Chapter 9
Conclusion

This thesis has presented a framework for the information dynamics of distributed computation in complex systems. We summarise the main contributions of this work in Sect. 9.1, and suggest directions for future exploration in Sect. 9.2.

9.1 Summary of Main Contributions

The results presented in this thesis have upheld the hypothesis in Sect. 1.1 that *with the ability to describe and locally quantify distributed computation in terms of information storage, transfer and modification, we will be better able to understand distributed computation in nature and its sources of complexity*. In this section, we describe our contribution to the fundamental understanding of distributed computation in complex systems.

9.1.1 Framework for the Information Dynamics of Distributed Computation

The primary contribution of this thesis is the first *complete framework to quantify the information dynamics of distributed computation*. That is, the framework quantifies computation in terms of the component operations on information: storage, transfer and modification. The framework has a particular focus on the dynamics of these operations on a local scale in space and time within a system.

There are three key properties of the framework which underpin its novelty:

1. With an **information-theoretic basis**, the framework captures non-linear effects and is applicable to any type of dynamic process (i.e. including both discrete and continuous valued states);

J. T. Lizier, *The Local Information Dynamics of Distributed Computation in Complex Systems*, Springer Theses, DOI: 10.1007/978-3-642-32952-4_9,

2. The approach is directly relevant to the language in which distributed computation in complex systems is normally described, via the concepts of **memory, communications and processing**, since these directly map to the operations of information storage, transfer and modification;
3. The framework focuses on the **local dynamics** of these operations on information. While averaged or system-wide measures have their place in providing summarised results, the local focus in space and time is vital for understanding the nature of each measure and providing insights about system behaviour that averaged measures cannot.

In the next three Sects. (9.1.2–9.1.4), we will describe the specific contributions made regarding each individual operation on information. Many of the measures presented in this framework are original contributions; all of the measures discussed are examined on a local scale here for the first time.

Incorporating the measures together in a single framework allowed us to provide insights that would not be possible with separate investigation. For example, we demonstrated how the component operations interrelate in the computation of the next state of a given variable. We also found that establishing the context of the past history of the destination was at the heart of the perspective of distributed computation, and was critical for accurate quantification of each operation on information. Also, the use of a single framework allowed us to provide broader insights into the fundamental nature of distributed computation, e.g. regarding cellular automata as described in Sect. 9.1.5.

9.1.2 Measuring Information Storage

We described how information storage is quantified in terms of either total storage (via the existing measure *excess entropy*) or the amount of storage currently in use (via the new measure *active information storage*). We presented the first *localisation* of both measures, allowing us to contrast the insights they provide on distributed computation. We also made clear the manner in which information storage in a distributed computation can be implemented using an agent's environment as the storage medium.

9.1.3 Measuring Information Transfer

We described how information transfer is quantified using the existing measure *transfer entropy*. We introduced a number of variants to this measure in order to capture subtly different concepts; most notably, we introduced the *complete transfer entropy* to measure information transfer taking into account how the given source interacts with other causal sources in acting on the destination. We described how to measure the transfer entropy on a *local scale* in time and space, allowing us to provide insights on how its parameters should be set, describe the relationship between its variants, and demonstrate its alignment with the popularly understood notion of information transfer (see Sect. 9.1.5). We also showed how properly establishing the past history of the destination was critical to separating information storage and transfer.

Additionally, we described the differentiation between the concepts of information transfer and causal information flow. The distinctions revealed here are particularly pertinent because of the large degree of confusion surrounding these concepts in the literature. This result was only possible using our local perspective, which including localising the existing *information flow* measure.

9.1.4 Measuring Information Modification

We described how information storage and transfer are combined in the operation of information modification, and introduced the measure *separable information* to quantitatively identify non-trivial information modification events on a local scale within a system.

We also outlined how to quantify information destruction within a distributed system, introducing a measure for *information destruction*. Using localisations of both these measures, we described the distinction between the concepts of information modification and destruction.

9.1.5 Quantitative Understanding of Information Dynamics in CAs

Our framework provided the first direct quantitative evidence for several important long-held conjectures regarding the facilitation of computation in cellular automata (CAs) via emergent structures. That is, we showed that blinkers implement information storage, moving particles (gliders and domain walls) are dominant information transfer agents, and particle collisions are information modification events. This demonstrated that our quantitative framework aligned with the popularly-understood concepts of memory, communication and processing.

CAs are a critical proving ground for any theory regarding the nature of distributed computation [1, 2], and these results suggest significant implications for our fundamental understanding of distributed computation and the dynamics of complex systems. The importance of our application to CAs is underlined because many other natural and artificial systems have been observed to process information using similar emergent coherent structures [3, 4]. We also emphasise that these insights were only possible using the *local* perspective introduced for these concepts here.

The application of the framework to CAs aligned well with other methods of spatiotemporal *filtering* for complex structure (e.g. [5–9]). Obviously, filtering of coherent structure is not a new concept. However, our work is distinct in that it provides *several different profiles of the system corresponding to each type of computational structure* (and indeed one view for each information transfer channel or direction). This approach allows more refined filtering, and is unique in providing quantitative evidence regarding the computational role of these emergent structures. Furthermore, this approach is distinct in that it: provides continuous rather than discrete values (like [7] and [9]); does not follow an arbitrary spatial preference (unlike [6] and [9]) but rather the flow of time only; is automatically obtained (unlike the approaches manually crafted for specific CA rules in [5, 10]) and like [7] does not require a new filter for every CA (though the probability distribution functions must be recalculated individually). Finally, by focussing on these operations on information it highlights subtly different parts of emergent structure to other filters (generally highlighting less of the structure).

9.1.6 Measuring Computational Properties in Phase Transitions in Networks

We presented the first analysis of information dynamics in order-chaos phase transitions in networks, finding that information storage and transfer were *maximised near the critical phase* of two different network types. While the finite-sized systems exhibited approximate phase transitions, we described reasons why this interesting result might be expected to be generalised in similar[1] order-chaos phase transitions.

The results were of particular interest because of the network models chosen for study: random Boolean networks (RBNs, a model of gene regulatory networks) and a model of cascading failures. In particular, the use of RBNs was important as they had been a focus for conjecture that computational properties were maximised near the critical phase (in alignment with the *edge of chaos* hypothesis). More generally, we revealed several interesting ways in which underlying network topology drives information dynamics. Indeed, several leading authors in network science suggest that dynamics are the next frontier in this domain [11–13], so the promising results from information dynamics here are notable in suggesting it as a generally-applicable candidate for further investigation.

9.1.7 Methodology for Studying Coherent Information Structure

We also demonstrated that the maximisation of information storage and transfer near order-chaos phase transitions is not a universal result that can be expected from any type of system exhibiting ordered and chaotic variants (in particular CAs).

[1] Similar phase transitions being those with transitions controlled by a single order-chaos parameter, whether those transitions cause a discontinuous or smooth change in properties.

Instead, we observed that coherent information structure is a defining feature of complex computation and presented a *methodology for studying coherent information structure*. Importantly, our approach identifies both clear and "hidden" coherent structure in complex computation, most notably reconciling conflicting interpretations of the complexity in CA rule 22.

9.1.8 Demonstrated Application Areas for Information Dynamics

Finally, we demonstrated the utility of the framework for information dynamics in two key application areas. Together, these showed its flexibility to different data types and ability to provide useful insights to practical problems.

We presented a *method for inferring directed interregional information structure* in multivariate data sets, e.g. time-series brain imaging data. The method is unique in combining the features of directional, non-linear, model-free analysis, on a regional level, capturing the results of interaction of multiple sources, and being robust to relatively small data sets. We demonstrated the efficacy of the method by applying it to an fMRI data set, revealing a tiered information structure that correlates well with the cognitive task the subjects were performing.

We also reported the first use of the transfer entropy as a fitness function for *guiding self-organisation*. This example demonstrated that the approach can induce the emergence of useful coherent information structure in a system, which could only be revealed by examining local information dynamics.

9.2 Directions for Future Work

Certainly there are ways in which **the experiments reported here can be directly expanded** for deeper analysis. For example, we described in Sect. 6.3 several ways in which the analysis of the phase transition in RBNs could be expanded, including investigating the effects of noise. There are also several directions outlined in Sect. 7.4 in which our analysis of coherent information structure should be pursued further, in particular in examining other measures of structure in the information state-space and the relationship of coherent information structure to overall complexity.

In addition to revisiting these experiments though, new work is required to build on the achievements of this thesis in both theoretical and practical directions.

Arguably the most important direction for *theoretical work* is to investigate the **relationship between the topology of networks and their information dynamics**. As discussed in Chap. 6, most leading authors in network science suggest that [11–13] the next great leaps in that field will be produced from understanding *time-series dynamics* and how they are coupled with network topology. We described the manner in which the research landscape suggests that the dynamics of distributed computation has the potential to be the widely-anticipated framework of choice for the study of time-series dynamics in networks. Two key reasons for this are: that the approach is generic and can be applied to any type of time-series dynamics; and that the language of computation pervades description of the time-series dynamics of networks. We have demonstrated important preliminary results in applying our framework to analyse computation in networks in Chap. 6. More work is required here though, in particular a thorough investigation of how information dynamics are imparted from underlying network topology. Such investigations will establish for example whether special topologies such as small-world and scale-free networks are distinguished from others by their computational properties, and how local topological structures [14] relate to local computational capabilities.

Further theoretical work is also required to establish **how the framework for information dynamics relates to other approaches and fields**, and to clarify the measurement of several important concepts in distributed computation. A primary example here is to establish the relationship with ϵ-machines and statistical complexity from *computational mechanics* [7, 15–19]. Certainly the relationship between excess entropy and statistical complexity is well-established [18, 20–22] (indeed the excess entropy originated in computational mechanics). The perspective of *distributed* computation here would be novel in considering how information storage and transfer together related to the overall statistical complexity. This would involve focussing on the light-cone formulation of computational mechanics, which considers how the next state of an agent (and its causal descendants) depends on the causal contributors to that agent [17]. Another interesting direction for exploration would be to examine whether the framework for distributed computation can be usefully altered to apply to the underlying internal causal states of the variables in a distribution computation. Similarly, just as computational mechanics is exploring measuring information storage in quantum computation (e.g. the quantum excess entropy [23]), our framework should be extended for application to *distributed quantum computation*. Also, we note work considering quantifying *interaction structures* [24]: i.e. investigating kth order statistical dependencies between variables that cannot be reduced to dependencies between $k - 1$ of them. New work is required to quantify similar interaction structures in the context of distributed computation (i.e. examining how many source information sources are irreducibly interacting to produce an outcome). This work should also establish how this is related to the distributed operations on information (especially information modification), and whether the concept can be quantified on a local scale in space and time.

We have demonstrated a number of promising practical results in the applications of the framework to date demonstrated in this thesis (e.g. in Chap. 8). That being said, there is much more scope for quantifying computation and producing both interesting and useful insights in **applying the framework in practical settings**. Such applications will not only provide useful insights in the domain under consideration, but also build momentum for further use of the approach.

A key application area will be **computational neuroscience**. In this domain there is an abundance of time-series imaging data, and powerful capability for computational analysis, yet the road forward to specifically understand distributed computation in the brain is unclear. As explored in Sect. 8.2, information dynamics offers potential for ground-breaking insights in providing analysis of space-time information patterns, and revealing how the brain is computing. We have demonstrated utility of the approach in this context by revealing directed information structure supporting a cognitive task in Sect. 8.2; future work will include applying our method to other cognitive tasks. We will also examine other data types, particularly those with shorter time scales which allow more direct conclusions about neuronal interactions. Furthermore, we will seek to expand the application of information dynamics here, in particular in examining information storage and modification in addition to transfer. Similarly, we will examine the information dynamics on a *local scale in time* as well as space in brain-imaging data. This could lead to the identification of travelling coherent information structures in the cortex (as described in [25]). The local perspective will also address questions such as "how much information is transferred from region A to B *at time* t?", specifically revealing the information dynamics associated with particular cognitive tasks. We will explore whether this direct approach to revealing space-time information interactions can improve on inferences of shared information such as those in [26]. Additionally, we will investigate the use of the framework to infer effective networks [27] on the level of individual variables (e.g. voxels) rather than regions. Building on the use of the transfer entropy alone (e.g. [28, 29]), this could be performed using the inference method described in Appendix E to determine the sources contributing to a node's computation of its next state.

Another important application area will be in **guiding self-organisation**. As described in Sects. 2.5 and 8.3, we that a promising approach to this type of system design is the use of measures of the information dynamics of distributed computation. This is primarily because any task we wish the system to solve involves a distributed computation, so focussing our guidance on providing the fundamental building blocks of the computation is a direct way to allow that computation to emerge. We reported preliminary results indicating that evolving to maximise information transfer on local links can lead to the emergence of useful coherent information structure on a global level. Future work will include examining the use of information dynamics to guide other types of self-organised systems, e.g. collective motion or flocking. An interesting domain will be examining how information dynamics can guide network topology, for example whether our understanding of the information dynamics of cascading failures in Chap. 6 can be applied to design power grids to avoid these events. From a theoretical perspective, this scope of our work in guided self-organisation needs to be significantly expanded to consider information storage and modification also, and to establish what types of properties can usefully produced by processes of evolution or adaptation to maximise each of them. More importantly, the approach needs to investigate how the information dynamics can be used *together* to guide the emergence of universal computation.[2] Intricate tasks will require such arbitrarily-complex computation (facilitated by bidirectional

[2] Indeed, whether these measures can be used to determine the capability of a distributed system for universal computation, or capability of other levels of computational complexity [16], needs to be established.

information transfer, storage structures and modification events) as distinct from computation that exhibits only one type of operation.

References

1. M. Mitchell, in *Computation in cellular automata selected review*, ed. by T. Gramss, S. Bornholdt, M. Gross, M. Mitchell, T. Pellizzari, non-standard computation, (Verlagsgesellschaft, Weinheim, 1998), pp. 95–140
2. J. Von Neumann, *Theory of self-reproducing automata* ed. by A.W. Burks (University of Illinois Press, Urbana, 1966)
3. D. Peak, J.D. West, S.M. Messinger, K.A. Mott, Evidence for complex, collective dynamics and emergent, distributed computation in plants. Proc. Nat. Acad. Sci. USA. **101**(4), 918–922 (2004)
4. M. Mitchell, J.P. Crutchfield, P.T. Hraber, Evolving cellular automata to perform computations: mechanisms and impediments. Physica D **75**, 361–391 (1994)
5. P. Grassberger, New mechanism for deterministic diffusion. Phys. Rev. A **28**(6), 3666 (1983)
6. J.E. Hanson, J.P. Crutchfield, The attractor-basin portait of a cellular automaton. J. Stat. Phys. **66**, 1415–1462 (1992)
7. C.R. Shalizi, R. Haslinger, J.-B. Rouquier, K.L. Klinkner, C. Moore, Automatic filters for the detection of coherent structure in spatiotemporal systems. Phys. Rev. E **73**(3), 036104 (2006)
8. A. Wuensche, Classifying cellular automata automatically: finding gliders, filtering, and relating space-time patterns, attractor basins, and the Z parameter. Complex **4**(3), 47–66 (1999)
9. T. Helvik, K. Lindgren, M.G. Nordahl, Local information in one-dimensional cellular automata, in *Proceedings of the International Conference on Cellular Automata for Research and Industry*, ed. by P.M. Sloot, B. Chopard, A.G. Hoekstra. Amsterdam, ser. Lecture Notes in Computer Science, vol. 3305 (Springer, Berlin, 2004), pp. 121–130
10. P. Grassberger, Information content and predictability of lumped and distributed dynamical systems. Phys. Scr. **40**(3), 346 (1989)
11. D.J. Watts, *Six Degrees: The Science of a Connected Age* (Norton, New York, 2003)
12. A.-L. Barabási, Scale-free networks: a decade and beyond. Science **325**(5939) 412–413 (2009)
13. M. Mitchell, Complex systems: Network thinking. Artif. Intell. **170**(18), 1194–1212 (2006)
14. M. Piraveenan, M. Prokopenko, A.Y. Zomaya, Local assortativeness in scale-free networks. Europhys. Lett. **84**(2), 28002 (2008)
15. J.P. Crutchfield, K. Young, Inferring statistical complexity. Phys. Rev. Lett. **63**(2), 105 (1989)
16. J.P. Crutchfield, The calculi of emergence: computation, dynamics and induction. Physica D **75**(1–3), 11–54 (1994)
17. C.R. Shalizi, *Causal architecture, complexity and self-organization in time series and cellular automata*, Ph.D. Dissertation, University of Wisconsin-Madison, 2001
18. C.R. Shalizi, J.P. Crutchfield, Computational mechanics: pattern and prediction, structure and simplicity. J. Stat. Phys. **104**, 817–879 (2001)
19. C.R. Shalizi, K.L. Shalizi, R. Haslinger, Quantifying self-organization with optimal predictors. Phys. Rev. Lett. **93**(11), 118701 (2004)
20. J.P. Crutchfield, C.J. Ellison, J.R. Mahoney, Time's barbed arrow: irreversibility, crypticity, and stored information. Phys. Rev. Lett. **103**(9), 094101 (2009)
21. C. Ellison, J. Mahoney, J. Crutchfield, Prediction, retrodiction, and the amount of information stored in the present. J. Stat. Phys. **136**(6), 1005–1034 (2009)
22. K. Wiesner, M. Gu, E. Rieper, V. Vedral, Information erasure lurking behind measures of complexity, 2009, arXiv:0905.2918v1. url: http://arxiv.org/abs/0905.2918
23. J.P. Crutchfield, K. Wiesner, Intrinsic quantum computation. Phys. Lett. A **372**(4), 375–380 (2008)

24. T. Kahle, E. Olbrich, J. Jost, N. Ay, Complexity measures from interaction structures. Phys. Rev. E **79**(2), 026201 (2009)
25. P. Gong, C. van Leeuwen, Distributed dynamical computation in neural circuits with propagating coherent activity patterns. PLoS Comput. Biol. **5**(12), e1000611 (2009)
26. C.S. Soon, M. Brass, H.-J. Heinze, J.-D. Haynes, Unconscious determinants of free decisions in the human brain. Nat. Neurosci. **11**(5), 543–545 (2008)
27. K.J. Friston, Functional and effective connectivity in neuroimaging: a synthesis. Hum. Brain Mapp. **2**, 56–78 (1994)
28. C.J. Honey, R. Kotter, M. Breakspear, and O. Sporns, Network structure of cerebral cortex shapes functional connectivity on multiple time scales, in Proc. Nat. Acad. Sci. **104**(24), 10240–10245 (2007)
29. J.T. Lizier, M. Piraveenan, D. Pradhana, M. Prokopenko, L. S. Yaeger, Functional and structural topologies in evolved neural networks, in *Proceedings of the European Conference on Artificial Life (ECAL)*, ed. by G. Kampis, I. Karsai, E. Szathmáry. Budapest, Hungary, ser. Lecture Notes in Computer Science, vol. 5777 (Springer, Berlin, 2011), pp. 140–147