further purposes to form further social machines. For example, a machine with the purpose of making a choice may support a machine formed around the purpose of spreading the belief in each possible choice. This is what we mean by self-similar. Such a part might model either a single person or a coherent community. The experiences held by these parts might include the few seconds of experience that arise while a person considers whom is walking up the gravel path to their house, or the few generations of communal experience resulting from people asking if vaccines are safe.

A social machine is not its habitat. In section 11.5.2 we noticed that by considering a social machine to exist only while its cohesion is maintained by its ongoing operation that we can't say that a subreddit, Facebook group, or Twitter hashtag *is* a social machine. Instead we proposed that the medium that hosts an ecosystem of social machines maintains memories, archives and services that, together with other machines in the ecosystem, form a social machine's habitat. Subreddits, Facebook groups, and Twitter hashtags then, are places in this habitat where social machines tend to form. This separation of the mediating infrastructure from the social machines that exist across it was not something we saw in the literature. We feel that separating these is key to understanding the nature of social machines: an engineer does not build a social machine but a habitat for social machines to form in.

12.2 Initial theory and adjustments made during research

We started our work with certain theories about social machines and communities. Our theories changed often before settling on what is presented in this thesis. We highlight some of these changes in this section.

It is easier to reason about activity than the medium supporting it. Although going into this thesis with an interest in process philosophy, we started with a view of a social machine as an object-like system dependent on object-like participants. Unable to clearly define membership boundaries—for example, as participants' rates of interaction vary—we then tried the idea that the social machine is a process that uses these participants to do something; but were unable to form an ecosystem of interacting social machines with this conception. Finally, we settled on the idea that a social machine exists where we wish to attribute it as supporting and causing some strand of activities. For example, if a community can plough, then we infer the presence of a social machine that can plough. In retrospect, computer scientists use this form of abstraction all the time—although it was hard to ratify this with our assumption that a social machine should potentially participate in other social machines, which, as discussed above, required us to ultimately remove a place for absolute containment from our supporting theoretical model.

Boundaries imply potential connection rather than enclosure. This suggests that we should shift our view of a boundary from something that separates parts, to something that connects them. The default relationship between strands is no relationship, other than perhaps a connection to something broader that situates them in a way that they cannot interact directly—like atoms acting as mass holding together the planet that holds them in place. With containment gone, the view of an echo chamber as a community surrounded by boundaries that reflect back what its members say is harder to hold. Instead, we see conversations, or stories, at many scales, with larger, slower changing, conversations acting as a backdrop to the faster, more topical, conversations that in turn carry these larger

ones. What we saw perhaps as waves reflecting around a closed pool, is more like the chanting of many voices into a void; and what we attributed to the echo, is instead the forming of a habitat where certain elements of experience are easy to recall and express easily together, and whose presence inhibits competing arguments and therefor the open interaction with those making them.

Boundaries take the form of overlaps between parts rather than network edges. Although viewing communities as connected to form a network can provide a very good representation of an ecosystem of interacting communities, this approach cannot fully capture situations in which a person is a member of more than two communities, or in which a member has open boundaries through which they may interact with other members. Instead, in the model described here, we found that communities are better viewed as overlapping where they share members. During our work we started with this view, then gave up on it because the models it resulted in were too large when expanded out using our simulation tools. We returned to it when we found it to be much better for representing the kinds of situations just described.

Linking back to the first paragraph in this section: more particularly, we found that a model in which communities overlap where they share strands of activity, is even better than a model where they share the object-like things, such as people, we might attribute those activities to.

People connect communities as much as communities connect people. We started the work with, and still of course maintain, a view of a community as something that connects people together allowing them to interact. The underlying model presented in this thesis however draws our attention to the overlap between communities; and at these overlaps are people. As such, people may act to propagate content out from and gather content into a community; or they may act as a path for strands in other parts of themselves, or in people they know, to participate in a community or, equivalently—depending on your perspective—as a path for that community to spread through as it includes these strands into itself.

Social interactions provide a medium, or carrier, for more mechanistic interactions. We formed the view that a community or social machine, is not best envisioned as people interacting socially, but by strands of activity that interact through a medium provided by people interacting socially. Some such strands may interact in a more mechanistic manor that might sometimes be considered as expressing logical reason or as performing work. We can model interactions as composite experiences to which strands attributed to the socially interacting parties all contribute. Some elements of the composite experience may form part of the social interaction itself, while others express the purpose driving it, and some actually enact the more mechanistic interaction. These more mechanistic strands of activity continue to require the attention of their, more complex, social hosts in order to interact with strands in other hosts.

An ecosystem of overlapping social machines is more than a web of separate social machines. We started with the theory that an ecosystem of social machines forms primarily as a way of gathering and sharing the online attention of potential participants—a collection of social machines providing online activities is more able to hold our attention than just one. This view is analogous to the model behind the PageRank algorithm describing how people follow links between web pages. In it we might see, for example, participants drawn to a discussion group online and then linked off to another. On a web of static sites, the path of sites a person has visited may affect their future browsing—as those buying and selling the attention of those people, for example, know well. The same is of course true for the interactive sites that can be said to host interfaces to social machines. Now however, while browsing between these, a person may become a participant in some and be drawn back when they are able to participate, having, for example, found content that they might share, or having learned a transferable skill, or found compatible belief. Where social machines share elements of compatible purpose, or where elements of an activity can be learned and applied across them, those social machines will find advantage to sharing participants—up to the level that they compete for an overly limited resource such as attention, head space or control over incompatible belief. A person participating in a number of such machines connects them in a way that is just as real as the physical

link between web sites.

We moved from a view of a web of pages hosting largely separate social machines, to an ecosystem of potentially overlapping social machines. This feels analogous to the way that providing common mechanisms for linking data on pages affords a shift from the web to the semantic web. Here though, connections between the machines are made through these participants rather than coded as links between elements of data.

12.3 Social interaction facilitated network technology

Except for the view just presented above, there is little in the above discussion or our theory that would not, at first glance, apply equally off the web, as on it. For example, just as people might shift their attention between online communities or social machines, they may move between work and home, or between social groups, or even shift focus between different purposes without moving location; and, as they do, connect their strands of experience to those in these situations.

Although we did not focus on this during the development of this theory, there are clearly differences between social interaction facilitated by network technology and interaction without it. To ask what our theory says about these differences, we start, in section 12.3.1, by showing how a social interaction may be broken down with the ontology and represented as a composite experience including the, more personal, experience of each interactant; and how, in turn, these personal experiences may be represented as composite experiences that include elements of experience from contributing strands attributed to an interactant.

Following this we will argue, with some help from this model of a social interaction, that (i) there are some fairly obvious differences to online channels mediating interactions from offline ones; (ii) the structure of online communities may be different, in part as an adaptation to (i); and that, (iii) these changes affect the distribution of types of interactions that participants experience. There is the potential for feedback around these three as (iii) may in turn effect (i); this might be true, for example, if we become better habituated to online than offline communication or if online communities out compete

offline communities.

We discuss the first in section 12.3.2, and possible effects due to these changes in section 12.3.3. To feed into this discussion, Kraut et al. describe the big differences when we take these communities online[46, page 2]:

> The promise of online communities is that they break the barriers of time, space, and scale that limit offline interactions. People with unusual medical conditions can get social support from others who share their condition but live far away, and they can do so whenever they need it rather than only at a weekly or monthly scheduled meeting. On ravelry.com, knitters can share patterns with thousands more people than they could stitch with in person.

12.3.1 Representing a social interaction with the ontology

We assume that social interaction can be broken down into discrete interactions. To fill out the theory presented above, that social interactions provide a medium, or carrier, for more mechanistic interactions, we must represent social interactions as experiences in our model. When applying the ontology to describe an experience, it is important to think of an experience such as a social interaction not just as a record of what went into forming it, but from the perspective of how it came to form and what occurred during its formation. For example, an experience may be affected by the structure of the strands contributing to it, and this may or may not, be reflected in the experience itself.

A social interaction between a collection of interactants can be represented with our ontology as a composite experience; possible forming a strand of such experiences localised to a particular ongoing social situation, for example, the presence of people at a table. A shared social experience includes the more localised, personal, experiences formed by each interactant during the interaction. These personal experiences, in turn, include the, even more localised, experiences formed as some of the strands attributed to that interactant participate in the interaction.

We've shown models that represent aspects of normative binding during the experience of a social interaction. During these interactions the interactants' personal experiences

include both the thing being shared, alongside the purpose for the thing being shared. Other aspects of these personal experiences would in reality include further social elements such as perhaps introductions and small talk; and of course some effect resulting from facts and feelings connected to the other participants and indeed the situation affording the social interaction itself. This might include identifiable features of the other interactants, such as their name, and perhaps the reason for the social situation formed in the first place—perhaps related to the shared purpose.

12.3.2 Personal difference in network-facilitated interaction

We make some observations about differences in online versus face-to-face interactions as our theory changed and developed.

Reduction of inter-personal and social cues mediating strand interaction. In the previous section we described social interactions as providing a medium for more mechanistic interactions of strands of activity. In the majority of our online interactions, the interaction of strands is not mediated through face-to-face or other conventional social interactions. There is presumably a reduction of objective and explicit cues that might, in face-to-face communication, carry emotional affect, tells of trust, or intention. For example, cues normally carried by intonation and tone, intensity, cadence and facial expression, or body language may be missing online. We saw this kind of disconnect while proposing simple models supporting different levels of shared intentionality, in which the channel mediating strand interaction is limited by an infant's stage of development.

Increased anonymity. Reflecting our model of a forum as affording both peer-to-peer and larger-scale anonymous interaction, participants presumably spend more time being anonymous online than they would offline—where anonymity may be provided by a pseudonym or proxy, or just by being one of a huge number of indistinguishable participants. When strands interact they do so, in these cases, during social interactions that result in experiences that don't carry this context. So, despite strand interactions existing with the help of participant's social abilities, the many social norms that might

mediate the interaction of these strands, for example, in a local community where the people might see personal repercussions for allowing certain strands to express themselves, are absent. In the other direction, the fact that the content of a message is associated with something 'trustable' rather than with a bunch of advertising executives, political spin doctors, foreign agents, or server farms running increasingly sophisticated AI, would change the experience of a twitter post for example—and therefor potentially the reception of the post's content.

Separation of online from offline community membership. Given the nature of online interactions it is easy for these to be hidden, deliberately or not, from our offline family and friends. There is more chance for us to experience strands of activity that only interact online, and to become part of communities that have very little chance to overlap with our offline communities.

Interaction with communities unbounded by time, space and scale. Without the analysis of how personal interactions affect community structure, we can jump ahead and ask how the breaking of barriers of time, space, and scale that limit offline interactions effect people: (i) participants may move far more quickly between communities in which they interact; (ii) they may move between a far larger pool of communities; (iii) they need not always pass through certain communities to reach others; (iv) these communities and their members may be anywhere and intersect parts of cultural regions physical unreachable offline; (v) and participants may interact coherently within much larger communities (for example speaking to more people, hearing more people and forming consensus with more people).

Communities may come to the participant rather than a participant to communities. Scrolling on Reddit, for example, we might skim over scores of posts a minute and connect with the community behind something that catches our attention at the tap of a finger. We don't need to walk between places where these communities meet, or visit a community via its url, or click on a link to it: these communities come to us.

12.3.3 Effects on community structure.

We've itemised some differences in how social interactions online may differ at a personal level from offline ones, and now propose how some of these differences might effect community structure.

Possible effects on community structure of the reduction of inter-personal cues mediating strand interaction. The possible reduction of inter-personal cues and increased anonymity of people interacting online—supposed above—would effect the nature of online social interactions. These interactions lack elements of experience that would otherwise indicate how participants in a conversation feel about a message. What effect might this have within an ecosystem of overlapping communities subject to a form of natural selection?

Assume that experiences with a common shared reaction are more compelling than those without. This suggests that communities that routinely offer a clear and predictable reaction to certain kinds of content they share will outcompete those that don't. This, in turn, suggests that communities with more homogenous views may be more viable; and suggests that the range of content they share will be limited to that where the common reaction is easily knowable. These limitations suggest that these form of communities offer a narrowed exposure to people's fuller lives. These effects might drive cultural polarisation.

Abstraction from physical space. We did not consider differences between online and offline social interaction during the main development of the thesis. The abstract social machine theory and the model suitable for describing overlapping communities of strands of experience that supports this, seems suitable for offline and online interaction. The model has no explicit place for physical location or distance, and nor does it have place for physical networks limiting connectivity and latency, for example. With its absence of an explicit place for physical extension, the model is actually more capable of modelling online social interaction than offline. Here, a final missing element is the delay between strands caused by the speed of light, and other causes of network latency. Although this is of huge importance in some situations, such as for stock trading and therefor the propagation of

news events monitored by algorithmic traders, it need not impact live conversation.

Coherent across more people during faster rates of change. Now, when we say that a collection of strands acting in some forums as a single strand is coherent, this means that those strands are seen to move together. This does not mean that they move in lock-step, but that, like the self-reinforcing particles forming the wave packet of a soliton, they stay together as they form experiences together; such that they all experience them in the same order and closely enough in time that they don't break up into multiple strands. Without network technology this coherence can be maintained over large distances if the rate of experience is slow enough. With network technology this coherence can potentially exist for real-time conversation at a global scale of distance. For example, the coherence of a state of activity, such as holding a changing system of beliefs and keeping this relevant, can be maintained across larger groups of people for faster changing belief systems online than it can offline.

Coherence within an online community then is not limited by physical distance, but by the time it takes for strands within the community to share a new experience with enough of each other that the community as a whole maintains the coherence of its ongoing state of activities. Where the strands' hosts have other interests or needs this will form a limit to cohesion, as this impacts the time it takes for an experience to percolate between parts of the community. In our model then, all the strands forming a larger coherent strand need not be situated in the same forum, but the forums they do constitute must overlap enough that this cohesion is maintained.

12.4 Applications of theory

Building successful online communities. In section 9.7 we introduced the body of knowledge summarised in Robert Kraut et al's book "Building Successful Online Communities. Many of the design claims proposed by this community are empirically-grounded by observations and experiments. Modelling scenarios that capture these would exercise our theoretical model and uncover more realistic mechanisms governing how

experiences form in communities. In turn we feel that our theoretical model could form a common language for i) describing situations in which observations are made; ii) defining the design claim's observable conditions and for testing for the expected outcome; and iii) building communities and for designing frameworks that encompass these designs. With strong enough models of a given scenario—real or under design—the theoretical model might also be used to prove that a design claim is true; and to prove that a scenario exhibits the conditions that will cause the expected outcome.

Basis for mixing positive and interpretive social science approaches. Our theory describing ecosystems of overlapping communities includes a space for the subjective experience of people. As such it could be used as a framework for sociologists to combine data-driven positive approaches with participatory interpretive approaches to studying systems of people.

Basis of framework for growing human-driven human-technology hybrid organisation. The ability to first analyse existing structures of communities and their activities, and to then build models that can both predict and, sometimes, perform the activities that parts of these communities perform, provides a useful process for automating aspects of these communities and their activities—especially if the same underlying model can be used for both. For example, a model might predict the next steps a community member is likely to take, and offer a plan for them to follow, or to perform some steps automatically. Where these predictions are based on the activities of more experienced participants they can be seen as providing training to participants themselves, and not just the technology supporting them.

Basis for creating transferable social machine specifications. Mechanistic models of computation may apply well to some social machines. For example, as described by Malone et al.[48], members of crowdsourcing communities may act as a pool of non-specialist workers for performing activities that can run concurrently such as looking for galaxies; or, where a pool of specialists are needed on a more ad-hoc basis, we might draw

an analogue to dynamic dispatch. Strands that overlap within a social machine and an expert connect the activities of that social machine to other strands accessible to the expert. These might carry specialist knowledge to fold into shared experiences formed as the machine operates—or they might contribute memories of past experience, or recent, or ongoing experiences. Participants that have joined more than one community might act to connect strands across these—like a facilitator, or broker, connecting people who can help each other.

The operation of some social machines might be specified in a way that they can be deployed deliberately without direct knowledge of which people or technology they might be deployed into. The abilities and skills needed by participants might be specified in the form of experiences or components of composite experiences. Larger experiences might include paths of such activities possibly defined by workflows. The specification might include a deployment process that includes sequences of operations designed to train participants with the capacity to form these experiences.

Chapter 13

Conclusion and future work

13.1 Conclusion 183
13.2 Future work 189

13.1 Conclusion

We started this work by asking how we might best understand ecosystems of overlapping communities, and focused this inquiry by looking particularly at the phenomena of social machines arising, or deployed, within such ecosystems. We took defining characteristics of the phenomena of social machines from existing conceptions. We started our theory, by assuming, among other things, that a representation of these communities, and of social machines, should be self-similar. To explore the feasibility of this and to understand the repercussions we developed a simple model of this ecosystem hoping that something with these characteristics would arise. Our conception of an idealised social machine, as emergent structures with these characteristics grew while doing this. In some cases the behaviours required of these structures seemed inescapable and in these cases we added them to the characteristics required of a social machine in our theory.

Theory. The theory comprises a collection of characteristics that we propose are required of any social machine that is to form part of an ecosystem. These characteristics are listed in section 11.1. In this theory an ecosystem of overlapping communities forms through the agency of some underlying medium. This medium might comprise anything. Although our focus has been on a medium comprising people and infrastructure, it might include technical automota or workflows for example—and in fact the idealised model that we developed alongside the theory is a medium that should, if executed, be capable of hosting social machines. This medium supports the sharing of experience between parts of itself. Driven by its desire to form experiences, and its social nature to bind to parts of itself that it has shared experiences with, this medium interacts with itself to form an ecosystem of overlapping communities—or social machines. The social machine ecosystem is abstracted from the medium it exists in. For example if a social machine solves a problem, then it is not important which parts contributed to that solution—other than that this may effect how the parts bind.

There is a strong implication of considering a social machine to exist only while its cohesion is maintained by its ongoing operation: we can't say that a subreddit, Facebook group, or Twitter hashtag *is* a social machine—but yet they clearly relate. As part of hosting this ecosystem of social machines the medium maintains archives and services that, together with other machines in the ecosystem, form a social machine's *habitat*. These archives include for example Wikipedia pages, and social media posts and comments attributed to a subreddit, Facebook group, or Twitter hashtag, as well as much broader corpuses of knowledge. These services will include an implementation of an interaction model that supports a social machine; and they might include collections of physical workflow fragments such as moving a person between locations in a city or performing part of a scientific data collection. Subreddits, Facebook groups, and Twitter hashtags then, are *places* in this habitat where social machines tend to form. This idea that a social machine may exist independently of its mediating infrastructure really opens up the way that we think about them—particularly in terms of their mobility and independence.

We also found strong implications of the requirement that a social machine be self-

similar. We explored these initially while developing the idealised model summarised in section 11.2.1. As described in section 11.2.2 this requirement constrains the topology of a social machine ecosystem by prohibiting the absolute containment of one social machine within another. It also leads to the idea that a social machine acts as a platform that can spawn new social machines. In this conception a social machine is social, not just in the sense that it depends on the social nature of its supporting medium for its operation, but that it forms always part of at least two other social machines in its connection to a broader ecosystem. To support this ecosystem we also require that the social machine, working in its capacity as a platform, can potentially introduce social machines with compatible binding processes such that they might coalescence. In our model, described in more detail below, where multiple social machines offer the same binding process to participants they will tend to favour the largest. This results in a network effect that will add larger structures to an ecosystem.

Models. We propose that this theory is best represented by a modelling part that is more generic than a social machine. A *forum* comprises a collection of forums and is part of a collection of forums. To add to this circular definition we assert the existence of forums that, when interacting with other forums, act as if they support no internal interaction. These *strands* interact through shared experience, with a model defining how past experiences build on each other to shape future experiences. To support this assertion the rules governing the interaction of these strands hosted within a network of forums must give rise naturally in turn to forums with this strand nature. We reason that this can occur if: closely connected forums similarly situated in a broader network may act like a single forum in this broader network; strands in a forum that have similar enough past experiences may form flocks that move together along paths of experience; and strands sharing experiences strengthen the ties between the forums that host them. We developed an illustrative set of rules and show that they have this property; and then applied this to model a scenario that recreates Axelrod's model of the dissemination of culture.

The approach we used to define scenarios and to then represent them as Petri nets is described in the text. The Petri net models are executable and this leads to something

interesting. If a Petri net implementation of the model in the text is executed some forums will persist; and if they do it will be because they support an ongoing process that binds their participants together through a tendency to bind to the source of shared experiences. Such a persisting forum would be a model of a social machine that might arise in the scenario.

Given the broad scope of our theory of social machines we could in fact say that this persisting forum *was* a social machine. In section 7.3 we showed that the idealised model supports the emergence of separate regions of culture. Each region is held together by an ongoing state of binding as its members chat about certain beliefs they share. However these structures were not technically social machines according to the theory. The most complete Petri net implementation of the model did not support features added after chapter 7. This Petri net model lacks the behaviours of chapter 8 that allow forums to overlap rather than comprise an explicit network. This is not required by the theory for the emergent structure to be a social machine. However the behaviours added in chapters 9 and 10 are required by the theory in order that the resulting structures support an ecosystem of social machines: the Petri net did not bind all contributors to an experience to their sources of motivation, and nor did it create a new forum for each composite experience.

A modelled scenario can be represented as single coloured stochastic Petri net. This mathematical form supports a wealth of analytical and statistical tools. For example, the stochastic rate equations governing the dynamics of large homogeneous aggregates of strands or forums can be approximated by continuous time differential equations.

Confidence in ontology, model and theory. The theoretical model of overlapping communities was developed alongside an ontology of categories of parts and behaviours. To build confidence in the ontology, we used it to break down some real world example scenarios into categories of connected parts and behaviours:

- We proposed models of scenarios exhibiting the ability to share intentionality at three stages of infant development described by Tomasello et al. [88]. (Section 5.2.1.)

- We broke down the social machine, and its operations, identified by Tarte et al. in their case study of the life of the collection of images created by citizen scientists that culminated in the recognition of the "pea galaxy"[82]. We identified classes of strand of activity and the forums associated with their interaction before the collection started; and then described how the resulting structure changed during the first few moments of the collection's life. (Sections 5.2.2 and 9.1.1.)

- We broke down generic instances of a Subreddit, Facebook group and Wikipedia entry in a similar way; and then briefly considered Wikipedia as a single social machine working on a single interconnected corpus of knowledge. (Section 5.2.3.)

These examples give us confidence that the ontology, and likely the model built using this, can be used to represent scenarios in which social machines are present—at least as understood by the authors of those works. More particularly, the examples relate each category in the taxonomy to phenomena identified by others.

To gain confidence in the theoretical model, we used it to represent scenarios exhibiting well-studied phenomena:

- We built models that exhibited the flocking behaviours we needed to advance the argument of the thesis. The scenarios these described, when interpreted as containing social machines, exhibit the dynamics captured by the HTP model of Tinati et al. [84]. Particularly, the simulation resulted in the emergence of a system of four overlapping social machines bound by a common element of three otherwise differing agendas; this system shifted between phases of stable activity as the advancement of these agendas impacted one another. (Section 5.5.)

- We showed that a model of a behaviour that mutually binds participants together whenever one of them advances a purpose they share, captures the observable conditions sufficient, according to a design claim of Kraut et al., to result in those peoples' increased normative commitment to a community—which provides a well studied example of the phenomena represented by the model. (Section 9.7.)

- We used the model in section 7.3 to model a scenario that behaved as Axelrod's model predicts.

These simulation results are based on implementations that fully reflect some aspects of the final model, and an approximation to some others. We extended this model in the later chapters of the thesis to result in a theory of overlapping communities .

We have confidence in this resulting theory of overlapping communities, having compared it with examples from reality and having used it to reason about some common phenomena in section 11.3. Particularly we confirmed that the structure and dynamics of the model share familiar features with the interaction models of the dominant social media brands—suggesting a suitability for our more narrowed purpose of explaining phenomena described as social machines. We showed that it meets some obvious criteria to be suitable for modelling an ecosystem of forums; and that a natural way to add a behaviour where strands bind to the source motivating an experience tends to result in the well studied binding mechanism described in section 3.2.3

As described at the start of chapter 9 this approach led to one of the core contributions in our theory of social machines: that they exist only in a state of ongoing binding through a dynamic that binds them to a source of shared experience. We compared this to other ideas about the life cycle of social machines in section 12.1.

Application. The theory presented comprises a collection of characteristics required of a social machine that taken together allow it to form part of an ecosystem that in turn supports its existence. We describe a number of possible applications both in chapter 1 and section 13.2. Two things are worth reiterating.

To an engineer this theory might be used to form a vision of what a system or ecosystem they are building might look like. The theory provides a language with which to formalise this system's operation, and requirements for building a habitat that might host the social machines that comprise it.

Models derived form this work might may be used in different ways. As we've seen, a model of a scenario may be used to create possible sequences of events. Such synthesis

is ultimately useful not just for exploring ideas and generating data for testing analysis methods but for the analysis of real world data itself: many techniques for fitting models to data require an iterative approach in which models are tuned to minimise the difference between the data they produce and real data.

13.2 Future work

Before advancing this work it will be important to confirm that the work to date might gain traction in the broader community. To this end we suggest a few ways that the work might be applied and tested and also some ways in which it might be tied more closely to some specific disciplines. We have left some smaller threads hanging that might provide opportunities to incrementally advance the model and theory. We mention those next and then discuss some bigger pieces of work that might bring things together more tightly.

Application of work in the broader community. These items are suggestions to apply and situate this work within the academic disciplines it connects with.

- **Webs science.** In section 11.3 we began to apply our theory and idealised model to Reddit's subreddits, Twitter's hashtags and Facebook's groups. There is an opportunity to do this in more detail. We might also ask if the work helps us understand echo chambers. Tarte et al. [82] propose that one element of a social machine is that it forms an identity in that members and non-members may tell stories about it. It would be interesting to ask if the theory supports some real instances of this phenomena. It would be useful to extend the characteristics used to describe a social machine in [77] to capture the way a machine fits into the broader ecosystem of social machines.

- **Engineering.** We could ask if the characteristics required of a social machine by our theory form a useful set of requirements or design principles.

- **Social science.** We might work with others to model some real scenarios off the web. The theory might be compared and contrasted to actor network theory and

other common models used by sociologists. Society tends to include status in social relationships between people—we could ask if this can be represented in a model of underlying interaction in which there is no such explicit role of status. Following from this we might ask if the model aligns with models of influence and control.

- **Computer science.** We might revisit the theory and idealised model while explicitly pushing away a tendency to assume that a social machine's medium includes people. For example just replacing the term 'state of mind' with 'state' would go a long way to doing this. This might include a comparison of the proposed social machine life cycle with the concepts of spawning and forking processes. It would be interesting if our model of experience is compatible with the generic unrestricted grammar of the Turing machine. In chapter 12 we briefly discuss agency in our theory and idealised model; and we might compare the characteristics of our theoretical social machine with those of agents [52]. Social machines might also relate to AI institutions.

- **Archiving and AI.** In many of the examples we described, participants in social machines are bound around a purpose of classifying or labelling particular classes of content. There might be something to learn by viewing Reddit for example as an ecosystem that not only labels content into categories but that discovers how those categories overlap resulting in a taxonomy.

- **Engineering, business or scientific workflows.** We suggested in section 4.3 that we might switch out the single labelled transition system in the formalisation of our idealised model to a collection of state machines with start and end states. These might define workflows—or more particularly workflow fragments, as a strand would need to enter the start state of each fragment from outside. These fragments might describe a collection of operations performed on a scientific instrument for example. This could form the basis for modelling and designing facilities whose operations can be described by social machines that operate such workflows.

There are further examples for this future work in the applications listed in chapter 1.

Next incremental steps. There are many incremental advances that might be made to this work. Advancing some of these might also advance the theory and model.

- We did not implement the final idealised model using Petri nets because the models grew too quickly. Using Petri nets is rigorous but models do get very large when expanded out into uncoloured Petri nets for execution. Custom code could handle much larger models and run a lot more quickly.
- The closed form analytical properties of our models could be usefully investigated.
- The tools of information theory could help investigate mechanisms that force apart social machines with too little focus.
- Analytical metrics could be further refined. For example, we started developing analytical metrics for measuring how coherent a group of strands is—the average time for all pairs of strands to reach each other's state seems useful. The software described in appendix A calculates this metric. However the structure of some experience graphs means that some pairs could never end up together and applying this metric must therefore be coupled with clustering strands into groups.

Explicit changes to theory and model. These items are suggestions for future extensions to the model and theory.

- A single state of mind in our models could be considered as a state of activity in which something is maintained. In this way we might expand a state to a collection of paths of experience that can hold an ensemble of strands together.
- Content might be treated no differently from experience. See section 4.5.3 for a discussion about extending the experience graph to support composite content.
- Our models are missing the back-tracking that is implicit within Whitehead's model and in the way we explore ideas. As such our models were limited to cycles in which strands have many chances to cohere together—addressing this may relate to the life

cycle of strands and fully embracing the idea that a strand is just a limiting form of persisting forum of simpler strands.

- We have considered a forum as an ideal state but all we require of the forum is that it support a coherent thread of experience. We might infer the presence of a forum even if it takes a while for a shared state of mind to spread through it—for the thread to maintain cohesion this rate of percolation needs only be faster than the rate at which this coherent thread of experience is changing. Also, we have not looked in depth at phenomena that span multiple time scales: for example, patterns of phrases reused by a community might vary more slowly than the stories told with them; and the slowly changing corpus of knowledge accessible to a community of scientists changes more slowly than that used by just part of this community.

- See section 10.4 for a discussion about long and short term binding.

- Alfred Whitehead's work was very much the launch pad for our approach here. There is an opportunity to return to this and more thoroughly relate the two.

References

[1] *A Standards-based, Open and Privacy-aware Social Web*. https://www.w3.org/2005/Incubator/socialweb/XGR-socialweb-20101206/. (Visited on 08/08/2018).

[2] Tilak Agerwala. "Putting Petri Nets to Work". In: *Computer* 12.12 (1979), pp. 85–94.

[3] Eduardo G. Altmann, Janet B. Pierrehumbert, and Adilson E. Motter. "Beyond Word Frequency: Bursts, Lulls, and Scaling in the Temporal Distributions of Words". In: *PLOS ONE* 4.11 (Nov. 2009), e7678. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0007678. (Visited on 02/15/2019).

[4] Rajeev Alur et al. "Hybrid Automata: An Algorithmic Approach to the Specification and Verification of Hybrid Systems". In: *Hybrid Systems*. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg, Oct. 1992, pp. 209–229. DOI: 10.1007/3-540-57318-6_30. (Visited on 07/11/2018).

[5] Robert Axelrod. "The Dissemination of Culture: A Model with Local Convergence and Global Polarization". In: *Journal of Conflict Resolution* 41.2 (Apr. 1997), pp. 203–226. ISSN: 0022-0027, 1552-8766. DOI: 10.1177/0022002797041002001. (Visited on 05/24/2020).

[6] Judit Bar-Ilan and Noa Aharony. "Twelve Years of Wikipedia Research". In: *The 2014 ACM Conference*. New York, New York, USA: ACM Press, 2014, pp. 243–244.

[7] Tim Berners-Lee. *The Scale-free Nature of the Web*. 1998.

[8] Tim Berners-Lee and Mark Fischetti. *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor*. Vol. 1. Harper, 1999.

[9] Tim Berners-Lee and Lalana Kagal. "The Fractal Nature of the Semantic Web". In: *AI Magazine* 29.3 (Sept. 2008), p. 29. ISSN: 0738-4602, 0738-4602. DOI: 10.1609/aimag.v29i3.2161. (Visited on 06/12/2019).

[10] David Bohm. *Wholeness and the Implicate Order*. Routledge, 1980.

[11] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. "Statistical Physics of Social Dynamics". In: *Reviews of Modern Physics* 81.2 (May 2009), pp. 591–646. DOI: 10.1103/RevModPhys.81.591. (Visited on 07/09/2018).

[12] J Chodak and M Heiner. "Spike - a Command Line Tool for Continuous, Stochastic & Hybrid Simulation of (Coloured) Petri Nets". In: *Proc. 21th German Workshop on Algorithms and Tools for Petri Nets (AWPN 2018)*. University of Augsburg, Oct. 2018, pp. 1–6.

[13] Andy Clark. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press, Dec. 2008. ISBN: 978-0-19-986885-8. (Visited on 07/04/2018).

[14] M.E. Crovella and A. Bestavros. "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes". In: *IEEE/ACM Transactions on Networking* 5.6 (1997), pp. 835–846. ISSN: 10636692. DOI: 10.1109/90.650143. (Visited on 07/01/2020).

[15] Mark d'Inverno et al. "Communicating Open Systems". In: *Artificial Intelligence* 186 (July 2012), pp. 38–94. ISSN: 0004-3702. DOI: 10.1016/j.artint.2012.03.004. (Visited on 07/05/2018).

[16] David De Roure et al. *Observing Social Machines Part 1: What to Observe?* Tech. rep. Apr. 2013.

[17] David De Roure et al. "Observing Social Machines Part 1: What to Observe?" In: *Proceedings of the 22Nd International Conference on World Wide Web*. WWW '13 Companion. New York, NY, USA: ACM, 2013, pp. 901–904. ISBN: 978-1-4503-2038-2. DOI: 10.1145/2487788.2488077. (Visited on 02/23/2018).

[18] David De Roure et al. "Observing Social Machines Part 2: How to Observe". In: *The ACM Web Science Conference*. New York, New York, USA: ACM Press, 2015, pp. 1–5.

[19] David De Roure et al. “Towards a Cyberphysical Web Science: A Social Machines Perspective on PokéMon GO!” In: *Proceedings of the 10th ACM Conference on Web Science*. WebSci ’19. New York, NY, USA: Association for Computing Machinery, 2019, pp. 65–69. ISBN: 978-1-4503-6202-3. DOI: 10.1145/3292522.3326043.

[20] Morris H. DeGroot. “Reaching a Consensus”. In: *Journal of the American Statistical Association* 69.345 (Mar. 1974), p. 118. ISSN: 01621459. DOI: 10.2307/2285509. JSTOR: 2285509. (Visited on 03/07/2018).

[21] Anhai Doan, Raghu Ramakrishnan, and Alon Y. Halevy. “Crowdsourcing Systems on the World-Wide Web”. In: *Commun. ACM* 54.4 (Apr. 2011), pp. 86–96. ISSN: 0001-0782. DOI: 10.1145/1924421.1924442. (Visited on 07/02/2018).

[22] Mohammed Elkoutbi and Rudolf K Keller. “Modeling Interactive Systems with Hierarchical Colored Petri Nets”. In: (), p. 7.

[23] Reddit employee. *Introducing Image Uploading Beta*. https://www.reddit.com/r/changelog/comments/4kuk2j/reddit_change_introducing_image_uploading_beta/. May 2016.

[24] Qian Gao et al. “A Multiscale Approach to Modelling Planar Cell Polarity in Drosophila Wing Using Hierarchically Coloured Petri Nets”. In: *Proceedings of the 9th International Conference on Computational Methods in Systems Biology - CMSB ’11*. Paris, France: ACM Press, 2011, p. 209. ISBN: 978-1-4503-0817-5. DOI: 10.1145/2037509.2037538. (Visited on 10/30/2018).

[25] Michael Gelfond and Vladimir Lifschitz. “Action Languages”. In: (1998).

[26] Giangiacomo Gerla and Annamaria Miranda. “Inclusion and Connection in Whitehead’s Point-Free Geometry”. In: 2006.

[27] M. Girvan and M. E. J. Newman. “Community Structure in Social and Biological Networks”. In: *Proceedings of the National Academy of Sciences* 99.12 (June 2002), pp. 7821–7826. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.122653799. (Visited on 02/15/2019).

[28] M. Girvan and M. E. J. Newman. "Community Structure in Social and Biological Networks". In: *Proceedings of the National Academy of Sciences* 99.12 (June 2002), pp. 7821–7826. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.122653799. (Visited on 07/09/2018).

[29] Philip Goff, William Seager, and Sean Allen-Hermanson. "Panpsychism". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2022. Metaphysics Research Lab, Stanford University, 2022.

[30] Roger Guimera et al. "Self-Similar Community Structure in a Network of Human Interactions". In: *Physical review E* 68.6 (2003), p. 065103.

[31] Alan G. Hawkes. "Spectra of Some Self-Exciting and Mutually Exciting Point Processes". In: *Biometrika* 58.1 (1971), pp. 83–90. ISSN: 0006-3444, 1464-3510. DOI: 10.1093/biomet/58.1.83. (Visited on 11/05/2020).

[32] Monika Heiner et al. "Snoopy – A Unifying Petri Net Tool". In: *Application and Theory of Petri Nets*. Ed. by Serge Haddad and Lucia Pomello. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2012, pp. 398–407. ISBN: 978-3-642-31131-4.

[33] Andreas Hepp, ed. *Connectivity, Networks and Flows: Conceptualizing Contemporary Communications*. The Hampton Press Communication Series. Communication, Globalization and Cultural Identity. Cresskill, NJ: Hampton Press, Inc, 2008. ISBN: 978-1-57273-856-0.

[34] Hubert J. Hermans, Harry J. Kempen, and Rens J. Van Loon. "The Dialogical Self: Beyond Individualism and Rationalism." In: *American Psychologist* 47.1 (1992), pp. 23–33. ISSN: 1935-990X, 0003-066X. DOI: 10.1037/0003-066X.47.1.23. (Visited on 05/22/2020).

[35] Hubert J.M. Hermans. "The Dialogical Self: Toward a Theory of Personal and Cultural Positioning". In: *Culture & Psychology* 7.3 (Sept. 2001), pp. 243–281. ISSN: 1354-067X, 1461-7056. DOI: 10.1177/1354067X0173001. (Visited on 10/22/2019).

[36] John H. Holland. *Signals and Boundaries : Building Blocks for Complex Adaptive Systems*. Cambridge, Mass., UNITED STATES: MIT Press, 2012. ISBN: 978-0-262-30589-1.

[37] Pierre Jacob. "Intentionality". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2019. Metaphysics Research Lab, Stanford University, 2019.

[38] William James. *Some Problems of Philosophy*. Vol. 7. Harvard University Press, 1979.

[39] Kurt Jensen. "Coloured Petri Nets and the Invariant-Method". In: *Theoretical Computer Science* 14.3 (Jan. 1981), pp. 317–336. ISSN: 0304-3975. DOI: 10.1016/0304-3975(81)90049-9. (Visited on 01/23/2019).

[40] Emily M. Jin, Michelle Girvan, and M. E. J. Newman. "Structure of Growing Social Networks". In: *Physical Review E* 64.4 (Sept. 2001), p. 046132. DOI: 10.1103/PhysRevE.64.046132. (Visited on 07/09/2018).

[41] Luiz Antonio Joia and Carla Danielle Soares. "Social Media and the Trajectory of the "20 Cents Movement" in Brazil: An Actor-Network Theory-based Investigation". In: *Telematics and Informatics* 35.8 (2018), pp. 2201–2218. ISSN: 0736-5853. DOI: 10.1016/j.tele.2018.08.007.

[42] Michael L. Katz and Carl Shapiro. "Systems Competition and Network Effects". In: *Journal of Economic Perspectives* 8.2 (June 1994), pp. 93–115. ISSN: 0895-3309. DOI: 10.1257/jep.8.2.93. (Visited on 08/08/2018).

[43] G. Klein, B. Moon, and R.R. Hoffman. "Making Sense of Sensemaking 1: Alternative Perspectives". In: *IEEE Intelligent Systems* 21.4 (July 2006), pp. 70–73. ISSN: 1541-1672. DOI: 10.1109/MIS.2006.75. (Visited on 06/23/2021).

[44] Klein, Moon, and Hoffman. "Making Sense of Sensemaking 2: A Macrocognitive Model". In: *IEEE Intelligent Systems* 21.5 (Oct. 2006), pp. 88–92. ISSN: 1541-1672. DOI: 10.1109/MIS.2006.100. (Visited on 06/23/2021).

[45] G. T. Kneebone. “The Application of Mathematics to the Natural World”. In: *Mathematical Logic and the Foundations of Mathematics : An Introductory Survey*. 1963.

[46] Robert E Kraut and Paul Resnick. *Building Successful Online Communities: Evidence-based Social Design*. Mit Press, 2012.

[47] Bruno Latour et al. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford university press, 2005.

[48] T. W. Malone, R. Laubacher, and C. Dellarocas. “The Collective Intelligence Genome”. In: *IEEE Engineering Management Review* 38.3 (2010), pp. 38–52. ISSN: 0360-8581. DOI: 10.1109/EMR.2010.5559142.

[49] Thomas W Malone, Robert Laubacher, and Chrysanthos Dellarocas. “Harnessing Crowds: Mapping the Genome of Collective Intelligence”. In: (2009).

[50] Hans Meinhardt. “Turing’s Theory of Morphogenesis of 1952 and the Subsequent Discovery of the Crucial Role of Local Self-Enhancement and Long-Range Inhibition”. In: *Interface Focus* 2.4 (Aug. 2012), pp. 407–416. ISSN: 2042-8898, 2042-8901. DOI: 10.1098/rsfs.2011.0097. (Visited on 03/27/2018).

[51] S. R. L. Meira et al. “The Emerging Web of Social Machines”. In: *2011 IEEE 35th Annual Computer Software and Applications Conference*. July 2011, pp. 26–27. DOI: 10.1109/COMPSAC.2011.12.

[52] Nicholas R Jennings Michael Wooldridge. “Intelligent Agents: Theory and Practice”. In: *doi.org* (May 2008), pp. 1–38.

[53] George A Miller. “WordNet: A Lexical Database for English”. In: *Communications of the ACM* 38.11 (1995), pp. 39–41.

[54] Marvin Minsky. *The Society of Mind*. New York: Simon and Schuster, 1986. ISBN: 978-0-671-60740-1.

[55] Michael K. Molloy. “Performance Analysis Using Stochastic Petri Nets”. In: *IEEE Transactions on computers* 9 (1982), pp. 913–917.

[56] Dave Murray-Rust and Dave Robertson. "LSCitter: Building Social Machines by Augmenting Existing Social Networks with Interaction Models". In: *Proceedings of the 23rd International Conference on World Wide Web*. WWW '14 Companion. New York, NY, USA: ACM, 2014, pp. 875–880. ISBN: 978-1-4503-2745-9. DOI: 10.1145/2567948.2578832. (Visited on 07/05/2018).

[57] Dave Murray-Rust et al. "Towards Executable Representations of Social Machines". In: *Diagrammatic Representation and Inference*. Ed. by Peter Chapman et al. Cham: Springer International Publishing, 2018, pp. 765–769. ISBN: 978-3-319-91376-6.

[58] Leandro Marques do Nascimento et al. "A New Architecture Description Language for Social Machines". In: *Proceedings of the 23rd International Conference on World Wide Web*. WWW '14 Companion. New York, NY, USA: ACM, 2014, pp. 873–874. ISBN: 978-1-4503-2745-9. DOI: 10.1145/2567948.2578831. (Visited on 07/06/2018).

[59] Nikolic, I. "Co-Evolutionary Method For Modelling Large Scale Socio-Technical Systems Evolution". PhD thesis. NGInfra / TU Delft, Nov. 2009. (Visited on 02/15/2019).

[60] Randal S. Olson and Zachary P. Neal. "Navigating the Massive World of Reddit: Using Backbone Networks to Map User Interests in Social Media". In: *PeerJ Computer Science* 1 (May 2015), e4. ISSN: 2376-5992. DOI: 10.7717/peerj-cs.4. (Visited on 07/11/2019).

[61] Kevin Page and David De Roure. "Trajectories through Social Machines". In: *Building Web Observatories Workshop*. Paris, May 2013.

[62] Manoj Parameswaran and Andrew Whinston. "Research Issues in Social Computing". In: *Journal of the Association for Information Systems* 8.6 (June 2007). ISSN: 1536-9323.

[63] C. Adam Petri and W. Reisig. "Petri Net". In: *Scholarpedia* 3.4 (2008), p. 6477. DOI: 10.4249/scholarpedia.6477.

[64] Carl Adam Petri and Wolfgang Reisig. "Petri Net". In: *Scholarpedia* 3.4 (Apr. 2008), p. 6477. ISSN: 1941-6016. DOI: 10.4249/scholarpedia.6477. (Visited on 01/24/2019).

[65] Martin Potthast, Benno Stein, and Robert Gerling. "Automatic Vandalism Detection in Wikipedia". In: *Advances in Information Retrieval*. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg, Mar. 2008, pp. 663–668. DOI: 10.1007/978-3-540-78646-7_75. (Visited on 07/04/2018).

[66] Alexander J. Quinn and Benjamin B. Bederson. "Human Computation: A Survey and Taxonomy of a Growing Field". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. New York, NY, USA: ACM, 2011, pp. 1403–1412. ISBN: 978-1-4503-0228-9. DOI: 10.1145/1978942.1979148. (Visited on 07/02/2018).

[67] Filippo Radicchi et al. "Defining and Identifying Communities in Networks". In: *Proceedings of the National Academy of Sciences* 101.9 (Mar. 2004), pp. 2658–2663. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.0400054101. (Visited on 07/09/2018).

[68] Craig W Reynolds. "Flocks, Herds and Schools: A Distributed Behavioral Model". In: *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*. 1987, pp. 25–34.

[69] Marian-Andrei Rizoiu et al. "A Tutorial on Hawkes Processes for Events in Social Media". In: *arXiv:1708.06401 [cs, stat]* (Oct. 2017). arXiv: 1708.06401 [cs, stat]. (Visited on 11/05/2020).

[70] Giacomo Rizzolatti and Laila Craighero. "The Mirror-Neuron System". In: *Annu. Rev. Neurosci.* 27 (2004), pp. 169–192.

[71] Tucmas C. Schelling. "Models of Segregation". In: *American Economic Review* 59.2 (May 1969), p. 488. ISSN: 00028282. (Visited on 02/18/2019).

[72] Jen Schradie. "Networked: The New Social Operating System By Lee Rainie and Barry Wellman MIT Press. 2012. 358 Pages." In: *Social Forces* 94.3 (Mar. 2016), e89–e89. ISSN: 0037-7732. DOI: 10.1093/sf/sou075. (Visited on 07/02/2018).

[73] M. A. Serrano, M. Boguna, and A. Vespignani. "Extracting the Multiscale Backbone of Complex Weighted Networks". In: *Proceedings of the National Academy of Sciences* 106.16 (Apr. 2009), pp. 6483–6488. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.0808904106. (Visited on 07/10/2020).

[74] Nigel Shadbolt et al. *The Theory and Practice of Social Machines*. Lecture Notes in Social Networks. Cham: Springer International Publishing, 2019. ISBN: 978-3-030-10889-2. DOI: 10.1007/978-3-030-10889-2. (Visited on 04/14/2023).

[75] Nigel R Shadbolt et al. "Towards a Classification Framework for Social Machines". In: *Proceedings of the 22Nd International Conference on World Wide Web*. New York, NY, USA: ACM, 2013, pp. 905–912.

[76] Donald W Sherburne. "A Key to Whitehead's Process and Reality". In: (1966).

[77] Paul Smart, Elena Simperl, and Nigel Shadbolt. "A Taxonomic Framework for Social Machines". In: *Social Collective Intelligence*. Ed. by Daniele Miorandi et al. Cham: Springer International Publishing, 2014, pp. 51–85. DOI: 10.1007/978-3-319-08681-1_3. (Visited on 04/17/2018).

[78] Paul R. Smart and Nigel R. Shadbolt. "Social Machines". In: *Encyclopedia of Information Science and Technology*. Ed. by Mehdi Khosrow-Pour. IGI Global, Aug. 2014, pp. 6855–6862. ISBN: 978-1-4666-5888-2. (Visited on 03/28/2018).

[79] John Kenly Smith. "The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology. Edited by Wiebe E. Bijker, Thomas P. Hughes, and Trevor Pinch. Cambridge, Mass.: MIT Press, 1987. 405 Pp. Illustrations, Charts, Notes, Bibliography, and Index. $35.00." In: *Business History Review* 62.2 (1988), pp. 341–342. ISSN: 2044-768X, 0007-6805. DOI: 10.2307/3116018. (Visited on 02/15/2019).

[80] Chaoming Song, Shlomo Havlin, and Hernán A. Makse. "Self-Similarity of Complex Networks". In: *Nature* 433.7024 (Jan. 2005), pp. 392–395. ISSN: 1476-4687. DOI: 10.1038/nature03248.

[81] Michele Starnini, Mattia Frasca, and Andrea Baronchelli. "Emergence of Metapopulations and Echo Chambers in Mobile Agents". In: *Scientific Reports* 6.1 (Oct. 2016). ISSN: 2045-2322. DOI: 10.1038/srep31834. arXiv: 1603.04789. (Visited on 10/05/2018).

[82] Ségolène Tarte et al. "Archetypal Narratives in Social Machines". In: *The ACM Web Science Conference*. New York, New York, USA: ACM Press, 2015, pp. 1–10.

[83] Ségolène M. Tarte, David De Roure, and Pip Willcox. "Working out the Plot: The Role of Stories in Social Machines". In: *Proceedings of the 23rd International Conference on World Wide Web*. WWW '14 Companion. New York, NY, USA: ACM, 2014, pp. 909–914. ISBN: 978-1-4503-2745-9. DOI: 10.1145/2567948.2578839. (Visited on 08/08/2018).

[84] Ramine Tinati et al. "The HTP Model: Understanding the Development of Social Machines". In: *Proceedings of the 22nd International Conference on World Wide Web*. Rio de Janeiro Brazil: ACM, May 2013, pp. 921–926. ISBN: 978-1-4503-2038-2. DOI: 10.1145/2487788.2488081. (Visited on 05/02/2023).

[85] T. Tiropanis et al. "The Web Science Observatory". In: *IEEE Intelligent Systems* 28.2 (Mar. 2013), pp. 100–104. ISSN: 1541-1672. DOI: 10.1109/MIS.2013.50.

[86] Thanassis Tiropanis et al. "Network Science, Web Science, and Internet Science". In: *Communications of the ACM* 58.8 (July 2015), pp. 76–82. ISSN: 00010782. DOI: 10.1145/2699416. (Visited on 08/07/2018).

[87] Michael Tomasello and Malinda Carpenter. "Shared Intentionality". In: *Developmental Science* 10.1 (Jan. 2007), pp. 121–125. ISSN: 1363755X, 14677687. DOI: 10.1111/j.1467-7687.2007.00573.x. (Visited on 05/22/2020).

[88] Michael Tomasello et al. “Understanding and Sharing Intentions: The Origins of Cultural Cognition”. In: *Behavioral and Brain Sciences* 28.5 (Oct. 2005), pp. 675–691. ISSN: 0140-525X, 1469-1825. DOI: 10.1017/S0140525X05000129. (Visited on 03/21/2023).

[89] Nathan Tompkins et al. “Testing Turing’s Theory of Morphogenesis in Chemical Cells”. In: *Proceedings of the National Academy of Sciences* 111.12 (Mar. 2014), pp. 4397–4402. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.1322005111. (Visited on 03/27/2018).

[90] A. M. Turing and F. R. S. “The Chemical Basis of Morphogenesis”. In: *Phil. Trans. R. Soc. Lond. B* 237.641 (Aug. 1952), pp. 37–72. ISSN: 0080-4622, 2054-0280. DOI: 10.1098/rstb.1952.0012. (Visited on 03/27/2018).

[91] Paola Valdivia et al. “Analyzing Dynamic Hypergraphs with Parallel Aggregated Ordered Hypergraph Visualization”. In: *IEEE Transactions on Visualization and Computer Graphics* 27.1 (2021), pp. 1–13. DOI: 10.1109/TVCG.2019.2933196.

[92] Luis von Ahn et al. “reCAPTCHA: Human-Based Character Recognition via Web Security Measures”. In: *Science* 321.5895 (2008), pp. 1465–1468. ISSN: 0036-8075. JSTOR: 20144795. (Visited on 08/08/2018).

[93] Robert Walton and David De Roure. “Modelling Web Based Socio-Technical Systems Through Formalising Possible Sequences of Human Experience”. In: *Proceedings of the 10th ACM Conference on Web Science*. WebSci ’19. New York, NY, USA: ACM, 2019, pp. 301–310. ISBN: 978-1-4503-6202-3. DOI: 10.1145/3292522.3326049. (Visited on 07/09/2019).

[94] Robert Walton and David De Roure. “Petri Nets for Modelling Communal Flocking along Paths of Possible Experience”. In: *13th ACM Web Science Conference 2021*. WebSci ’21. New York, NY, USA: Association for Computing Machinery, 2021, pp. 53–57. ISBN: 978-1-4503-8525-1. DOI: 10.1145/3462741.3466668.

[95] Duncan J Watts and Steven H Strogatz. “Collective Dynamics of ‘Small-World’ Networks”. In: *nature* 393.6684 (1998), p. 440.

[96] Alfred North Whitehead. *An Enquiry Concerning the Principles of Natural Knowledge*. Cambridge University Press, 1919.

[97] Alfred North Whitehead. *The Concept of Nature*. Vol. 62. 1919.

[98] Alfred North Whitehead and Donald W Sherburne. *Process and Reality*. Macmillan New York, NY, 1957.

[99] "Causal Processes". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2008. Metaphysics Research Lab, Stanford University, 2008.

Appendix A

Modelling software

This appendix includes a brief description of the software we developed to build and explore scenarios using various revisions of our model. Scenarios, although defined with this software through Python code, have an underlying representation as a single coloured stochastic Petri net. This software builds on the Snoopy suite of Petri net modelling software[1] [32]; although this is closed source its authors have a long history of collaborating with others and were helpful when contacted.

We found that Snoopy's interactive software was very useful when first developing concrete Petri net implementations of our idealised model. To configure this model by hand however to represent particular scenarios was tedious and error prone. We developed Python code that provides a way to define a scenario using a more accessible representation and that then 'compiles' this representation into a form recognised by the Snoopy software suite. This representation is a Python implementation of the various sets of objects and graphs used in the text. The python code uses Snoopy's scriptable component, Spike, for simulation [12].

During this work we created two variations of this software:

1. The package *occmodeller*[2] was developed to support the work published in [93]. The approach used in this paper is superseded by the approach presented in this thesis; and so is this software. It creates appealing 3d representations of traces resulting

[1] http://www-dssz.informatik.tu-cottbus.de/DSSZ/Software/Snoopy
[2] https://github.com/robwalton/occmodeller

from parts interacting through sharing experience. Interactive traces can be viewed online[3]; figure A.1 shows a static image of one of these traces,

2. A newer package[4] was used initially to support the work presented in chapter 5 and published in [94]. It has been expanded to also support the work described in chapters 6 and 7. All simulation results shown in this thesis can be recreated with this software. The random aspects of these simulations can be determined by specifying an initial random seed; this results in completely repeatable simulations and simplifies the workflow as data from simulation runs need not be stored. The repository includes instructions for reproducing these results, and an interactive Streamlit[5] web application for viewing models, configuring initial conditions, executing and viewing results.

We briefly survey the second of these packages here. Its repository contains a README file including more detailed instructions.

Specifying models. Models comprise a description of the experience generator, a description of the spatial structure of forums, initial conditions in the form of a marking, constants effecting the rate equations, and default parameters for simulation. Listing 1 contains code showing how a model might be specified and used. Here, a model is instantiated with a function taking arguments that are particular to the situation it represents. The rate constants may then be tuned and in this case a random seed set that will fully determine every random choice made during execution. Listing 2 contains the function used to specify the model. The model elements just described are evident. Some knowledge about the output file produced by Spike is not encapsulated; as shown in `EXPORT_ARGS`.

Viewing models. Models specified in this way can be explored using standard Python tools. For example, the repository includes a Jupyter notebook used to generate all the

[3] https://robwalton.github.io/posts/2019/websci19/
[4] https://github.com/robwalton/social
[5] https://streamlit.io

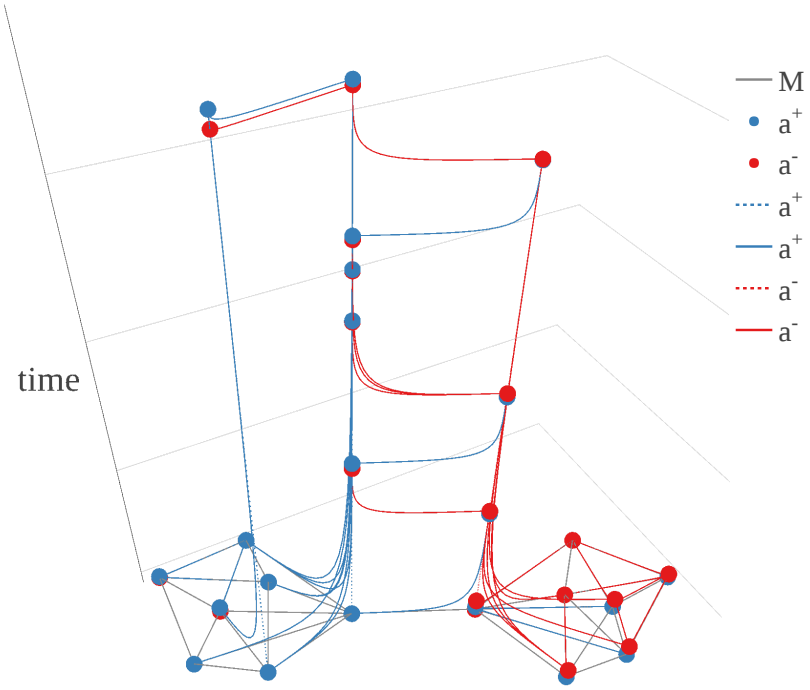


Figure A.1: Static image of a trace created with the occmodeler software. This trace shows strands changing their mind about an issue. There are two communities bridged by one pair of strands. These strands occasionally convince one another to shift their belief, but each is quickly pulled inline with their community's consensus. With time increasing upwards each sphere indicates an experience; and the lines incident indicating which historic experiences informed it.

traces in this thesis. The repository also includes a server and web application to explore models. Callables that return an executable model, such at that shown in listing 2, are added to the server to form a library of models. Figure A.2 shows the web application used to load and explore a model specified in this way. This screenshot shows a drop down list of available models on the left. The python documentation for the selected model is shown on the right. Notice that in this example this takes two arguments. These arguments may be tuned under the 'Model args' options shown on the left. Below the model's documentation is a view of the model's current experience generator graph. A very useful feature of the server is that the python code may be edited live in just the same way that its arguments can be tweaked; there is no need to restart the server. These features together form a nice interactive environment for developing models.

Tuning and executing models. The page containing the view of the model can be scrolled to reveal controls for executing it. The screenshot in figure A.3 shows the web application being used to tune and execute a model and view results. The application includes some generic tools for plotting traces showing the number of tokens in places and the firing rate of transitions.

Customising the application. The server can be started with a flag that customises its operation for Axelrod's models. This injects special code for configuring the model and for analysing and viewing the results. Other than this the workflow remains the same. These changes are implemented in a simple way that might act as a template for others looking to customise the software. The Streamlit framework we used for the software makes this simple: the code looks largely like a simple Python script, there is no need to understand Web development to make things work, and changes can be seen on the fly. Those seeking to extend things below this scripting level will be helped by the full suite of unit tests in the package and the use of standard Python development tools.

Software limitations. We were limited by this software stack to executing quite small models. This is because Spike expands models out to a non-coloured Petri net before executing. The number of strands in forums can be increased as required, but increasing the number of forums and the size of the generator graph results in very rapidly growing nets. This a downside to keeping the mathematics tractable. More customised approaches could be optimised to execute this form of simulation faster, at larger scale, and more concurrently.

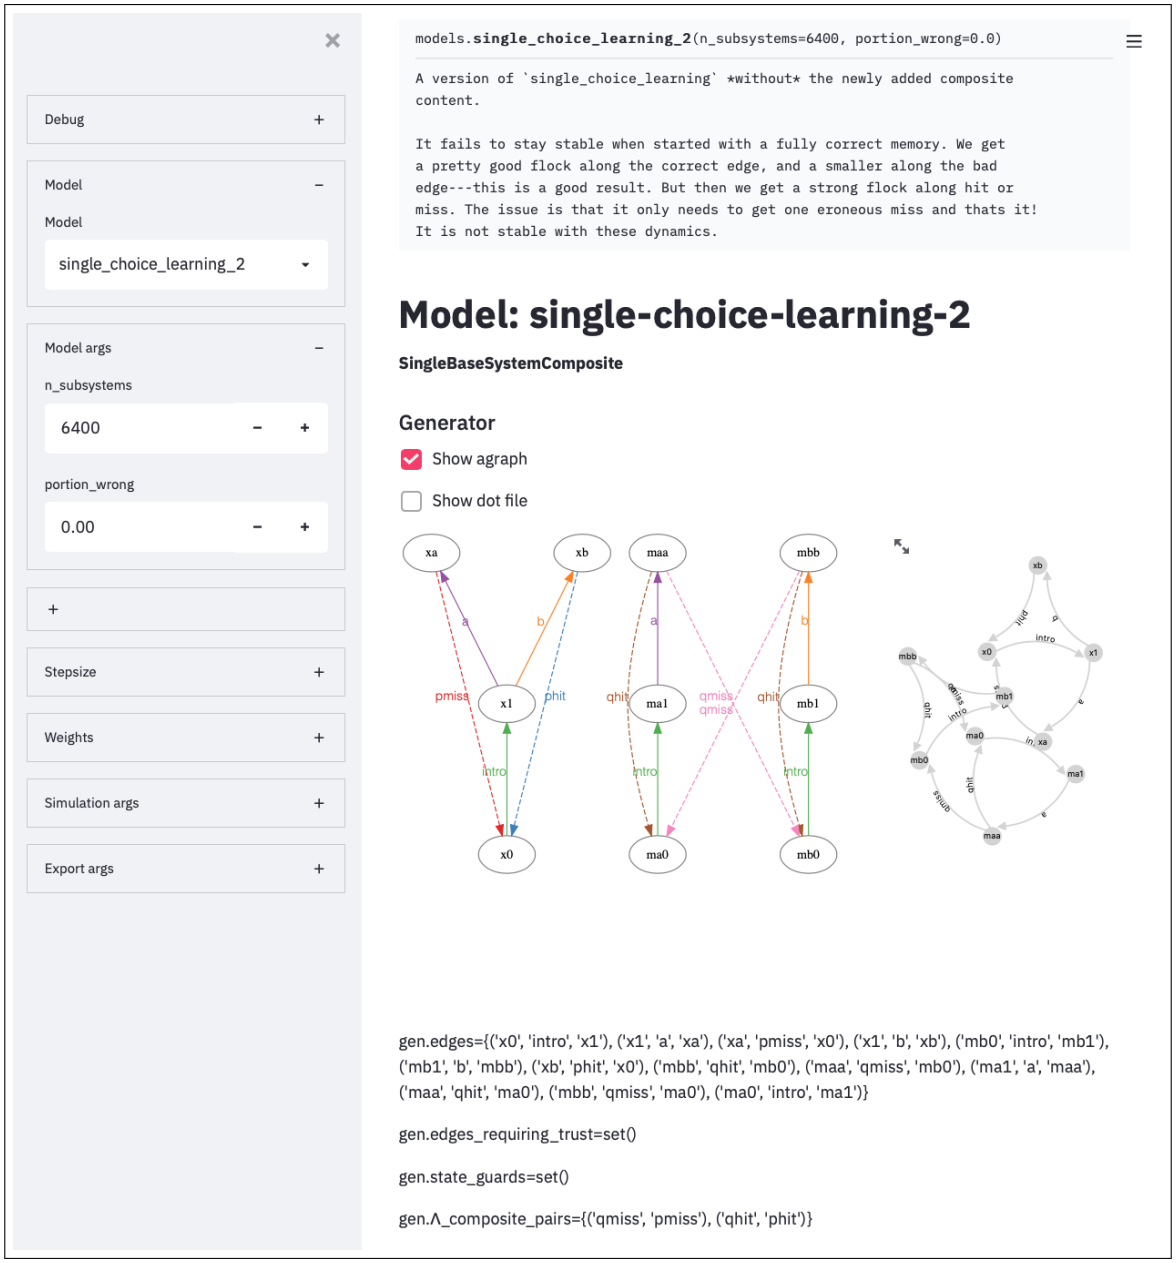


Figure A.2: Screenshot of web application showing the experience generator of a model selected from the library. This model might perform reinforcement learning with some changes to the underlying rate equations.



Figure A.3: Screenshot of web application showing model tuning and execution. This run, with the given parameters and seed, exactly reproduces the results shown in the bottom trace of figure 5.3 (although the trace colours differ).

```
import models
import basis

model = models.evenly_spread_cycles_on_grid(
    n_cycles=1, n_alphabet=2, n_rows=3, n_cols=4, n_subsystems=3600
)
model.rates = model.Rates(
    wself=0.1, wintra=1, wstick=1, delta=0.2, winter=1,
)
model.runargs.seed = 1

res = model.sim()
```

Listing 1: Code to reproduce the experiment of section 6.3.2.

```
import pyspike.call

import basis
import generators
import routers
import markings


EXPORT_ARGS = pyspike.call.ExportArgs(
    uncoloured_place_export_list=["Xe_.*", "X_.*", "Ae_.*", "Trust.*"],
    uncoloured_transition_export_list=["experience"],
)


def evenly_spread_cycles_on_grid(
    n_cycles=1, n_alphabet=2, n_rows=2, n_cols=4, n_subsystems=3600
) -> basis.MultipleBaseSystems:

    gen = generators.cycles(
        n_alphabet, n_cycles, overlapping=True, consecutive=False
    )
    routing = routers.grid_2d(n_rows, n_cols)
    model = basis.MultipleBaseSystems(
        gen=gen,
        routing=routing,
        initial_X={},
        n_subsystems=n_subsystems,
        rates=basis.MultipleBaseSystems.Rates(),
        sim_args=pyspike.call.SimArgs(0, 0.1, 100),
        export=EXPORT_ARGS,
    )
    marking, _ = markings.evenly_spread(
        n_subsystems, model.routing.S, model.gen.X
    )
    model.initial_X = marking
    return model
```

Listing 2: Code backing listing 1 that fully defines the model used in section 6.3.2.

Appendix B

Log of changes made during corrections and their rationale

I made the following changes in response to your suggested corrections. These changes are grouped under headings from the notice of minor corrections. Where I've given no explicit rationale the rationale is to address the given correction.

1. Nature of the theoretical model we're building

- Added section 1.2 section 'Developing a theory of social machines' to address this. This caused duplication in the intro chapter that needed resolution:
  - Merged non-duplicate parts of the old section 1.2 'Requirements of social machine theory' into here. Pulled the remaining discussion on Turing machines out into a new section 1.4 'Social machines and Turing machines' (as didn't want to loose it and it sets the scene quite well for what is coming).
  - Also moved in non-duplicated parts of old section 1.4.1 'Vision from requirements' into here and removed it; dropping some bits which were preemptive to the development.
  - Split old section '1.4.1 Method and choice of modelling tools' and '1.4.2 Argument and presentation' out into their own sections (because they were only parts of section 1.4 remaining), and removed duplication introduced above.

- Added 'Basis for mixing positive and interpretive social science approaches' paragraph to a new section 12.4 'Applications of theory' in the final 'Discussion' chapter. This is a better place than section 1.2 as this was not the primary purpose of the thesis (which you encouraged me to identify during the viva).

- Added 'Basis of framework for growing human-machine hybrid organisation' paragraph to the same new section. This also cluttered the primary purpose of the thesis whose discussion it showed up in.

2. Sources of assumptions

- Added section 1.3 'Theory and model domain and assumptions'. This was done with some care and ties into the text, both back to the background chapters and forward into the development chapters. I might have added this into the text itself, but pulling it forward felt right and brings the narrative to the surface better.

- Expanded the background info on shared intentionality from section 3.1.1 into the new section 3.1.2. Given its importance to the work there was not enough about this. This was my source standing in for a fuller review of the nature of social interaction and it is something we use while addressing the next batch of suggested corrections—to justify the underlying model choices, and as an example. It was a source justifying two core aspects of our model: our ability to form i) coupled states of activity, and ii) shared focus into the environment or into the content of our shared experiences.

3. Provide clarification of terminology, examples to ground terms, and connections of simulations to known dynamics. While addressing this I tended to make changes to the front of the chapters.

- To address your concern about the use of the word 'individual' I started by carefully replacing it with the term 'strand' or 'strand of experiences' throughout the thesis. This term draws attention to the nature of this part as something that exists in flux and also feels less permanent and indivisible than 'individual'.

- In the new section 1.2 just introduced, I added the idea that part of the work of building a theoretical model is to develop an ontology over a chosen domain of study. I then added better descriptions of the terms used in the guise of this. This choice kept corrections from impacting existing text anywhere but in chapter introductions.

- Added section 4.1 'Ontology' to the start of the chapter on 'strands of experience'. This defines the terms here in a way that was missing before—and it does so with reference to the newly-added source of assumptions identified in section 1.3. This section includes some smaller inline examples, but the bulk of the examples you suggested are in the next chapter on 'Collective experience in forums'.

- Added section 5.1 'Ontology' to the start of the chapter on 'Collective experience in forums'. This defines the bulk of the remaining categories of parts and their relationships used in the theory, except for 'binding' that will follow. Later chapters add some characteristics to these parts and their relationships.

- Added section 5.2 'Examples' to the start of the chapter on 'Collective experience in forums'. This includes the peas galaxy case study in some detail and examples of a sampling of more generic social machines given with reference to Maloy's crowdsourcing dimensions. The peas galaxy example is extended into the chapter on binding.

- Added section 5.5 'Discussion of experiment results' to the end of the chapter on 'Collective experience in forums'. This includes a much better explanation of how to interpret the results based around a scenario of story telling. It also explains why the flocking that we see is important to the thesis. These scenarios may be viewed as containing social machines and we apply the HTP model of Tinati et al. [84] as evidence that the dynamics we see exist in reality.

- Merged some signposting into the first few paragraphs of chapter 6 and chapter 7 indicating that the models developed within, although important to the results of the thesis, are an approximation to the main modelling approach. From the perspective

of these corrections this allows us to defer a more formal definition of 'binding' until chapter 9 when we settle on strands binding to forums of strands rather than forums binding to forums.

Continued this sign posting in chapter 8 where we return the overlapping forums representation: Tweak first sentences to point back to section 5.1.1 where we introduced overlapping forums and reiterate that the previous two chapters included a useful approximation that is equivalent in some cases and potentially useful at scale.

- Added section 9.1 'Ontology' section for the chapter on 'Persistence through ongoing binding'. This better defines 'binding' and includes real-world evidence of its existence by adding to the green peas example started in section 5.2. Added an example to this chapter as part of the following correction (4).

- Added section 10.1 'Ontology' section for the chapter on the 'Life cycle of persisting forums'.

4. Draw on "building successful online communities" literature. I could have done a lot more here but to do so would have been very involved. I proved that the theory can be connected to this body of knowledge and used part of this knowledge as evidence supporting a nuanced binding mechanism discussed in the thesis.

- Added section 9.7 'Example: normative commitment as binding' as an example at end of the 'Persistence through ongoing binding' chapter. Introduced this body of knowledge here including the process the community uses to develop it.

  Showed that a scenario we modelled in this chapter represents the observable conditions sufficient, according to a design claim of Kraut et al., to achieve the outcome of normative commitment to the community—which provides a well studied example of the phenomena we characterise as binding.

- Added 'Building successful online communities' paragraph to a new section 12.4 'Applications of theory' in the final 'Discussion' chapter. I felt this was a better

place than future work, and that I could not have defended this claim at the start of chapter 1 where I have some more generic applications motivating the need for a theory.

5. What profound things did I learn about the social interaction facilitated by network technology? In perhaps a clumsy reading of this welcome chance for a fresh look at the work, I took this suggestion pretty literally:

- Added section 12.2 'Initial theory and adjustments made during research'.
- Added section 12.3 'Social interaction facilitated by network technology' to end of final discussion chapter. This includes, as an example, a the breaking down of generic social interaction into a composite experience.

Updated abstract and conclusion. To reflect the changes made above I:

- Tweaked the start of the conclusion and added to the section 'Confidence in ontology, model and theory': i) examples showing how the ontology can be used to break down real world phenomena (many small examples in text and the bigger green peas examples and others); and ii) connections between the models we built and the behaviours that we saw running them to well-studied dynamics and a design claim that equates observable phenomena to a social outcome.
- Revisited the abstract to better reflect the contents of the thesis.

Changes not suggested in corrections

- Updated acknowledgements to include more professional and academic discussions that informed the work.
- Tweaked the paragraphs at the start of the thesis that set the scene for our interest in this as they muddied the water as to what the single purpose of the thesis is—something you drew my attention to in the viva. I had previously presented the use of social machines as something to build on to reach the goal of modelling ecosystems; instead:

> In this thesis we ask how to best represent these ecosystems of short and long-lived social situations in a way that we can reason about their operation. ... We focus this inquiry by looking particularly at phenomena of social interaction arising within the social machines ... .

- Made a pass fixing typos and wordy sentences found while reading the thesis prior to the viva.
- Removed a negative judgement about the impact of a source I used—as am not in a position to judge and it was not needed.