1

A framework for the local information dynamics of distributed computation in complex systems

Joseph T. Lizier[1,2], Mikhail Prokopenko[1] and Albert Y. Zomaya[2]

[1]CSIRO Computational Informatics, Locked Bag 17, North Ryde, NSW 1670, Australia
[2]School of Information Technologies, The University of Sydney, NSW 2006, Australia

Summary. The nature of distributed computation has often been described in terms of the component operations of universal computation: *information storage*, *transfer* and *modification*. We review the first complete framework that quantifies each of these individual information dynamics on a local scale within a system, and describes the manner in which they interact to create non-trivial computation where "the whole is greater than the sum of the parts". We describe the application of the framework to cellular automata, a simple yet powerful model of distributed computation. This is an important application, because the framework is the first to provide quantitative evidence for several important conjectures about distributed computation in cellular automata: that blinkers embody information storage, particles are information transfer agents, and particle collisions are information modification events. The framework is also shown to contrast the computations conducted by several well-known cellular automata, highlighting the importance of information coherence in complex computation. The results reviewed here provide important quantitative insights into the fundamental nature of distributed computation and the dynamics of complex systems, as well as impetus for the framework to be applied to the analysis and design of other systems.

1.1 Introduction

The nature of distributed computation has long been a topic of interest in complex systems science, physics, artificial life and bioinformatics. In particular, emergent complex behavior has often been described from the perspective of computation within the system (Mitchell 1998b,a) and has been postulated to be associated with the capability to support universal computation (Langton 1990; Wolfram 1984c; Casti 1991).

In all of these relevant fields, distributed computation is generally discussed in terms of "memory", "communication", and "processing". Memory refers to the storage of information by some variable to be used in the future of its time-series process. It has been investigated in coordinated motion in modular robots (Prokopenko et al. 2006), in the dynamics of inter-event distribution times (Goh and Barabási 2008), and in synchronization between coupled systems (Morgado et al. 2007). Communication

refers to the transfer of information between one variable's time-series process and another; it has been shown to be of relevance in neuroscience (Wibral et al. 2011; Lindner et al. 2011; Marinazzo et al. 2012) and in other biological systems (e.g. dipole-dipole interaction in microtubules (Brown and Tuszynski 1999), and in signal transduction by calcium ions (Pahle et al. 2008)), social animals (e.g. schooling behavior in fish (Couzin et al. 2006)), agent-based systems (e.g. the influence of agents over their environments (Klyubin et al. 2005), and in inducing emergent neural structure (Lungarella and Sporns 2006)). Processing refers to the combination of stored and/or transmitted information into a new form; it has been discussed in particular for biological neural networks and models thereof (Kinouchi and Copelli 2006; Atick 1992; Sánchez-Montañés and Corbacho 2002; Yamada and Aihara 1994) (where it has been suggested as a potential biological driver), and also regarding collision-based computing (e.g. (Jakubowski et al. 1997; Adamatzky 2002), and including soliton dynamics and collisions (Edmundson and Enns 1993)).

Significantly, these terms correspond to the component operations of Turing universal computation: **information storage**, **information transfer** (or transmission) and **information modification**. Yet despite the obvious importance of these **information dynamics**, until recently there was no framework for either quantifying them individually or understanding how they interact to give rise to distributed computation.

Here, we review the first complete framework (Lizier et al. 2007, 2008c, 2012c, 2010, 2012b; Lizier and Prokopenko 2010; Lizier 2013) which quantifies each of these information dynamics or component operations of computation within a system, and describes how they inter-relate to produce distributed computation. We refer to the *dynamics* of information for two key reasons here. First, this approach describes the composition of information in the *dynamic state update* for the time-series process of each variable within the system, in terms of how information is stored, transferred and modified. This perspective of state updates brings an important connection between information theory and dynamical systems. Second, the approach focuses on the *dynamics* of these operations on information on a *local scale* in space and time within the system. This focus on the local scale is an important one. Several authors have suggested that a complex system is better characterized by studies of its local dynamics than by averaged or overall measures (Shalizi et al. 2006; Hanson and Crutchfield 1992), and indeed here we believe that quantifying and understanding distributed computation will necessitate studying the information dynamics and their interplay on a local scale in space and time. Additionally, we suggest that the quantification of the individual information dynamics of computation provides three *axes of complexity* within which to investigate and classify complex systems, allowing deeper insights into the variety of computation taking place in different systems.

An important focus for discussions on the nature of distributed computation have been cellular automata (CAs) as model systems offering a range of dynamical behavior, including supporting complex computations and the ability to model complex systems in nature (Mitchell 1998b). We review the application of this framework to CAs here because there is very clear *qualitative* observation of emergent structures representing information storage, transfer and modification therein (Langton 1990; Mitchell 1998b). CAs are a critical proving ground for any theory on the nature of distributed

computation: significantly, Von Neumann was known to be a strong believer that "a general theory of computation in 'complex networks of automata' such as cellular automata would be essential both for understanding complex systems in nature and for designing artificial complex systems" (Mitchell (1998b) describing Von Neumann (1966)).

Information theory provides the logical platform for our investigation, and we begin with a summary of the main information-theoretic concepts required. We provide additional background on the qualitative nature of distributed computation in CAs, highlighting the opportunity which existed for our framework to provide quantitative insights. Subsequently, we consider each component operation of universal computation in turn, and describe how to quantify it locally in a spatiotemporal system. As an application, we review the measurement of each of these information dynamics at every point in space-time in several important CAs. We show that our framework provided the first complete quantitative evidence for a well-known set of conjectures on the emergent structures dominating distributed computation in CAs: that blinkers provide information storage, particles provide information transfer, and particle collisions facilitate information modification. Furthermore, we describe the manner in which our results implied that the *coherence* of information may be a defining feature of complex distributed computation. Our findings are significant because these emergent structures of computation in CAs have known analogues in many physical systems (e.g. solitons and biological pattern formation processes, coherent waves of motion in flocks), and as such this work will contribute to our fundamental understanding of the nature of distributed computation and the dynamics of complex systems. We finish by briefly reviewing the subsequent application of the framework to various complex systems, including in analyzing flocking behavior and in a computational neuroscience setting.

1.2 Information-theoretic preliminaries

Information theory (Shannon 1948; Cover and Thomas 1991; MacKay 2003) is an obvious tool for quantifying the information dynamics involved in distributed computation. In fact, information theory has already proven to be a useful framework for the design and analysis of complex self-organized systems (e.g. see (Prokopenko et al. 2009)).

We begin by reviewing several necessary information theoretic quantities, including several measures explicitly defined for use with time-series processes. We also describe *local* information-theoretic quantities - i.e. the manner in which information-theoretic measures can be used to describe the information content associated with single observations.

1.2.1 Fundamental quantities

The fundamental quantity is the Shannon **entropy**, which represents the uncertainty associated with any measurement x of a random variable X (using units in bits):

$$H_X = -\sum_x p(x) \log_2 p(x). \tag{1.1}$$

The **joint entropy** of two (or more) random variables *X* and *Y* is a generalization to quantify the uncertainty of the joint distribution of *X* and *Y*:

$$H_{X,Y} = -\sum_{x,y} p(x,y) \log_2 p(x,y). \tag{1.2}$$

The **conditional entropy** of *X* given *Y* is the average uncertainty that remains about *x* when *y* is known:

$$H_{X|Y} = -\sum_{x,y} p(x,y) \log_2 p(x \mid y). \tag{1.3}$$

The **mutual information** (MI) between *X* and *Y* measures the average reduction in uncertainty about *x* that results from learning the value of *y*, or vice versa:

$$I_{X;Y} = \sum_{x,y} p(x,y) \log_2 \frac{p(x,y)}{p(x)p(y)}. \tag{1.4}$$

$$I_{X;Y} = H_X - H_{X|Y} = H_Y - H_{Y|X}. \tag{1.5}$$

One can also describe the MI as measuring the information contained in X about Y (or vice versa).

The **conditional mutual information** between *X* and *Y* given *Z* is the mutual information between *X* and *Y* when *Z* is known:

$$I_{X;Y|Z} = H_{X|Z} - H_{X|Y,Z} \tag{1.6}$$

$$= H_{Y|Z} - H_{Y|X,Z}. \tag{1.7}$$

Importantly, the conditional MI $I_{X;Y|Z}$ can be larger or smaller than the unconditioned $I_{X;Y}$ (MacKay 2003); it is reduced by redundant information held by Y and Z about X, and increased by synergy between Y and Z about X (e.g. where X is the result of an exclusive-OR or XOR operation between Y and Z).

1.2.2 Measures for time-series processes

Next, we describe several measures which are explicitly defined for time-series processes X.

The **entropy rate** is the limiting value of the rate of change of the joint entropy over *k* consecutive values of a time-series process X, (i.e. measurements $\mathbf{x}_n^{(k)} = \{x_{n-k+1}, \ldots, x_{n-1}, x_n\}$, up to and including time step n, of the random variable $\mathbf{X}_n^{(k)} = \{X_{n-k+1}, \ldots, X_{n-1}, X_n\}$), as *k* increases (Cover and Thomas 1991; Crutchfield and Feldman 2003):

$$H_{\mu X} = \lim_{k \to \infty} \frac{H_{\mathbf{X}_n^{(k)}}}{k} = \lim_{k \to \infty} H'_{\mu X}(k), \tag{1.8}$$

$$H'_{\mu X}(k) = \frac{H_{\mathbf{X}_n^{(k)}}}{k}, \tag{1.9}$$

where the limit exists. Note that $\mathbf{X}_n^{(k)}$ is a k-dimensional *embedding vector* of the *state* of X (Takens 1981). A related definition is given by the limiting value of the conditional entropy of the next value of X (i.e. measurements x_{n+1} of the random variable X_{n+1}) given knowledge of the previous k values of X (i.e. measurements $\mathbf{x}_n^{(k)}$ of the random variable $\mathbf{X}_n^{(k)}$):

$$H_{\mu X} = \lim_{k\to\infty} H_{X_{n+1}|\mathbf{X}_n^{(k)}} = \lim_{k\to\infty} H_{\mu X}(k), \tag{1.10}$$

$$H_{\mu X}(k) = H_{\mathbf{X}_{n+1}^{(k+1)}} - H_{\mathbf{X}_n^{(k)}}, \tag{1.11}$$

again, where the limit exists. This can also be viewed as the uncertainty of the next state $\mathbf{x}_{n+1}^{(k)}$ given the previous state $\mathbf{x}_n^{(k)}$, since x_{n+1} is the only non-overlapping quantity in $\mathbf{x}_{n+1}^{(k)}$ which is capable of carrying any conditional entropy. Cover and Thomas (1991) point out that these two quantities correspond to two subtly different notions. These authors go on to demonstrate that for stationary processes X, the limits for the two quantities $H'_\mu(X)$ and $H_\mu(X)$ exist (i.e. the average entropy rate converges) and are equal. For our purposes in considering information dynamics, we are interested in the latter formulation $H_\mu(X)$, since it explicitly describes how one random variable X_{n+1} is related to the previous instances $\mathbf{X}_n^{(k)}$.

Grassberger (1986b) first noticed that a slow approach of the entropy rate to its limiting value was a sign of complexity. Formally, Crutchfield and Feldman (2003) use the conditional entropy form of the entropy rate (1.10)[1] to observe that at a finite block size k, the difference $H_{\mu X}(k) - H_{\mu X}$ represents the information carrying capacity in size k-blocks that is due to correlations. The sum over all k gives the total amount of structure in the system, quantified as the **effective measure complexity** or **excess entropy** (measured in bits):

$$E_X = \sum_{k=0}^{\infty} \left[H_{\mu X}(k) - H_{\mu X} \right]. \tag{1.12}$$

The excess entropy can also be formulated as the mutual information between the semi-infinite past and semi-infinite future of the system:

$$E_X = \lim_{k\to\infty} I_{\mathbf{X}_n^{(k)};\mathbf{X}_{n+1}^{(k^+)}}, \tag{1.13}$$

where $\mathbf{X}_{n+1}^{(k^+)} = \{X_{n+1}, X_{n+2}, \ldots, X_{n+k}\}$ is the random variable (with measurements $\mathbf{x}_{n+1}^{(k^+)} = \{x_{n+1}, x_{n+2}, \ldots, x_{n+k}\}$) referring to the k future values of the process X (from time step $n+1$ onwards). This interpretation is known as the *predictive information* (Bialek et al. 2001), as it highlights that the excess entropy captures the information in a process' past which is relevant to predicting its future.

1.2.3 Local information-theoretic measures

Finally, we note that the aforementioned information-theoretic quantities are *averages* over all of the observations used to compute the relevant probability distribution func-

[1] $H_{\mu X}(k)$ here is equivalent to $h_\mu(k+1)$ in (Crutchfield and Feldman 2003).

tions (PDFs). One can also write down *local* or pointwise measures for each of these quantities, representing their value for one specific observation or configuration of the variables (x, y, z) being observed. The average of a local quantity over all observations is of course the relevant average information-theoretic measure.

Primarily, the **Shannon information content** or **local entropy** of an outcome x of measurement of the variable X is (MacKay 2003):

$$h(x) = -\log_2 p(x). \tag{1.14}$$

Note that by convention we use lower-case symbols to denote local information-theoretic measures throughout this chapter. The quantity $h(x)$ is simply the information content attributed to the specific symbol x, or the information required to predict or uniquely specify that value. Less probable outcomes x have higher information content than more probable outcomes, and we have $h(x) \geq 0$. Specifically, the Shannon information content of a given symbol x is the *code-length* for that symbol in an optimal encoding scheme for the measurements X, i.e. one that produces the minimal expected code length.[2]

Now, note that although the PDF $p(x)$ is *evaluated* for $h(x)$ locally at the given observation x, it is *defined* using all of the available (non-local) observations of the variable X which would go into evaluation of the corresponding $H(X)$. That is to say, we define a certain PDF $p(x)$ from all given measurements of a variable X: we can measure local entropies $h(x)$ by evaluating $p(x)$ for a given observation x, or we can measure average entropies $H(X)$ from the whole function $p(x)$, and indeed we have $H(X) = \langle h(x) \rangle$ when the expectation value is taken over $p(x)$.

Similarly, we have the **local conditional entropy** $h(x \mid y) = -\log_2 p(x \mid y)$ with $H(X \mid Y) = \langle h(x \mid y) \rangle$.

Next, the **local mutual information** (Fano 1961) for a specific observation (x, y) is the information held in common between the specific values x and y:

$$i(x; y) = h(x) - h(x \mid y), \tag{1.15}$$

$$= \log_2 \frac{p(x \mid y)}{p(x)}. \tag{1.16}$$

The local mutual information is the difference in code lengths between coding the value x in isolation (under the optimal encoding scheme for X), or coding the value x given y (under the optimal encoding scheme for X given Y). Similarly, we have the **local conditional mutual information**:

$$i(x; y \mid z) = h(x \mid z) - h(x \mid y, z), \tag{1.17}$$

$$= \log_2 \frac{p(x \mid y, z)}{p(x \mid z)}. \tag{1.18}$$

[2] This "optimal code-length" may specify non-integer choices; full discussion of the implications here, practical issues in selecting integer code-lengths, and block-coding optimisations are contained in (Cover and Thomas 1991, Chapter 5).

Indeed, the form of $i(x;y)$ and $i(x;y \mid z)$ are derived directly from four postulates by Fano (1961, ch. 2): once-differentiability, similar form for conditional MI, additivity (i.e. $i(\{y_n, z_n\}; x_n) = i(y_n; x_n) + i(z_n; x_n \mid y_n)$), and separation for independent ensembles. This derivation means that $i(x;y)$ and $i(x;y \mid z)$ are uniquely specified, up to the base of the logarithm. Of course, we have $I(X;Y) = \langle i(x;y) \rangle$ and $I(X;Y \mid Z) = \langle i(x;y \mid z) \rangle$, and like $I(X;Y)$ and $I(X;Y \mid Z)$, the local values are symmetric in x and y.

Importantly, $i(x;y)$ may be positive or negative, meaning that one variable can either positively inform us or actually *misinform* us about the other. An observer is misinformed where, conditioned on the value of y the observed outcome of x was *relatively* unlikely as compared to the unconditioned probability of that outcome (i.e. $p(x \mid y) < p(x)$). Similarly, $i(x;y \mid z)$ can become negative where $p(x \mid y,z) < p(x \mid z)$.

Applied to time-series data, local measures tell us about the *dynamics* of information in the system, since they vary with the specific observations in time, and local values are known to reveal more details about the system than the averages alone (Shalizi 2001; Shalizi et al. 2006).

1.3 Cellular automata

1.3.1 Introduction to Cellular Automata

Cellular automata (CA) are discrete dynamical systems consisting of an array of cells which each synchronously update their discrete value as a function of the values of a fixed number of spatially neighboring cells using a uniform rule. Although the behavior of each individual cell is very simple, the (non-linear) interactions between all cells can lead to very intricate global behavior, meaning CAs have become a classic example of self-organized complex behavior. Of particular importance, CAs have been used to model real-world spatial dynamical processes, including fluid flow, earthquakes and biological pattern formation (Mitchell 1998b).

The neighborhood of a cell used as inputs to its update rule at each time step is usually some regular configuration. In 1D CAs, this means the same range *r* of cells on each side and including the current value of the updating cell. One of the simplest variety of CAs – 1D CAs using binary values, deterministic rules and one neighbor on either side ($r = 1$) – are known as the *Elementary CAs*, or *ECAs*. Example evolutions of ECAs from random initial conditions may be seen in Fig. 1.2a and Fig. 1.6a. For more complete definitions of CAs, including the definition of the Wolfram rule number convention for specifying update rules, see Wolfram (2002).

Wolfram (1984c, 2002) sought to classify the asymptotic behavior of CA rules into four classes: I. Homogeneous state; II. Simple stable or periodic structures; III. Chaotic aperiodic behavior; and IV. Complicated localized structures, some propagating. Much conjecture remains as to whether these classes are quantitatively distinguishable, e.g. see Gray (2003), however they do provide an interesting analogy (for discrete-state and time) to our knowledge of dynamical systems, with classes I and II representing

ordered behavior, class III representing chaotic behavior, and class IV representing complex behavior and considered as lying between the ordered and chaotic classes.

More importantly though, the approach seeks to characterize complex behavior in terms of emergent structure in CAs, regarding *gliders*, *particles* and *domains*. Qualitatively, a domain may described as a set of background configurations in a CA, for which any given configuration will update to another such configuration in the set in the absence of any disturbance. Domains are formally defined within the framework of computational mechanics (Hanson and Crutchfield 1992) as spatial process languages in the CA. Particles are qualitatively considered to be moving elements of coherent spatiotemporal structure. Gliders are particles which repeat periodically in time while moving spatially (repetitive non-moving structures are known as *blinkers*). Formally, particles are defined within the framework of computational mechanics as a boundary between two domains (Hanson and Crutchfield 1992); as such, they can also be termed as domain walls, though this is typically used with reference to aperiodic particles.

These emergent structures are more clearly visible when the CA is filtered in some way. Early filtering methods were hand-crafted for specific CAs (relying on the user knowing the pattern of background domains) (Grassberger 1983, 1989), while later methods can be automatically applied to any given CA. These include: ϵ-machines (Hanson and Crutchfield 1992), input entropy (Wuensche 1999), local information (Helvik et al. 2004), and local statistical complexity (Shalizi et al. 2006). All of these filtering techniques produce a *single* filtered view of the structures in the CA: our measures of local information dynamics will present several filtered views of the distributed computation in a CA, separating each operation on information. The ECA examples analyzed in this chapter are introduced in Section 1.3.3.

1.3.2 Computation in Cellular Automata

CAs can be interpreted as undertaking distributed computation: it is clear that "data represented by initial configurations is processed by time evolution" (Wolfram 1984c). As such, computation in CAs has been a popular topic for study (see Mitchell (1998b)), with a particular focus in observing or constructing (Turing) universal computation in certain CAs. An ability for universal computation is defined to be where "suitable initial configurations can specify arbitrary algorithm procedures" in the computing entity, which is capable of "evaluating any (computable) function" (Wolfram 1984c). Wolfram (1984c,a) conjectured that all class IV complex CAs were capable of universal computation. He went on to state that prediction in systems exhibiting universal computation is limited to explicit simulation of the system, as opposed to the availability of any simple formula or "short-cut", drawing parallels to the halting problem for universal Turing machines (Wolfram 1984c,a) which are echoed by Langton (1990) and Casti (1991). (Casti extended the analogy to undecidable statements in formal systems, i.e. Gödel's Theorem). The capability for universal computation has been proven for several CA rules, through the design of rules generating elements to (or by identifying elements which) specifically provide the component operations required for universal computation: information storage, transmission and modification. Examples here in-

clude most notably the Game of Life (Conway 1982) and ECA rule 110 (Cook 2004); also see Lindgren and Nordahl (1990) and discussions by Mitchell (1998b).

The focus on elements providing information storage, transmission and modification pervades discussion of all types of computation in CAs, e.g. (Adamatzky 2002; Jakubowski et al. 2001). Wolfram (1984a) claimed that in class III CAs information propagates over an infinite distance at a (regular) finite speed, while in class IV CAs information propagates at an irregular speed over an infinite range. Langton (1990) hypothesized that complex behavior in CAs exhibited the three component operations required for universal computation. He suggested that the more chaotic a system becomes the more information transmission increases, and the more ordered a system becomes the more information it stores. Complex behavior was said to occur at a phase transition between these extremes requiring an intermediate level of both information storage and transmission: if information propagates too well, *coherent information decays into noise*. Langton elaborates that transmission of information means that the "dynamics must provide for the propagation of information in the form of signals over arbitrarily long distances", and suggests that particles in CAs form the basis of these signals. To complete the qualitative identification of the elements of computation in CAs, he also suggested that blinkers formed the basis of information storage, and collisions between propagating (particles) and static structures (blinkers) "can modify either stored or transmitted information in the support of an overall computation". Rudimentary attempts were made at quantifying the average information transfer (and to some extent information storage), via mutual information (although as discussed later this is a symmetric measure not capturing directional transfer). Recognizing the importance of the emergent structures to computation, several examples exist of attempts to automatically identify CA rules which give rise to particles and gliders, e.g. (Wuensche 1999; Eppstein 2002), suggesting these to be the most interesting and complex CA rules.

Several authors however criticize the aforementioned approaches of attempting to classify CAs in terms of their generic behavior or "bulk statistical properties", suggesting that the wide range of differing dynamics taking place across the CA makes this problematic (Hanson and Crutchfield 1992; Mitchell 1998b). Gray (2003) suggests that there there may indeed be classes of CAs capable of more complex computation than universal computation alone. More importantly, Hanson and Crutchfield (1992) criticize the focus on universal computational ability as drawing away from the ability to identify "generic computational properties", i.e. a lack of ability for universal computation does not mean a CA is not undertaking any computation at all. Alternatively, these studies suggest that analyzing the rich space-time dynamics *within* the CA is a more appropriate focus. As such, these and other studies have analyzed the *local* dynamics of intrinsic or other specific computation, focusing on particles facilitating the transfer of information and collisions facilitating the information processing. Noteworthy examples here include: the method of applying filters from the domain of computational mechanics by Hanson and Crutchfield (1992); and analysis using such computational mechanics filters of CA rules selected via evolutionary computation to perform classification tasks by Mitchell et al. (1994, 1996). Related are studies which deeply investigate the nature of particles and their interactions, e.g. particle types and

their interaction products identified for particular CAs (Mitchell et al. 1996; Boccara et al. 1991; Martinez et al. 2006), and rules established for their interaction products by Hordijk et al. (2001).

Despite such interest, until recently there was no complete framework that locally quantifies the individual information dynamics of distributed computation within CAs or other systems. In this review, we describe how the information dynamics can be locally quantified within the spatiotemporal structure of a CA. In particular, we describe the dynamics of how information storage and information transfer interact to give rise to information processing. Our approach is not to quantify computation or overall complexity, nor to identify universal computation or determine what is being computed; it is simply intended to quantify the component operations in space-time.

1.3.3 Examples of distributed computation in CAs

In this chapter, we review analysis of the computation carried out by several important ECA rules:

- Class IV complex rules 110 and 54 (Wolfram 2002) (see Fig. 1.4a and Fig. 1.2a), both of which exhibit a number of glider types and collisions. ECA rule 110 is the only proven computationally universal ECA rule (Cook 2004).
- Rules 22 and 30 as representative class III chaotic rules (Wolfram 2002) (see rule 22 in Fig. 1.6a);
- Rules 18 as a class III rule which contains domain walls against a chaotic background domain (Wolfram 1984b; Hanson and Crutchfield 1992).

These CAs each carry out an *intrinsic* computation of the evolution to their ultimate attractor and phase on it (see Wuensche (1999) for a discussion of attractors and state space in finite-sized CAs). That is to say, we view the attractor as the end point of an intrinsic computation by the CA – the dynamics of the transient to the attractor may contain information storage, transfer and modification, while the dynamics on the attractor itself can only contain information storage (since the attractor is either a fixed point or periodic process here). As such, we are generally only interested in studying computation during the transient dynamics here, as non-trivial computation processes.

We also examine a CA carrying out a "human-understandable" computational task. Rule ϕ_{par} is a 1D CA with range $r = 3$ (the 128-bit Wolfram rule number 0xfeedffdec1aaeec0eef000a0e1a020a0) that was evolved by Mitchell et al. (1994, 1996) to classify whether the initial CA configuration had a majority of 1's or 0's by reaching a fixed-point configuration of all 1's for the former or all 0's for the latter. This CA rule achieved a success rate above 70% in its task. An example evolution of this CA can be seen in Fig. 1.5a. The CA appears to carry out this computation using blinkers and domains for information storage, gliders for information transfer and glider collisions for information modification. The CA exhibits an initial emergence of domain regions of all 1's or all 0's storing information about local high densities of either value. Where these domains meet, a checkerboard domain propagates slowly (1 cell per time step) in both directions, transferring information of a *soft* uncertainty in this part of the CA. Some "certainty" is provided where the glider of the leading edge of

a checkerboard encounters a blinker boundary between 0 and 1 domains, which stores information about a *hard* uncertainty in that region of the CA. This results in an information modification event where the domain on the opposite side of the blinker to the incoming checkerboard is concluded to represent the higher density state, and is allowed to propagate over the checkerboard. This new information transfer associated with local decision of which is the higher density state has evolved to occur at a faster speed (3 cells per time step) than the checkerboard uncertainty; it can overrun checkerboard regions, and in fact collisions of opposing types of this strong propagation give rise to the (hard uncertainty) blinker boundaries in the first place. The final configuration is therefore the result of this distributed computation.

Quantification of the local information dynamics via these *three axes of complexity* (information storage, transfer and modification) will provide quite detailed insights into the distributed computation carried out in a system. In all of these CAs we expect local measures of information storage to highlight blinkers and domain regions, local measures of information transfer to highlight particles (including gliders and domain walls), and local measures of information modification to highlight particle collisions.

This will provide a deeper understanding of computation than single or generic measures of bulk statistical behavior, from which conflict often arises in attempts to provide classification of complex behavior. In particular, we seek clarification on the long-standing debate regarding the nature of computation in ECA rule 22. Suggestions that rule 22 is complex include the difficulty in estimating the metric entropy (i.e. temporal entropy rate) for rule 22 by Grassberger (1986b), due to "complex long-range effects, similar to a critical phenomenon" (Grassberger 1986a). This effectively corresponds to an implication that rule 22 contains an infinite amount of memory (see Section 1.4.1). Also, from an initial condition of only a single "on" cell, rule 22 forms a pattern known as the "Sierpinski Gasket" (Wolfram 2002) which exhibits clear fractal structure. Furthermore, rule 22 is a 1D mapping of the 2D Game of Life CA (known to have the capability for universal computation (Conway 1982)) and in this sense is referred to as "life in one dimension" (McIntosh 1990), and complex structure in the language generated by iterations of rule 22 has been identified by Badii and Politi (1997). Also, we reported in (Lizier et al. 2012b) that we have investigated the C_1 complexity measure (Lafusa and Bossomaier 2005) (an enhanced version of the variance of the input entropy (Wuensche 1999)) for all ECAs, and found rule 22 to clearly exhibit the largest value of this measure (0.78 bits to rule 110's 0.085 bits). On the other hand, suggestions that rule 22 is not complex include its high sensitivity to initial conditions leading to Wolfram (2002) classifying it as class III chaotic. Gutowitz and Domain (1997) claim this renders it as chaotic despite the subtle long-range effects it displays, further identifying its fast statistical convergence, and exponentially long and thin transients in state space (see Wuensche (1999)). Importantly, no coherent structure (particles, collisions, etc.) is found for typical profiles of rule 22 using a number of known filters for such structure (e.g. local statistical complexity (Shalizi et al. 2006)): this reflects the paradigm shift to an examination of local dynamics rather than generic, overall or averaged analysis. In our approach, we seek to combine this local viewpoint of the dynamics with a quantitative breakdown of the individual elements of computation, and we will review the application to rule 22 in this light.

1.4 Information Storage

In this section we review the methods to quantify information storage on a local scale in space and time, as presented in (Lizier et al. 2012c). We describe how total information storage used in the future is captured by excess entropy, and introduce active information storage to capture the amount of information storage that is currently in use. We review the application of local profiles of both measures to cellular automata.

1.4.1 Excess entropy as total information storage

Although discussion of information storage or memory in CAs has often focused on periodic structures (particularly in construction of universal Turing machines), information storage does not necessarily entail periodicity. The excess entropy Eq. (1.12, 1.13) more broadly encompasses all types of structure and memory by capturing correlations across all lengths of time, including non-linear correlations. It is quite clear from the predictive information formulation of the excess entropy Eq. (1.13) – as the information from a process' past that is contained its future – that it is a measure of the total information storage used in the future of a system.[3]

We use the term *univariate excess entropy*[4] to refer to measuring the excess entropy for individual variables X using their one-dimensional time-series process, i.e. $E_X = \lim_{k\to\infty} I_{\mathbf{X}_n^{(k)};\mathbf{X}_{n+1}^{(k^+)}}$ from Eq. (1.13). This is a measure of the *average* memory for *each variable* X. Furthermore, we use the term *collective excess entropy* to refer to measuring the temporal excess entropy for a collective of variables $\mathbf{X} = \{X_1, X_2, \ldots, X_m\}$ (e.g. a set of neighboring cells in a CA) using their two-dimensional time-series process. Considered as the mutual information between their joint past and future, i.e. a joint temporal predictive information:

$$E_{\mathbf{X}} = \lim_{k\to\infty} I_{\{\mathbf{X}_{1,n}^{(k)},\mathbf{X}_{2,n}^{(k)},\ldots,\mathbf{X}_{m,n}^{(k)}\};\{\mathbf{X}_{1,n+1}^{(k^+)},\mathbf{X}_{2,n+1}^{(k^+)},\ldots,\mathbf{X}_{m,n+1}^{(k^+)}\}}, \tag{1.19}$$

this is a measure of the *average* total memory stored in the collective (i.e. stored *collectively* by a set of cells in a CA). Collective excess entropy could be used for example to quantify the "undiscovered *collective memory* that may present in certain fish schools" (Couzin et al. 2006).

Grassberger (1986b,a) studied temporal entropy rate estimates for several ECAs in order to gain insights into their excess entropies. He revealed divergent collective excess entropy for a number of rules, including rule 22, implying a highly complex process. This case has been described by Lindgren and Nordahl (1988) as "a phenomenon which can occur in more complex environments", as with strong long-range

[3] In (Lizier et al. 2012c) we provide further comment on the relation to the statistical complexity (Crutchfield and Young 1989), which measures *all* information stored by the system which *may be used* in the future, while the *excess entropy* measures that information which *is used* by the system *at some point* in the future. The relation between the two concepts is covered in a more general mathematical context by Shalizi and Crutchfield (2001).

[4] Called "single-agent excess entropy" in (Lizier et al. 2012c).

correlations a semi-infinite sequence "could store an infinite amount of information about its continuation" (as per the predictive information form of the excess entropy Eq. (1.13)). On the other hand, infinite collective excess entropy can also be achieved for systems that only trivially utilise all of their available memory (e.g. simply copying cell values to the right when started from random initial states). Rule 22 was inferred to have $H_{\mu,N} = 0$ and infinite collective excess entropy, which was interpreted as a process requiring an infinite amount of memory to maintain an aperiodicity (Crutchfield and Feldman 2003).

In attempting to quantify *local* information dynamics of distributed computation here, our focus is on information storage for *single variables or cells* rather than the joint information storage across the collective. Were the univariate excess entropy found to be divergent (this has not been demonstrated), this may be more significant than for the collective case: divergent collective excess entropy implies that the collective is at least trivially utilizing all of its available memory (and even the chaotic rule 30 exhibits this), whereas divergent univariate excess entropy implies that all cells are individually highly utilizing the resources of the collective in a highly complex process. Again though, we emphasize that our focus is on local measures in *time* as well as space, which we present in the next section.

First we note that with respect to CAs, where each cell has only a finite number of values b and takes direct influence from only its single past value and the values of a finite number of neighbors, the meaning of (either average or local) information storage being greater $\log_2 b$ bits (let alone infinite) in the time series process of a single cell is not immediately obvious. Clearly, a cell in an ECA cannot store more than 1 bit of information in isolation. However, the *bidirectional* communication in CAs effectively allows a cell to store extra information in neighbors (even beyond the immediate neighbors), and to subsequently retrieve that information from those neighbors at a later point in time. While measurement of the excess entropy does not explicitly look for such *self-influence* communicated through neighbors, it is indeed the method by which a significant portion of information is channeled. Considering the predictive information interpretation in Eq. (1.13), it is easy to picture self-influence between semi-infinite past and future blocks being conveyed via neighbors (see Fig. 1.1a). This is akin to the use of stigmergy (indirect communication through the environment, e.g. see Klyubin et al. (2004)) to communicate with oneself.

A measurement of more than $\log_2 b$ bits stored by a cell on average, or indeed an infinite information storage, is then a perfectly valid result: in an infinite CA, each cell has access to an infinite amount of neighbors in which to store information which can later be used to influence its own future. Note however, that since the storage medium is shared by all cells, one should not think about the total memory as the total number of cells multiplied by this average. The total memory would be properly measured by the collective excess entropy, which takes into account the inherent redundancy here.

Following similar reasoning (i.e. that information may be stored and retrieved from one's neighbors), we note that a variable can store information regardless of whether it is causally connected with itself. Also, note that a variable can be perceived to store information simply as a result of how that variable is driven (Obst et al. 2013), i.e. where information is physically stored elsewhere in the system but recurs in the variable at

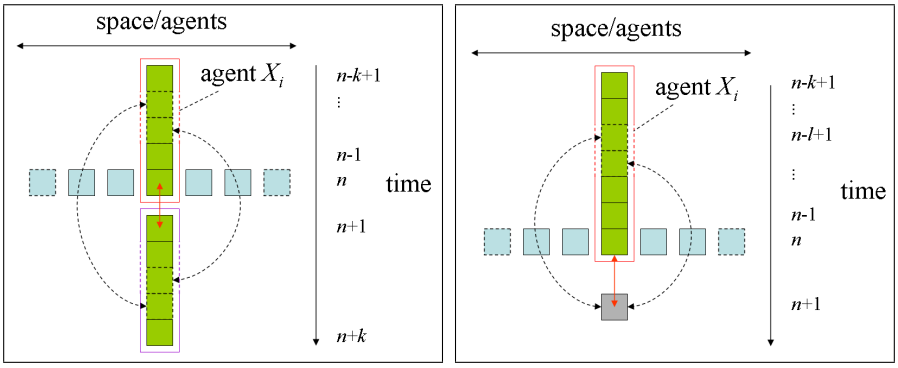

(a) Excess Entropy (b) Active Information Storage

Fig. 1.1: Measures of information storage in the time-series processes of single variables in distributed systems. (a) Excess Entropy: *total* information from the variable's past that is predictive of its future. (b) Active Information storage: the information storage that is *currently in use* in determining the next value of the variable. The stored information can be conveyed directly through the variable itself or via feedback from neighbors. (NB: This figure is reprinted from (Lizier et al. 2012c), Lizier, J. T., Prokopenko, M., and Zomaya, A. Y., Local measures of information storage in complex distributed computation, *Information Sciences*, 208:39–54, Copyright (2012), with permission from Elsevier.)

different time steps (e.g. see the description of information storage in feed-forward loop motifs in (Lizier et al. 2012a)).

1.4.2 Local excess entropy

We now shift focus to local measures of information storage, which have the potential to provide more detailed insights into information storage structures and their involvement in computation than single ensemble measures.

The *local excess entropy* is a measure of how much information a given variable is storing *at a particular point in time* (Shalizi 2001).[5] The local excess entropy $e_X(n+1)$ of a process is simply the local mutual information Eq. (1.16) of the semi-infinite past and future of the process X at the given time step $n+1$:

$$e_X(n+1) = \lim_{k \to \infty} \log_2 \frac{p(\mathbf{x}_n^{(k)}, \mathbf{x}_{n+1}^{(k^+)})}{p(\mathbf{x}_n^{(k)}) p(\mathbf{x}_{n+1}^{(k^+)})}. \tag{1.20}$$

Note that the excess entropy is the average of the local values, $E_X = \langle e_X(n) \rangle$. The limit $k \to \infty$ is an important part of this definition, since correlations at all time

[5] This is as per the original formulation of the local excess entropy by Shalizi (2001), however this presentation is for a single time-series rather than the light-cone formulation used there.

scales should be included in the computation of information storage. Since this is not computationally feasible in general, we retain the notation $e_X(n+1,k)$ to denote finite-k estimates of $e_X(n+1)$.

The notation is generalized for lattice systems (such as CAs) with *spatially-ordered* variables to represent the local excess entropy for cell X_i at time $n+1$ as:

$$e(i,n+1) = \lim_{k\to\infty} \log_2 \frac{p(\mathbf{x}_{i,n}^{(k)}, \mathbf{x}_{i,n+1}^{(k^+)})}{p(\mathbf{x}_{i,n}^{(k)})p(\mathbf{x}_{i,n+1}^{(k^+)})}. \tag{1.21}$$

Again, $e(i,n+1,k)$ is used to denote finite-k estimates of $e(i,n+1)$. Local excess entropy is defined for every spatiotemporal point (i,n) in the system. (Alternatively, the collective excess entropy can only be localized in time).

As a local mutual information, the local excess entropy may be positive or negative, meaning the past history of the cell can either positively inform us or actually misinform us about its future. An observer is misinformed where a given semi-infinite past and future are relatively unlikely to be observed together as compared to the product of their marginal probabilities. Another view is that we have misinformative values when $p(\mathbf{x}_{i,n+1}^{(k^+)} \mid \mathbf{x}_{i,n}^{(k)}) < p(\mathbf{x}_{i,n+1}^{(k^+)})$, meaning that taking the past $\mathbf{x}_{i,n}^{(k)}$ into account reduced the probability of the future which was observed $\mathbf{x}_{i,n+1}^{(k^+)}$.

1.4.3 Active information storage

The excess entropy measures the total stored information which will be used *at some point* in the future of the time-series process of a variable, possibly but not necessarily at the next time step $n+1$. In examining the local information dynamics of computation, we are interested in how much of the stored information is actually *in use* at the next time step. As we will see in Section 1.6, this is particularly important in understanding how stored information interacts with information transfer in information processing. As such, the **active information storage** A_X was introduced in (Lizier et al. 2012c) as the average mutual information between the (semi-infinite) past state of the process and its *next value*, as opposed to its whole (semi-infinite) future:

$$A_X = \lim_{k\to\infty} I(\mathbf{X}_n^{(k)}; X_{n+1}). \tag{1.22}$$

The **local active information storage** is then a measure of the amount of information storage in use by the process at a particular time-step $n+1$:

$$a_X(n+1) = \lim_{k\to\infty} \log_2 \frac{p(\mathbf{x}_n^{(k)}, x_{n+1})}{p(\mathbf{x}_n^{(k)})p(x_{n+1})}, \tag{1.23}$$

$$= \lim_{k\to\infty} \log_2 \frac{p(x_{n+1} \mid \mathbf{x}_n^{(k)})}{p(x_{n+1})}, \tag{1.24}$$

and we have $A_X = \langle a_X(n) \rangle$. We retain the notation $a_X(n+1,k)$ and $A_X(k)$ for finite-k estimates. Again, we generalize the measure for variable X_i in a lattice system as:

$$a(i, n+1) = \lim_{k \to \infty} \log_2 \frac{p(\mathbf{x}_{i,n}^{(k)}, x_{i,n+1})}{p(\mathbf{x}_{i,n}^{(k)}) p(x_{i,n+1})}, \tag{1.25}$$

and use $a(i, n+1, k)$ to denote finite-k estimates there, noting that the local active information storage is defined for every spatiotemporal point (i, n) in the lattice system.

The average active information storage will always be positive (as for the excess entropy), but is bounded above by $\log_2 b$ bits if the variable takes one of b discrete values. The local active information storage is not bound in this manner however, with values larger than $\log_2 b$ indicating that the particular past of an variable provides strong positive information about its next value. Furthermore, the local active information storage can be negative, where the past history of the variable is actually misinformative about its next value. An observer is misinformed where the past history and observed next value are relatively unlikely to occur together as compared to their separate occurrence.

1.4.4 Local information storage results

In this and subsequent results sections, we review the application of these local measures in (Lizier et al. 2007, 2008c, 2012c, 2010, 2012b; Lizier and Prokopenko 2010; Lizier 2013) to sample CA runs. As described earlier, we are interested in studying the non-trivial computation during the transient dynamics before an attractor is reached. Certainly it would be easier to study these information dynamics on attractors – since the dynamics there are cyclo-stationary (because the attractors in finite-length CAs involve only fixed or periodic dynamics) – however as described in Section 1.3.3 the computation there is trivial. To investigate the dynamics of the transient, we estimate the required probability distribution functions (PDFs) from CA runs of 10 000 cells, initialized from random values, in order to generate a large ensemble of transient automata dynamics. We retain only a relatively short 600 time steps for each cell, in order to avoid attractor dynamics and focus on quasi-stationary transient dynamics during that short time period. Alternatively, for ϕ_{par} we used 30 000 cells with 200 time steps retained. Periodic boundary conditions were used. Observations taken at every spatiotemporal point in the CA were used in estimating the required PDFs, since the cells in the CA are homogeneous variables and quais-stationarity is assumed over the relatively short time interval.

The results and the figures displayed here were produced using the open source *Java Information Dynamics Toolkit* (Lizier 2012), which can be used in Matlab/Octave and Python as well as Java. All results can be reproduced using the Matlab/Octave script `GsoChapterDemo2013.m` in the `demos/octave/CellularAutomata` example distributed with this toolkit. We make estimates of the measures with finite values of k, noting that the insights described here could not be attained unless a reasonably large value of k was used in order to capture a large proportion of the correlations. Determination of an appropriate value of k was discussed in (Lizier et al. 2012c), and in (Lizier et al. 2008c) for the related transfer entropy measure presented in Section 1.5. As a rule of thumb, k should at least be larger than the period of any regular background domain in order to capture the information storage underpinning its continuation.

We begin by examining the results for rules 54 and 110, which contain regular gliders against periodic background domains. For the CA runs described above, sample areas of the large CAs are shown in Fig. 1.2a and Fig. 1.4a, while the corresponding local profiles of $e(i, n, k = 8)$ generated are displayed in Fig. 1.2b and Fig. 1.4b, and the local profiles of $a(i, n, k = 16)$ in Fig. 1.2c and Fig. 1.4c. It is quite clear that positive information storage is concentrated in the vertical gliders or blinkers, and the domain regions. As expected, these results provide quantitative evidence that the **blinkers are the dominant information storage entities**. That the **domain regions contain significant information storage** should not be surprising, since as a periodic sequence its past does indeed store information about its future.

In fact, the local values for each measure form spatially and temporally periodic patterns in the domains, corresponding to the spatial and temporal periodicities exhibited in the underlying raw values. Certainly if the dynamics are only composed of a consistent domain pattern (which is deterministic when viewing single cells' time series), then for $a(i, n, k)$ for example we will always have $p(x_{n+1} \mid \mathbf{x}_n^{(k)}) = 1$ and if $p(x_{n+1})$ is balanced then $a(i, n, k)$ would be constant across the CA. However, the existence of discontinuities in the domain, e.g. gliders, reduces $p(x_{n+1} \mid \mathbf{x}_n^{(k)})$ here, and does so differently for each $\mathbf{x}_n^{(k)}$ configuration in the domain. Imbalances in $p(x_{n+1})$ can also contribute to differences in storage across the domain. These factors leads to the spatiotemporal periodicities of information storage that are observed in the domains.

While the local active information storage indicates a similar amount of stored information in use to compute each space-time point in both the domain and blinker areas, the local excess entropy reveals a larger *total* amount of information is stored in the blinkers. For the blinkers known as α and β in rule 54 (Hordijk et al. 2001) this is because the temporal sequences of the center columns of the blinkers (0-0-0-1, with $e(i, n, k = 8)$ in the range 5.01 to 5.32 bits) are more complex than those in the domain (0-0-1-1 and 0-1, with $e(i, n, k = 8)$ in the range 1.94 to 3.22 bits), even where they are of the same period. We have $e(i, n, k = 8) > 1$ bit here due to the distributed information storage supported by bidirectional communication (as discussed earlier). Such bidirectional communication is also critical to these periodic domain sequences being longer than two time steps – the maximum period that a binary cell could sustain in isolation (e.g. the period-7 domain in rule 110).

Another area of strong information storage appears to be the "wake" of the more complex gliders in rule 110 (see the glider at top right of Fig. 1.4b and Fig. 1.4c). This result aligns well with our observation (Lizier et al. 2008c) that the dynamics following the leading edge of regular gliders consists largely of "non-traveling" information. The presence of the information storage is shown by both measures, although the relative strength of the total information storage is again revealed only by the local excess entropy.

Negative values of $a(i, n, k = 16)$ for rules 54 and 110 are also visible in Fig. 1.2c and Fig. 1.4c. Interestingly, negative local components of local active information storage measure are concentrated in the traveling glider areas (e.g. γ^+ and γ^- for rule 54 (Hordijk et al. 2001)), providing a good spatiotemporal filter of these struc-

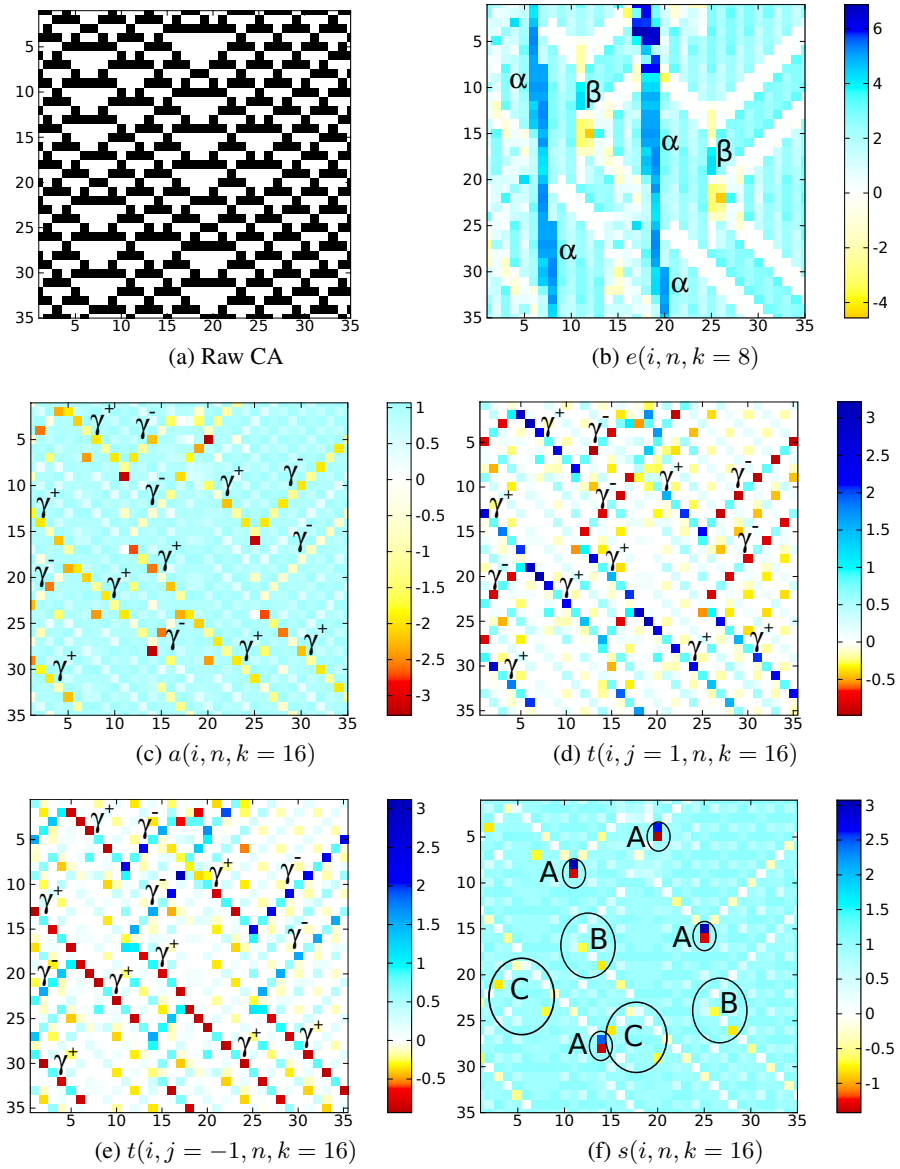

(a) Raw CA
(b) $e(i, n, k = 8)$
(c) $a(i, n, k = 16)$
(d) $t(i, j = 1, n, k = 16)$
(e) $t(i, j = -1, n, k = 16)$
(f) $s(i, n, k = 16)$

Fig. 1.2: Local information dynamics in **rule 54** for the raw values in (a) (black for "1", white for "0"). 35 time steps are displayed for 35 cells, and time increases down the page for all CA plots. All units are in bits. (b) Local excess entropy $e(i, n, k = 8)$; (c) Local active information storage $a(i, n, k = 16)$; Local apparent transfer entropy: (d) one cell to the right $t(i, j = 1, n, k = 16)$, (e) one cell to the left $t(i, j = -1, n, k = 16)$; (f) Local separable information $s(i, n, k = 16)$.

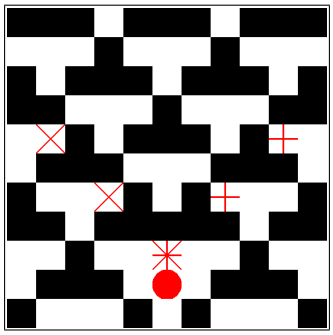

Fig. 1.3: Close up of raw values of rule 54. "x" and "+" mark some positions in the γ^+ and γ^- gliders respectively. Note their point of coincidence in collision type "A", with "•" marking the subsequent non-trivial information modification as detected using $s(i, n, k = 16) < 0$. (Reprinted with permission from (Lizier et al. 2010) J. T. Lizier, M. Prokopenko, and A. Y. Zomaya, "Information modification and particle collisions in distributed computation," *Chaos*, vol. 20, no. 3, p. 037109, 2010. Copyright 2010, AIP Publishing LLC.)

tures. This is because when a traveling glider is encountered at a given cell, the past history of that cell (being part of the background domain) is misinformative about the next value, since the domain sequence was more likely to continue than be interrupted. For example, see the marked positions of the γ gliders in Fig. 1.3. There we have $p(x_{n+1} \mid \mathbf{x}_n^{(k=16)}) = 0.25$ and $p(x_{n+1}) = 0.52$: since the next value occurs relatively infrequently after the given history, we have a misinformative $a(n, k = 16) = -1.09$ bits. This is juxtaposed with the points four time steps before those marked "x", which have the same history $\mathbf{x}_n^{(k=16)}$ but are part of the domain, with $p(x_{n+1} \mid \mathbf{x}_n^{(k=16)}) = 0.75$ and $p(x_{n+1}) = 0.48$ giving $a(n, k = 16) = 0.66$ bits, quantifying the positive information storage there. Note that the points with misinformative information storage are not necessarily those selected by other filtering techniques as part of the gliders: e.g. the finite state transducers technique (using left to right scanning by convention) by Hanson and Crutchfield (1997) would identify points 3 cells to the right of those marked "x" as part of the γ^+ glider.

The local excess entropy produced some negative values around traveling gliders, though these were far less localized on the gliders themselves and less consistent in occurrence than for the local active information storage. This is because the local excess entropy, as measure of total information storage into the future, is more loosely tied to the dynamics at the given spatiotemporal point. The effect of a glider encounter on $e(i, n, k)$ is smeared out in time, and in fact the dynamics may store more positive information in total than the misinformation encountered at the specific location of the glider. For example, glider pairs were observed in (Lizier et al. 2012c) to have positive total information storage, since a glider encounter becomes much more likely in the wake of a previous glider.

As another rule containing regular gliders against a periodic background domain, analysis of the raw values of ϕ_{par} in Fig. 1.5a provides similar results for $e(i, n, k = 5)$ (not shown, see (Lizier 2013)) and $a(i, n, k = 10)$ in Fig. 1.5b here. One distinction is that the blinker here contains no more stored information than the domain, since it is no more complicated. Importantly, we confirm the information storage capability of the blinkers and domains in this human understandable computation.

Another interesting example is provided by ECA rule 18, which contains domain walls against a seemingly irregular background domain. We measured the local information profiles for $e(i, n, k = 8)$ and $a(i, n, k = 16)$ in (Lizier et al. 2012c) (shown in that paper, but not here). Importantly, the most significant negative components of the local active information storage are concentrated on the domain walls: analogous to the regular gliders of rule 54, when a domain wall is encountered the past history of the cell becomes misinformative about its next value. There is also interesting information storage dynamics in the background domain for rule 18, discussed in detail in (Lizier et al. 2012c).

Finally, we examine ECA rule 22, suggested to have infinite collective excess entropy (Grassberger 1986b,a) but without any known coherent structural elements (Shalizi et al. 2006). For the raw values of rule 22 displayed in Fig. 1.6a, the calculated local excess entropy profile is shown in Fig. 1.6b, and the local active information storage profile in Fig. 1.6c. While information storage certainly occurs for rule 22, these plots provide evidence that there is no coherent structure to this storage. This is another clear example of the utility of examining local information dynamics over ensemble estimates, given the earlier discussion on collective excess entropy for rule 22.

In summary, we have demonstrated that the local active information storage and local excess entropy provide insights into information storage dynamics that, while often similar in general, are sometimes subtly different. While both measures provide useful insights, the local active information storage is the most useful in a real-time sense, since calculation of the local excess entropy requires knowledge of the dynamics an arbitrary distance into the future.[6] Furthermore, it also provides the most specifically localized insights, including filtering moving elements of coherent spatiotemporal structure. This being said, it is not capable of identifying the information source of these structures; for this, we turn our attention to a specific measure of information transfer.

1.5 Information Transfer

Information transfer refers to a directional signal or communication of dynamic information from a *source* to a *destination*. In this section, we review descriptions of how

[6] Calculation of $e(i, n, k)$ using local block entropies analogous to Eq. (1.12) would also require block entropies to be taken into the future to compute the same local information storage values. Without taking account of the dynamics into the future, we will not measure the information storage that *will* be used in the future of the process, but the information storage that is *likely* to be used in the future.

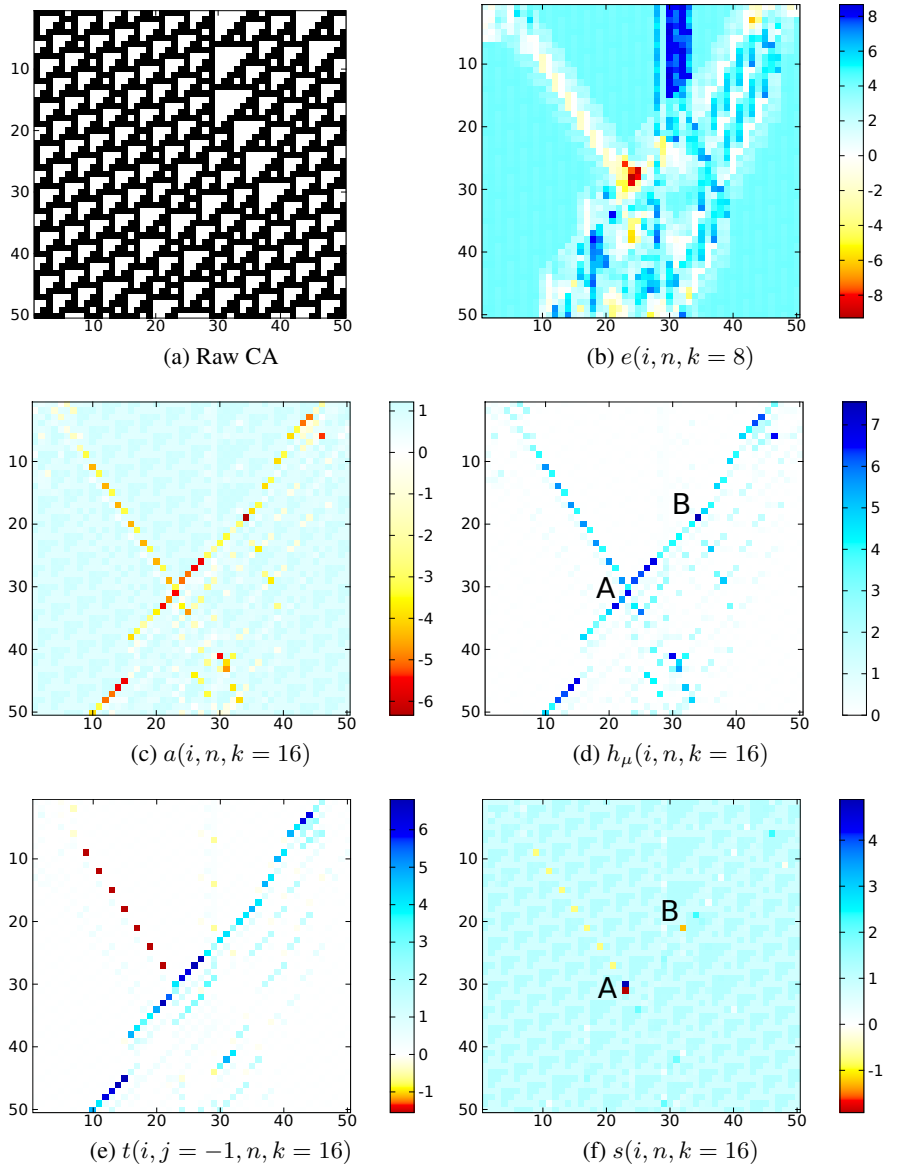

Fig. 1.4: Local information dynamics in **rule 110** for the raw values displayed in (a) (black for "1", white for "0"). 50 time steps are displayed for 50 cells, and all units are in bits. (b) Local excess entropy $e(i, n, k = 8)$; (c) Local active information storage $a(i, n, k = 16)$; (d) Local temporal entropy rate $h_\mu(i, n, k = 16)$; (e) Local apparent transfer entropy one cell to the left $t(i, j = -1, n, k = 16)$; (f) Local separable information $s(i, n, k = 16)$.

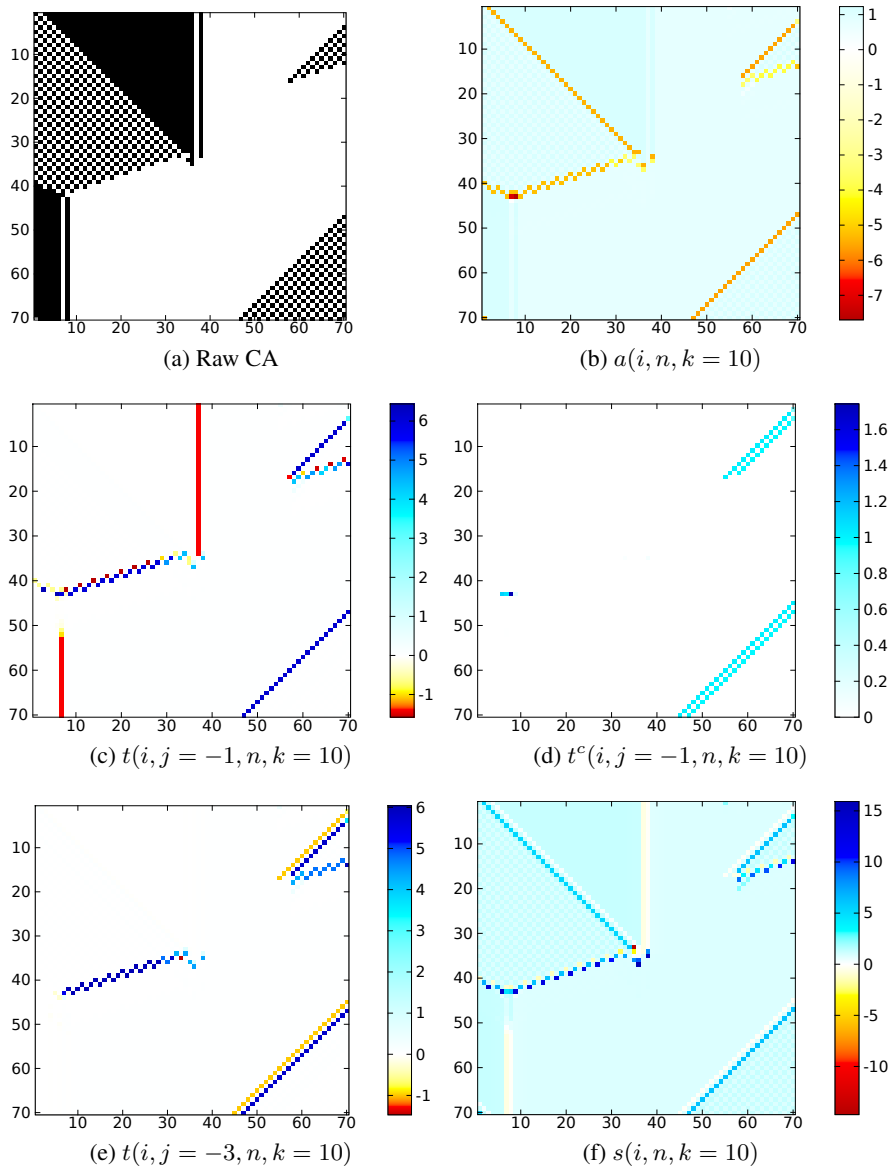

(a) Raw CA

(b) $a(i, n, k = 10)$

(c) $t(i, j = -1, n, k = 10)$

(d) $t^c(i, j = -1, n, k = 10)$

(e) $t(i, j = -3, n, k = 10)$

(f) $s(i, n, k = 10)$

Fig. 1.5: Local information dynamics in $r = 3$ **rule** $\phi_{\mathbf{par}}$ for the raw values displayed in (a) (black for "1", white for "0"). 70 time steps are displayed for 70 cells, and all units are in bits. (b) Local active information storage $a(i, n, k = 10)$; Local apparent transfer entropy: (c) one cell to the left $t(i, j = -1, n, k = 10)$, and (e) three cells to the left $t(i, j = -3, n, k = 10)$; (d) Local complete transfer entropy one cell to the left $t^c(i, j = -1, n, k = 10)$; (f) Local separable information $s(i, n, k = 10)$.

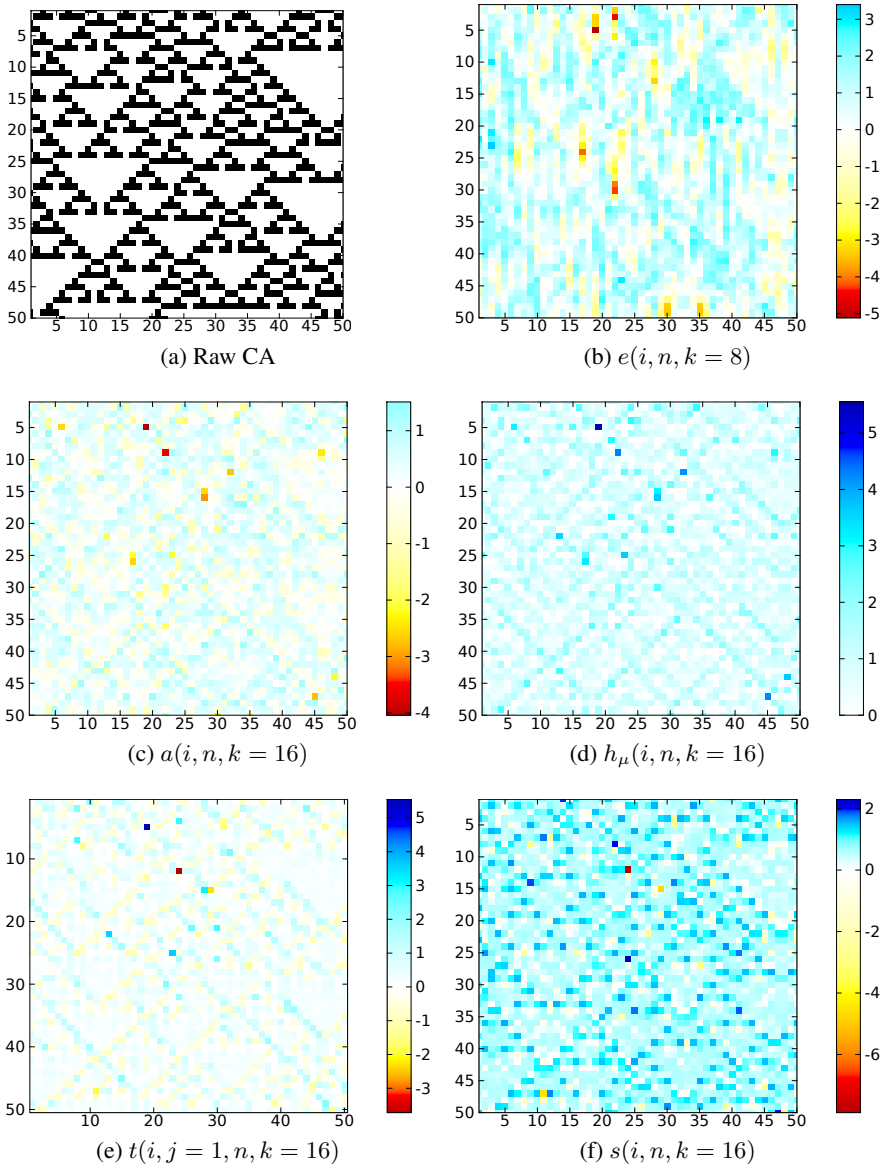

(a) Raw CA

(b) $e(i, n, k = 8)$

(c) $a(i, n, k = 16)$

(d) $h_\mu(i, n, k = 16)$

(e) $t(i, j = 1, n, k = 16)$

(f) $s(i, n, k = 16)$

Fig. 1.6: Local information dynamics in **rule 22** for the raw values in (a) (black for "1", white for "0"). 50 time steps displayed for 50 cells, and all units are in bits. (b) Local excess entropy $e(i, n, k = 8)$; (c) Local active information storage $a(i, n, k = 16)$; (d) Local temporal entropy rate $h_\mu(i, n, k = 16)$; (e) Local apparent transfer entropy one cell to the right $t(i, j = 1, n, k = 16)$; (f) Local separable information $s(i, n, k = 16)$.

to measure information transfer in complex systems from (Lizier et al. 2008c, 2010; Lizier 2013), and the associated application to several ECA rules.

1.5.1 Local transfer entropy

Schreiber (2000) presented **transfer entropy** as a measure for information transfer in order to address deficiencies in the previous de facto measure, mutual information (Eq. (1.4)), the use of which he criticized in this context as a symmetric measure of statically shared information. Transfer entropy is defined as the deviation from independence (in bits) of the state transition of an information destination *X* from the previous state of an information source *Y*:

$$T_{Y \to X}(k, l) = \sum_{w_n} p(w_n) \log_2 \frac{p(x_{n+1} \mid \mathbf{x}_n^{(k)}, \mathbf{y}_n^{(l)})}{p(x_{n+1} \mid \mathbf{x}_n^{(k)})}, \tag{1.26}$$

where w_n is the state transition tuple $(x_{n+1}, \mathbf{x}_n^{(k)}, \mathbf{y}_n^{(l)})$. This is shown diagrammatically in Fig. 1.7a. The transfer entropy will be zero if the next value of the destination is completely dependent on its past (leaving no information for the source to add), or if the state transition of the destination is independent of the destination. At the other extreme, it will be maximal if the state transition is completely specified by the source (in the context of the destination's past). As such, the transfer entropy is a *directional*, *dynamic* measure of information transfer. It is a *conditional* mutual information, casting it as the average information in the source about the next state of the destination conditioned on the destination's past. We have provided a thermodynamic interpretation of transfer entropy in (Prokopenko et al. 2013).

The role of the past state of the destination $\mathbf{x}_n^{(k)}$ is particularly important here. This past state can indirectly influence the next value via the source or other neighbors: this may be mistaken as an independent flow from the source here (Lizier et al. 2008c). In the context of distributed computation, this is recognizable as the *active information storage*. That is, conditioning on the destination's history $\mathbf{x}_n^{(k)}$ serves to eliminate the active information storage from the transfer entropy measurement. Yet any self-influence transmitted prior to these k values will not be eliminated: in (Lizier et al. 2008c) we suggested that the asymptote $k \to \infty$ is most correct for variables displaying non-Markovian dynamics. Just as the excess entropy and active information storage require $k \to \infty$ to capture all information storage, accurate measurement of the transfer entropy requires $k \to \infty$ to eliminate all information storage from being mistaken as information transfer. Further to these, even if the destination variable does display Markovian dynamics of order k, synergistic interactions between the source and the past of the destination beyond k time steps necessitate the use of a longer destination history to capture the information transfer, again leading us to $k \to \infty$ to capture all transfer. We describe other interpretations of the role of $\mathbf{x}_n^{(k)}$ in (Lizier and Mahoney 2013), including properly capturing the state transition of the destination and capturing the contribution of the source in the context of that state transition; which align with the above. The most generally correct form of the transfer entropy is therefore computed as:

$$T_{Y\to X}(l) = \lim_{k\to\infty} \sum_{w_n} p(w_n) \log_2 \frac{p(x_{n+1} \mid \mathbf{x}_n^{(k)}, \mathbf{y}_n^{(l)})}{p(x_{n+1} \mid \mathbf{x}_n^{(k)})}, \tag{1.27}$$

with $T_{Y\to X}(k,l)$ retained for finite-k estimates.

Also, we note that considering a source *state* $\mathbf{y}_n^{(l)}$ rather than a scalar y_n is most appropriate where the observations y mask a hidden causal process in Y, or where multiple past values of Y in addition to y_n are causal to x_{n+1}. Otherwise, where y_n is directly causal to x_{n+1}, and where it is the only direct causal source in Y (e.g. in CAs), we use only $l = 1$ (Lizier et al. 2008c; Lizier and Prokopenko 2010) and drop it from our notation here. Furthermore, note that one may use source-destination delays other than one time step, and indeed it is most appropriate to match any causal delay from Y to X (Wibral et al. 2013).

Next, we introduced the corresponding **local transfer entropy** at each observation n in (Lizier et al. 2008c):

$$t_{Y\to X}(n+1,l) = \lim_{k\to\infty} t_{Y\to X}(n+1,k,l), \tag{1.28}$$

$$t_{Y\to X}(n+1,k,l) = \log_2 \frac{p(x_{n+1} \mid \mathbf{x}_n^{(k)}, \mathbf{y}_n^{(l)})}{p(x_{n+1} \mid \mathbf{x}_n^{(k)})}. \tag{1.29}$$

The local transfer entropy describes the information added by a specific source state $\mathbf{y}_n^{(l)}$ about x_{n+1} in the context of the past of the destination $\mathbf{x}_n^{(k)}$. Of course, we have $T_{Y\to X}(k,l) = \langle t_{Y\to X}(n+1,k,l)\rangle$.

For lattice systems such as CAs with spatially-ordered variables, the local information transfer to agent X_i from X_{i-j} (across j cells to the right) at time $n+1$ is represented as:

$$t(i,j,n+1,l) = \lim_{k\to\infty} t(i,j,n+1,k,l), \tag{1.30}$$

$$t(i,j,n+1,k,l) = \log_2 \frac{p(x_{i,n+1} \mid \mathbf{x}_{i,n}^{(k)}, \mathbf{x}_{i-j,n}^{(l)})}{p(x_{i,n+1} \mid \mathbf{x}_{i,n}^{(k)})}. \tag{1.31}$$

This information transfer $t(i,j,n+1,k,l)$ to variable X_i from X_{i-j} at time $n+1$ is illustrated in Fig. 1.7a. Then $t(i,j,n,k,l)$ is defined for every spatiotemporal destination (i,n), for every information channel or direction j; sensible values for j correspond to causal information sources, i.e. for CAs, sources within the cell range $|j| \leq r$. Again, for homogeneous variables (with stationarity) it is appropriate to estimate the PDFs used in Eq. (1.31) from all spatiotemporal observations, and we write the average across homogeneous variables as $T(j,k) = \langle t(i,j,n,k)\rangle$.

Calculations conditioned on no other information contributors (as in Eq. (1.31))) are labeled as *apparent* transfer entropy (Lizier et al. 2008c). Local apparent transfer entropy $t(i,j,n,k)$ may be either positive or negative, with negative values occurring where (given the destination's history) the source element is actually misleading about the next value of the destination. In deterministic systems, this can only occur where another source is influencing the destination at that time. To counter that effect, the

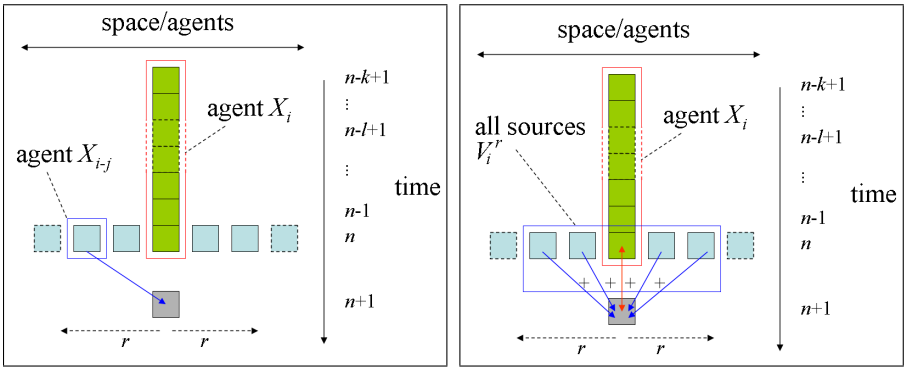

(a) Transfer Entropy (b) Separable Information

Fig. 1.7: (a) Transfer Entropy $t(i, j, n+1, k)$: information contained in the source cell X_{i-j} about the next value of the destination cell X_i at time $n+1$ in the context of the destination's past. (b) Separable information $s(i, n+1, k)$: information gained about the next value of the destination from separately examining each causal information source in the context of the destination's past. For CAs these causal sources are within the cell range r. (NB: Fig. 1.7a is reprinted with kind permission from Springer Science+Business Media: (Lizier 2013) Lizier, J. T. , *The Local Information Dynamics of Distributed Computation in Complex Systems*, Springer Theses, Springer, Berlin / Heidelberg, Copyright 2013. Fig. 1.7b is reprinted with permission from (Lizier et al. 2010) J. T. Lizier, M. Prokopenko, and A. Y. Zomaya, "Information modification and particle collisions in distributed computation," *Chaos*, vol. 20, no. 3, p. 037109, 2010. Copyright 2010, AIP Publishing LLC.)

transfer entropy may be conditioned on other possible causal information sources Z, to eliminate their influence from being attributed to the source in question Y (Schreiber 2000). We call this the **conditional transfer entropy** (Lizier et al. 2010), given (as a finite-k estimate) along with the **local conditional transfer entropy** as follows:

$$T_{Y \to X|Z}(k, l) = \left\langle t_{Y \to X|Z}(n+1, k, l) \right\rangle, \tag{1.32}$$

$$t_{Y \to X|Z}(n+1, k, l) = \log_2 \frac{p(x_{n+1} \mid \mathbf{x}_n^{(k)}, \mathbf{y}_n^{(l)}, z_n)}{p(x_{n+1} \mid \mathbf{x}_n^{(k)}, z_n)}. \tag{1.33}$$

Z may of course be multivariate, or be an embedded state vector $\mathbf{z}_n^{(m)}$ itself. Indeed, a special case involves conditioning on *all* sources jointly in the set of causal information contributors $\mathbf{V}_X$ to X, except for the source Y, i.e. $\mathbf{V}_X \setminus Y$. This gives the **complete transfer entropy** $T^c_{Y \to X}(k, l) = T_{Y \to X|\mathbf{V}_X \setminus Y}(k, l)$ (Lizier et al. 2008c). At time step n, this set $\mathbf{V}_X \setminus Y$ has joint state $\mathbf{v}_{x,y,n}$, giving the **local complete transfer entropy** (Lizier et al. 2008c):[7]

[7] Note that if past values of X are causal sources to the next value x_{n+1}, they can be included in $\mathbf{v}_{x,y,n}$, but this is irrelevant for complete TE since they are already conditioned on in $\mathbf{x}_n^{(k)}$.

$$t^c_{Y\to X}(n+1,k,l) = t_{Y\to X|\mathbf{V}_{X,Y}}(n+1,k,l), \tag{1.34}$$

$$\mathbf{v}_{x,y,n} = \{z_n \mid \forall Z \in \mathbf{V}_X \setminus Y\}. \tag{1.35}$$

For CAs the set of causal information contributors to X_i is the neighborhood $\mathbf{V}_i^r$ of X_i, and for the complete transfer entropy we condition on this set except for the source X_{i-j}: $\mathbf{V}_i^r \setminus X_{i-j}$. At time step n this set has joint value $\mathbf{v}^r_{i,j,n}$, giving the following expression for the local complete transfer entropy in CAs (Lizier et al. 2008c):

$$t^c(i,j,n+1,k) = \log_2 \frac{p\left(x_{i,n+1} \mid \mathbf{x}^{(k)}_{i,n}, x_{i-j,n}, \mathbf{v}^r_{i,j,n}\right)}{p\left(x_{i,n+1} \mid \mathbf{x}^{(k)}_{i,n}, \mathbf{v}^r_{i,j,n}\right)}, \tag{1.36}$$

$$\mathbf{v}^r_{i,j,n} = \{x_{i+q,n} \mid \forall q : -r \le q \le +r, q \ne -j, q \ne 0\}. \tag{1.37}$$

Again, the most correct form is $t^c(i,j,n+1)$ in the limit $k \to \infty$. In deterministic systems (e.g. CAs), complete conditioning renders $t^c(i,j,n) \ge 0$ because the source can only add information about the outcome of the destination.

1.5.2 Total information, entropy rate and collective information transfer

The total information required to predict the next value of any process X is the *local entropy* $h_X(n+1)$ Eq. (1.14). Similarly, the *local temporal entropy rate* $h_{\mu X}(n+1,k) = -\log_2 p(x_{n+1} \mid x^{(k)}_n)$ is the information to predict the next value of that process given that its past, and the entropy rate is the average of these local values: $H_{\mu X}(k) = \langle h_{\mu X}(n+1,k)\rangle$. For lattice systems we have $h_\mu(i,n+1,k)$. Now, the entropy can be considered as the sum of the active information storage and temporal entropy rate (Lizier et al. 2010, 2012c):

$$H_{X_{n+1}} = I_{X_{n+1};\mathbf{X}^{(k)}_n} + H_{X_{n+1}|\mathbf{X}^{(k)}_n}, \tag{1.38}$$

$$h_X(n+1) = a_X(n+1,k) + h_{\mu X}(n+1,k). \tag{1.39}$$

For deterministic systems (e.g. CAs) there is no intrinsic uncertainty, so the local temporal entropy rate is equal to the **local collective transfer entropy** (Lizier and Prokopenko 2010) and represents a collective information transfer: the information about the next value of the destination *jointly* added by the causal information sources in the context of the past of the destination. This suggested that the local collective transfer entropy (or simply the local temporal entropy rate $h_\mu(i,n,k)$ for deterministic systems) is likely to be a meaningful measure and filter for incoming information.

Also, we showed that the information in a destination variable can be expressed as a sum of incrementally conditioned mutual information terms, considering each of the sources iteratively (Lizier et al. 2010; Lizier and Prokopenko 2010). For ECAs, these expressions become:

$$h(i,n+1) = a(i,n+1,k) + t(i,j=-1,n+1,k) + t^c(i,j=1,n+1,k), \tag{1.40}$$

(and vice-versa in $j = 1, -1$). Clearly, this total information is not simply a simple sum of the active information storage and the apparent transfer entropy from each source, nor the sum of the active information storage and the complete transfer entropy from each source.

1.5.3 Local information transfer results

In this section, we review the application of the local apparent and complete transfer entropies, as well as the local entropy rate, to several ECA rules (Lizier et al. 2007, 2008c, 2010; Lizier 2013). We focus in particular here on the local apparent transfer entropy, whose profiles $t(i, j = 1, n, k = 16)$ (measuring transfer across one unit to the right per time step) are plotted for rules 54 (Fig. 1.2d) and 22 (Fig. 1.6e), with $t(i, j = -1, n, k = 16)$ (transfer across one unit to the left per time step) plotted for rules 54 (Fig. 1.2e), 110 (Fig. 1.4e) and ϕ_{par} (Fig. 1.5c).

Both the local apparent and complete transfer entropy highlight **particles as strong positive information transfer against background domains**. This is true for both regular gliders as well as domain walls in rule 18 (not shown here, see (Lizier et al. 2008c)). Importantly, the particles are measured as information transfer in their direction of macroscopic motion, as expected. As such, local transfer entropy provided the first quantitative evidence for the long-held conjecture that particles are the *dominant* information transfer agents in CAs. For example, at the "x" marks in Fig. 1.3 which denote parts of the right-moving γ^+ gliders, we have $p(x_{i,n+1} \mid \mathbf{x}_{i,n}^{(k=16)}, x_{i-1,n}) = 1.00$ and $p(x_{i,n+1} \mid \mathbf{x}_{i,n}^{(k=16)}) = 0.25$: there is a strong information transfer of $t(i, j = 1, n, k = 16) = 2.02$ bits here because the source (in the glider) added a significant amount of information to the destination about the continuation of the glider.

For ϕ_{par} we confirm the role of the gliders as information transfer agents in the human understandable computation, and demonstrate information transfer across multiple units of space per unit time step for fast-moving gliders in Fig. 1.5e. Interestingly, we also see in Fig. 1.5c ($j = -1$) and Fig. 1.5e ($j = -3$) that the apparent transfer entropy can attribute information transfer to several information sources, whereas the complete transfer entropy (see Fig. 1.5d) is more likely to attribute the transfer to the single causal source. We emphasize though that information transfer and causality are distinct concepts, as discussed in detail in (Lizier and Prokopenko 2010). This result also underlines that **the apparent and complete transfer entropies have a similar nature but are complementary** in together determining the next state of the destination (as in Eq. (1.40)). Neither measure is more correct than the other though – both are required to understand the dynamics fully. A more detailed example contrasting the two is studied for rule 18 in (Lizier et al. 2008c), showing that the complete TE detects transfer to X due to synergies between the source Y and conditioned variable Z, whereas the apparent TE does not.

We also examine the profiles of the local temporal entropy rate $h_\mu(i, n + 1, k)$ (which is equal to the local collective transfer entropy in these deterministic systems) here in Fig. 1.4d for rule 110 and Fig. 1.6d for rule 22. As expected, the local temporal entropy rate profiles $h_\mu(i, n+1, k)$ highlight particles moving in *each* relevant channel and are a useful *single* spatiotemporal filter for moving emergent structure. In fact,

these profiles are quite similar to the profiles of the negative values of local active information storage. This is not surprising given they are counterparts in Eq. (1.39): where $h_\mu(i, n+1, k)$ is strongly positive (i.e. greater than 1 bit), it is likely that $a(i, n+1, k)$ is negative since the local single cell entropy will average close to 1 bit for these examples. Unlike $a(i, n+1, k)$ however, the local temporal entropy rate $h_\mu(i, n+1, k)$ is never negative.

Note that while achieving the limit $k \to \infty$ is not computationally feasible, a large enough k was required to achieve a reasonable estimates of the transfer entropy; without this, as discussed earlier the active information storage was not eliminated from the transfer entropy measurements in the domains, and the measure did not distinguish the particles from the domains (Lizier et al. 2008c).

We also demonstrated (Lizier et al. 2008c) that while there is zero information transfer in an infinite periodic domain (since the dynamics there only involve information storage), there is a small non-zero information transfer in domains acting as a background to gliders, effectively indicating the *absence* of gliders. These small non-zero information transfers are stronger in the wake of a glider, indicating the absence of (relatively common) following gliders. Similarly, we note here that the local temporal entropy rate profiles $h_\mu(i, n+1, k)$ contain small but non-zero values in these periodic domains. Furthermore, there is interesting structure to the information transfer in the domain of rule 18, described in detail in (Lizier et al. 2008c). As such, while particles are the dominant information transfer agents in CAs, they are not the only transfer entities.

The highlighting of structure by local transfer entropy is similar to results from other methods of filtering for structure in CAs (Shalizi et al. 2006; Wuensche 1999; Hanson and Crutchfield 1992; Helvik et al. 2004), but subtly different in revealing the leading edges of gliders as the major transfer elements in the glider structures, and providing multiple profiles (one for each direction or channel of information transfer).

Also, a particularly relevant result for our purposes is the finding of negative values of transfer entropy for some space-time points in particles moving orthogonal to the direction of measurement in space-time. This is displayed for $t(i, j = 1, n, k = 16)$ in rule 54 (Fig. 1.2d), and $t(i, j = -1, n, k = 16)$ for rule 110 (Fig. 1.4e), and also occurs for rule 18 (see (Lizier et al. 2008c; Lizier 2013)). In general this is because the source, as part of the domain, suggests that this same domain found in the past of the destination is likely to continue; however since the next value of the destination forms part of the particle, this suggestion proves to be misinformative. For example, consider the "x" marks in Fig. 1.3 which denote parts of the right-moving γ^+ gliders. If we now examine the source at the right (still in the domain), we have $p(x_{i,n+1} \mid \mathbf{x}_{i,n}^{(k=16)}, x_{i+1,n}) = 0.13$, with $p(x_{i,n+1} \mid \mathbf{x}_{i,n}^{(k=16)}) = 0.25$ as before, giving $t(i, j = -1, n, k = 16) = -0.90$ bits: this is negative because the source (still in the domain) was misinformative about the destination.

Regarding the local information transfer structure of rule 22, we note similar results as for local information storage. There is much information transfer here (in fact the average value $T(j = 1, k = 16) = 0.19$ bits is greater than for rule 110 at 0.07 bits), although there is no coherent structure to this transfer. Again, this demonstrates the

utility of local information measures in providing more detailed insights into system dynamics than their global averages.

In this section, we have described how the local transfer entropy quantifies the information transfer at space-time points within a system, and provides evidence that particles are the dominant information transfer agents in CAs. We also described the collective transfer entropy, which quantifies the joint information contribution from all causal information contributors, and in deterministic systems is equal to the temporal entropy rate. However, we have not yet separately identified collision events in CAs: to complete our exploration of the information dynamics of computation, we now consider the nature of information modification.

1.6 Information Modification

Langton (1990) interpreted information modification as interactions between transmitted and/or stored information which resulted in a modification of one or the other. CAs provide an illustrative example, where the term *interactions* is generally interpreted to mean collisions of particles (including blinkers as information storage), with the resulting dynamics involving something other than the incoming particles continuing unperturbed. The resulting dynamics could involve zero or more particles (with an annihilation leaving only a background domain), and perhaps even some of the incoming particles. Given the focus on perturbations in the definition here, it is logical to associate a collision event with the modification of transmitted and/or stored information, and to see it as an information processing or decision event. Indeed, as an information processing event the important role of collisions in determining the dynamics of the system is widely acknowledged (Hordijk et al. 2001), e.g. in the ϕ_{par} density classification.

Attempts have previously been made to quantify information modification or processing in a system (Sánchez-Montañés and Corbacho 2002; Yamada and Aihara 1994; Kinouchi and Copelli 2006). However, these have either been too specific to allow portability across system types (e.g. by focusing on the capability of a system to solve a known problem, or measuring properties related to the particular type of system being examined), focus on general processing as movement or interpretation of information rather than specifically the modification of information, or are not amenable to measuring information modification at *local* space-time points *within* a distributed system.

In this section, we review the separable information (Lizier et al. 2010) as a tool to detect non-trivial information modification events, and demonstrate it as the first measure which filtered collisions in CAs as such. At the end of the section however, we describe criticisms of the separable information, and describe current efforts to develop new measures of information modification.

1.6.1 Local separable information

We begin by considering what it means for a particle to be *modified*. For the simple case of a glider, a modification is simply an alteration to the predictable periodic

pattern of the glider's dynamics. At such points, an observer would be surprised or misinformed about the next value of the glider, having not taken account of the entity about to perturb it. The intuition behind the separable information (Lizier et al. 2010) is that this interpretation is reminiscent of the earlier findings that local apparent transfer entropy $t(i, j, n)$ and local active information storage $a(i, n)$ were negative where the respective information sources were *misinformative* about the next value of the information destination (in the context of the destination's past for transfer entropy). Local active information storage was misinformative at gliders, and local apparent transfer entropy was misinformative at gliders traveling in the orthogonal direction to the measurement in space-time. This being said, one expects that the local apparent transfer entropy measured in the direction of glider motion will be *more informative* about its evolution than any misinformation conveyed from other sources. However, where the glider is modified by a collision with another glider, we would no longer expect the local apparent transfer entropy in its macroscopic direction of motion to remain informative about the dynamics. Assuming that the incident glider is also perturbed, the local apparent transfer entropy in its macroscopic direction of motion will also not be informative about the dynamics at this collision point. We expect the same argument to be true for irregular particles, or domain walls.

As such, we made the hypothesis that at the spatiotemporal location of a local information modification event or collision, *separate* inspection of each information source will *misinform* an observer overall about the next value of the modified information destination. More specifically, the information sources referred to here are the past history of the destination (via the local active information storage) and each other causal information contributor (examined in the context of the past history of the destination, via their local apparent transfer entropies).

We quantified the independent sum of information gained from separate observation of the information storage and information transfer contributors $Y \in V$ to a process X as the **local separable information** $s_X(n)$ (Lizier et al. 2010):

$$s_X(n) = a_X(n) + \sum_{Y \in V, Y \neq X} t_{Y \to X}(n). \tag{1.41}$$

$s_X(n, K)$ is used for finite-k estimates. For CAs, where the causal information contributors are homogeneously within the neighborhood r, we write the local separable information in lattice notation as:

$$s(i, n) = a(i, n) + \sum_{j=-r, j\neq 0}^{+r} t(i, j, n). \tag{1.42}$$

We use $s(i, n, k)$ to represent finite-k estimates, and show $s(i, n, k)$ in Fig. 1.7b.

As inferred earlier, we expected the local separable information to be *positive* or *highly separable* where separate observations of the information contributors are informative overall regarding the next value of the destination. This was be interpreted as a trivial information modification, because information storage and transfer are not interacting in any significant manner. More importantly, we expected the local separable

information to be *negative* at spatiotemporal points where an information modification event or collision takes place. Here, separate observations are misleading overall because a *non-trivial information modification* is taking place (i.e. the information storage and transfer are interacting).

Importantly, this formulation of non-trivial information modification aligns with the descriptions of complex systems as consisting of (a large number of) elements interacting in a *non-trivial* fashion (Prokopenko et al. 2009), and of emergence as where *"the whole is greater than the sum of its parts"*. "The whole" meant to refer to examining all information sources together; the whole is greater where all information sources must be examined together in order to receive positive information on the next value of the examined entity. The thinking behind the separable information was in the direction of measuring synergies between information storage and transfer sources, prior to the development of a proper framework for examining such synergies (Williams and Beer 2010), as discussed in Section 1.6.3.

1.6.2 Local separable information results

Next, we review the application of the separable information to several ECA rules from (Lizier et al. 2010). The simple gliders in ECA rule 54 give rise to relatively simple collisions which we focus on in our discussion here. Notice that the positive values of $s(i, n, k = 16)$ for rule 54 (displayed in Fig. 1.2f) are concentrated in the domain regions and at the stationary gliders (α and β). As expected, these regions are undertaking trivial computations only. The negative values of $s(i, n, k = 16)$ are also displayed in Fig. 1.2f, with their positions marked. The dominant negative values are clearly concentrated around the areas of collisions between the gliders, including collisions between the traveling gliders only (marked by "A") and between the traveling gliders and the stationary gliders (marked by "B" and "C").

Collision "A" involves the γ^+ and γ^- particles interacting to produce a β particle ($\gamma^+ + \gamma^- \rightarrow \beta$ (Hordijk et al. 2001)). The only information modification point highlighted is one time step below (or delayed from) that at which the gliders initially appear to collide (see close-up of raw values in Fig. 1.3). The periodic pattern in the past of the destination breaks there, however the neighboring sources are still able to support separate prediction of the value (i.e. $a(i, n, k = 16) = -1.09$ bits, $t(i, j = 1, n, k = 16) = 2.02$ bits and $t(i, j = -1, n, k = 16) = 2.02$ bits, giving $s(i, n, k = 16) = 2.95$ bits). This is no longer the case however where our measure has successfully identified the modification point; there we have $a(i, n, k = 16) = -3.00$ bits, $t(i, j = 1, n, k = 16) = 0.91$ bits and $t(i, j = -1, n, k = 16) = 0.90$ bits, with $s(i, n, k = 16) = -1.19$ bits suggesting a non-trivial information modification. A delay is also observed before the identified information modification points of collision types "B" and "C"; possibly these delays represent a time-lag of information processing. Not surprisingly, the results for these other collision types imply that the information modification points are associated with the creation of new behavior: in "B" and "C" these occur along the newly created γ gliders, and for "C" in the new α blinkers.

Importantly, weaker non-trivial information modification points continue to be identified at every second point along all the γ^+ and γ^- particles after the initial collisions. These can also be seen for a similar (right-moving) glider in rule 110 in Fig. 1.4f). This was unexpected from our earlier hypothesis. However, these events can be understood as non-trivial computations of the continuation of the glider in the *absence* of a collision; in effect they are virtual collisions between the real glider and the absence of an incident glider. Interestingly, this finding is analogous to the small but non-zero information transfer in periodic domains indicating the absence of gliders.

We also note that measurements of local separable information must be performed with a reasonably large value of k. Here, using $k < 4$ could not distinguish any information modification points clearly from the domains and particles, and even $k < 8$ could not distinguish *all* the modification points (results not shown). Correct quantification of information modification requires satisfactory estimates of information storage and transfer, and accurate distinction between the two.

We observe similar results in $s(i, n, k = 10)$ for ϕ_{par} (see Fig. 1.5f). Note that the collisions at the left and right of the figure do in fact contain significant negative values of $s(i, n, k = 10)$ – around 1 to 2 bits – however these are difficult to see in comparison to the much larger negative value at the collision in the centre of the diagram. These results confirm the particle collisions here as non-trivial information modification events, and this therefore completes the evidence for all of the conjectures about this human understandable computation.

The results for $s(i, n, k = 16)$ for ECA rule 110 (see Fig. 1.4f) are also similar to those for rule 54. Here, we have collisions "A" and "B" which show non-trivial information modification points slightly delayed from the collision in a similar fashion to those for rule 54. We note that collisions between some of the more complex glider structures in rule 110 (not shown) exhibit non-trivial information modification points which are more difficult to interpret, and which are even more delayed from the initiation of the collision. The larger delay is perhaps this is a reflection of the more complex gliders requiring more time steps for the processing to take place. An interesting result not seen for rule 54 is a collision where an incident glider is absorbed by a blinker, without any modification to the absorbing glider (not shown here, see (Lizier et al. 2010)). No information modification is detected for this absorption event by $s(i, n, k = 16)$: this is as expected because the information storage for the absorbing blinker is sufficient to predict the dynamics at this interaction.

As a further test of the measure, we examined collisions between the domain walls of rule 18; see (Lizier et al. 2010). We found that collisions between the domain walls were quite clearly highlighted as the dominant information modification events for this rule - importantly, this result provides evidence that collision of *irregular* particles are information modification events, as expected. The reader is referred to (Lizier et al. 2010) for further discussion of the information modification dynamics of rule 18.

We also apply $s(i, n, k = 16)$ to ECA rule 22, as displayed in Fig. 1.6f. As could be expected from our earlier results, there are many points of both positive and negative local separable information here. The presence of negative values implies the occurrence of non-trivial information modification, yet there does not appear to be any structure to these profiles. Again, this aligns well with the lack of coherent structure

found using the other measures in this framework and from the local statistical complexity profile of rule 22 (Shalizi et al. 2006).

1.6.3 Outlook for information modification

Here, we have reviewed the local separable information, which attempts to quantify information modification at each spatiotemporal point in a complex system. The separable information suggests that information modification events occur where the separable information is negative, indicating that separate or independent inspection of the causal information sources (in the context of the destination's past) is misleading because of non-trivial interaction between these sources. The local separable information was demonstrated to provide the first quantitative evidence that **particle collisions in CAs are the dominant information modification events** therein, and is capable of identifying events involving both creation and destruction.

With that said however, it has been shown that the separable information double-counts parts of the information in the next state of the destination (Flecker et al. 2011). This is clear, and so it is a heuristic more than a measure. Efforts to properly quantitatively define information modification, by combining information dynamics with the partial information decomposition approach (Williams and Beer 2010) to properly measure synergies between information storage and transfer, are ongoing and described in (Lizier et al. 2013). While the separable information is not a proper information-theoretic measure, it remains the only technique which has uniquely filtered particle collision events.

1.7 Importance of coherent computation

Our framework has proven successful in locally identifying the component operations of distributed computation. We then considered in (Lizier et al. 2012b) whether this framework can provide any insights into the overall complexity of computation. In other words, what can our results say about the difference in the complex computations of rules 110 and 54 as compared to rule 22 and others? We review those considerations in this section.

We observed that the *coherence* of local computational structure appears to be the most significant differentiator here. "Coherence" implies a property of sticking together or a logical relationship (Oxford English Dictionary 2008): in this context we use the term to describe a logical spatiotemporal relationship between values in local information dynamics profiles. For example, the manner in which particles give rise to similar values of local transfer entropy amongst spatiotemporal neighbors is coherent. From the spatiotemporal profiles presented here, we note that rules 54 and 110 exhibit the largest amount of coherent computational structure, with rule 18 containing a smaller amount of less coherent structure. Rules 22 and 30 (results for rule 30 not shown, see (Lizier et al. 2012b)) certainly exhibit all of the elementary functions of computation, but do not appear to contain any coherent structure to their computations. This aligns well with similar explorations of local information structure for these

rules, e.g. by Shalizi et al. (2006). Using language reminiscent of Langton's analysis (Langton 1990), we suggested that complex systems exhibit very *highly-structured coherent* computation in comparison to ordered systems (which exhibit coherence but minimal structure in a computation dominated by information storage) and chaotic systems (whose computations are dominated by rampant information transfer eroding any coherence).

Coherence may also be interpreted as a logical relationship *between* profiles of the individual local information dynamics (as three axes of complexity) rather than only within them. To investigate this possibility, Fig. 1.8 plots state-space diagrams of the local apparent transfer entropy for $j = 1$ versus local active information storage (after (Lizier et al. 2012b)). Each point in these diagrams represents the local values of each measure at one spatiotemporal point, thereby generating a complete state-space for the CA. Such state-space diagrams are known to provide insights into structure that are not visible when examining either measure in isolation; for example, in examining structure in classes of systems (such as logistic maps), Feldman et al. (2008) demonstrate that plotting average excess entropy versus entropy rate (while changing a system parameter) reveals loci of the two which are not clear from observing either in isolation. Here however we are looking at structure *within* a single system rather than across a class of systems.

The state-space diagram for rule 110 (Fig. 1.8a) exhibits interesting structure, with significant clustering around certain areas and lines in the state space, reflecting its status as a complex rule. (The two diagonal lines are upper limits representing the boundary condition $t^c(i, j = -1, n, k = 16) \geq 0$ for both destination states "0" and "1"). Rule 54 (Fig. 1.8b) exhibits similar structure in its state-space diagram. On the other hand, the example state space diagram for rule 30 (Fig. 1.8c) exhibits minimal structure (apart from the mathematical upper limit), with a smooth spread of points across the space reflecting its underlying chaotic nature. From the apparent absence of coherent structure in its space-time information profiles, one may expect state-space diagrams for rule 22 to exhibit a similar absence of structure to rule 30. As shown by Fig. 1.8d however this is not the case: the state-space diagram for rule 22 exhibits significant structure, with similar clustering to that of rules 110 and 54.

Importantly, the apparent information structure in the state-space diagrams lends some credence to the claims of complex behavior for rule 22 discussed in Section 1.3.3. However it is a very subtle type of structure, not complex enough to be revealed in the individual local information profiles shown here or by other authors (e.g. by Shalizi et al. (2006)). The structure does not appear to be coherent in these individual profiles, though the state space diagrams indicate a coherent relationship between the local information dynamics which may underpin coherent computation at other scales.

There are certain clues as to the type of coherence which may be displayed by rule 22. Fig. 1.6e does appear to have some traces of coherent transfer entities moving diagonally in space-time; however these seem to be distributed through the CA, seemingly without structure or interactions. More concretely, Grassberger (1983) observed that for rule 22, "there are (at least) four different sets of ordered states, corresponding to $S_i(t) = 0$ for all even/odd i and all even/odd t" – i.e. *rule 22 does have a domain pattern which self-replicates* (with four possible configurations, just offset from each

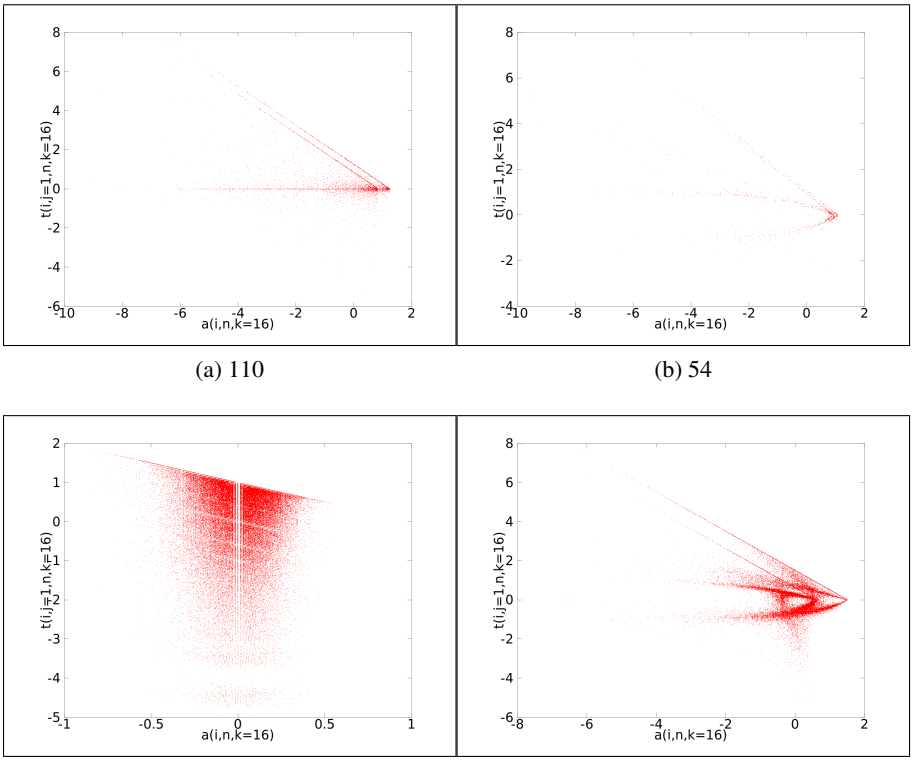

(a) 110 (b) 54

(c) 30 (d) 22

Fig. 1.8: State space diagrams of local transfer entropy (one step to the right) $t(i, j = 1, n, k = 16)$ versus local active information $a(i, n, k = 16)$ (both in bits) at the same space-time point (i, n) for several ECA rules: (a) 110, (b) 54, (c) 30 and (d) 22 (after (Lizier et al. 2012b)).

other in space and time). Indeed, ϵ-machines have been generated to recognize these domains (Crutchfield et al. 2013). Grassberger (1983) goes on to note that "In contrast to the ordered states of rule 18, these states however are unstable: after implanting a kink in an otherwise ordered state, the kink widens without limit, leaving behind it a seemingly disordered state." That is to say, *this domain pattern does not self-organise and it is not robust to perturbations*. This means that, despite the existence of such domains, they are highly unlikely to be found "in the wild" (i.e. when rule 22 is started from random initial states, as we have done for Fig. 1.6). One could also view this as inferring that "life in one dimension" (the perspective that rule 22 is a 1D projection of the 2D Game of Life (McIntosh 1990)) is less stable than "life in two dimensions".

Coming back to Fig. 1.8d – it is possible that these domain patterns, or small versions of them, are what is detected as a signature of coherent information structure by our methods above. Furthermore, emerging evidence suggests that rule 22 can be set

up in certain initial states which sustain such domains for a longer period, with certain stable domain walls (Crutchfield 2013), and that these domain walls are detected as information transfer by our methods. Our investigations in this area remain ongoing.

Given the subtlety of structure in the bounds of our analysis, and using our mutual information heuristics, at this stage we conclude that the behavior of this rule is less complex than that exhibited by rules 110 and 54. As such, we suggested that coherent information structure is a defining feature of complex computation, and explored a technique for inferring this property using local information dynamics. These state-space diagrams for local information dynamics produced useful visual results and were shown to provide interesting insight into the nature of computation in rule 22.

1.8 Conclusion

In this chapter, we have reviewed our complete quantitative framework for the information dynamics of distributed computation in complex systems. Our framework quantifies the information dynamics in terms of the component operations of universal computation: information storage, information transfer and information modification. Our framework places particular importance on examining computation on a local scale in space and time. While averaged or system-wide measures have their place in providing summarized results, this focus on the local scale is vital for understanding the information dynamics of computation and provides many insights that averaged measures cannot.

We reviewed the application of the framework to cellular automata, an important example because of the weight of previous studies on the nature of distributed computation in these systems. Significantly, our framework provided the first quantitative evidence for the widely accepted conjectures that blinkers provide information storage in CAs, particles are the dominant information transfer agents, and particle collisions are the dominant information modification events. In particular, this was demonstrated for the human-understandable density classification computation carried out by the rule ϕ_{par}. This is a fundamental contribution to our understanding of the nature of distributed computation, and provides impetus for the framework to be used for the analysis and design of other complex systems.

The application to CAs aligned well with other methods of filtering for complex structure in CAs. However, our work is distinct in that it provides several different views of the system corresponding to each type of computational structure. In particular, the results align well with the insights of computational mechanics, underlining the strong connection between these approaches.

From our results, we also observed that coherent local information structure is a defining feature of complex distributed computation, and used local information state-spaces to study coherent complex computation. Here, our framework provides further insight into the nature of computation in rule 22 with respect to the accepted complex rules 54 and 110. Certainly rule 22 exhibits all of the elementary functions of computation, yet (in line with Shalizi et al. (2006)) there is no apparent coherent structure to the profiles of its local information dynamics (“in the wild” at least). On the other

hand, state space views of the interplay between these local information dynamics reveal otherwise hidden structure. Our framework is unique in its ability to resolve both of these aspects. We conclude that rule 22 exhibits more structure than chaotic rules, yet the subtlety of this structure prevents it from being considered as complex as rules 110 and 54.

The major thrust of our work since the presentation of this framework was to apply it to other systems, because the information-theoretic basis of this framework makes it readily applicable as such. For example, we have used the measures in this framework to: quantitatively demonstrate coherent waves of motion in flocks and swarms as information cascades (Wang et al. 2012); evolve a modular robot for maximal information transfer between components, observing the emergence of glider-like information cascades (Lizier et al. 2008a); and to study interactions in robotic football and the relation of information measures to success on the field (Cliff et al. 2013). We have also inferred information structure supporting cognitive tasks using fMRI brain imaging data (Lizier et al. 2011a), and studied how the computational capabilities of artificial neural networks relate to underlying parameters and ability to solve particular tasks (Boedecker et al. 2012). We have also made more specific investigations of the relationship between underlying network structure and computational capabilities, including: revealing that intrinsic information storage and transfer capabilities are maximized near the phase transition in dynamics for random Boolean networks (Lizier et al. 2008b); showing that regular networks are generally associated with information storage, random networks with information transfer, and small-world networks exhibit a balance of the two (Lizier et al. 2011b); revealing that feedback and feedforward loop motifs determine information storage capability (Lizier et al. 2012a); and exploring how these information measures relate to synchronization capability of network structures (Ceguerra et al. 2011). We have also explored the relationship of the framework to the context of the observer (Lizier and Mahoney 2013), and provided thermodynamic interpretations of transfer entropy (Prokopenko et al. 2013) and related information-theoreic quantities (Prokopenko et al. 2011). And finally, we have begun reformulating our approach to information modification in seeking a proper measure rather than a heuristic (Lizier et al. 2013). Further developments in all of these directions are expected in the future, due to the utility of the framework.

Acknowledgements

The authors thank Melanie Mitchell for helpful comments and suggestions regarding an early version of this manuscript.

References

Adamatzky, A., editor (2002). *Collision-Based Computing*. Springer-Verlag, Berlin.

Atick, J. J. (1992). Could information theory provide an ecological theory of sensory processing? *Network: Computation in Neural Systems*, 3(2):213.

Badii, R. and Politi, A. (1997). Thermodynamics and Complexity of Cellular Automata. *Physical Review Letters*, 78(3):444.

Bialek, W., Nemenman, I., and Tishby, N. (2001). Complexity through nonextensivity. *Physica A: Statistical Mechanics and its Applications*, 302(1-4):89–99.

Boccara, N., Nasser, J., and Roger, M. (1991). Particlelike structures and their interactions in spatiotemporal patterns generated by one-dimensional deterministic cellular-Automaton rules. *Physical Review A*, 44(2):866–875.

Boedecker, J., Obst, O., Lizier, J. T., Mayer, N. M., and Asada, M. (2012). Information processing in echo state networks at the edge of chaos. *Theory in Biosciences*, 131(3):205–213.

Brown, J. A. and Tuszynski, J. A. (1999). A review of the ferroelectric model of microtubules. *Ferroelectrics*, 220:141–156.

Casti, J. L. (1991). Chaos, Gödel and truth. In Casti, J. L. and Karlqvist, A., editors, *Beyond belief: randomness, prediction and explanation in science*, pages 280–327. CRC Press, Boca Raton.

Ceguerra, R. V., Lizier, J. T., and Zomaya, A. Y. (2011). Information storage and transfer in the synchronization process in locally-connected networks. In *Proceedings of the 2011 IEEE Symposium on Artificial Life (ALIFE)*, pages 54–61. IEEE.

Cliff, O. M., Lizier, J. T., Wang, X. R., Wang, P., Obst, O., and Prokopenko, M. (2013). Towards quantifying interaction networks in a football match. In *Proceedings of the RoboCup'2013 Symposium*. To be published.

Conway, J. H. (1982). What is Life? In Berlekamp, E., Conway, J. H., and Guy, R., editors, *Winning ways for your mathematical plays*, volume 2, ch. 25, pages 927–962. Academic Press, New York.

Cook, M. (2004). Universality in Elementary Cellular Automata. *Complex Systems*, 15(1):1–40.

Couzin, I. D., James, R., Croft, D. P., and Krause, J. (2006). Social Organization and Information Transfer in Schooling Fishes. In Brown, C., Laland, K. N., and Krause, J., editors, *Fish Cognition and Behavior*, Fish and Aquatic Resources, pages 166–185. Blackwell Publishing.

Cover, T. M. and Thomas, J. A. (1991). *Elements of Information Theory*. Wiley-Interscience, New York.

Crutchfield, J. P. (2013). Private communication.

Crutchfield, J. P., Ellison, C. J., and Riechers, P. M. (2013). Exact complexity: The spectral decomposition of intrinsic computation. arXiv:1309.3792.

Crutchfield, J. P. and Feldman, D. P. (2003). Regularities Unseen, Randomness Observed: Levels of Entropy Convergence. *Chaos*, 13(1):25–54.

Crutchfield, J. P. and Young, K. (1989). Inferring statistical complexity. *Physical Review Letters*, 63(2):105–108.

Edmundson, D. E. and Enns, R. H. (1993). Fully 3-dimensional collisions of bistable light bullets. *Optics Letters*, 18:1609–1611.

Eppstein, D. (2002). Searching for spaceships. In Nowakowski, R. J., editor, *More Games of No Chance*, volume 42 of *MSRI Publications*, pages 433–453. Cambridge Univ. Press.

Fano, R. M. (1961). *Transmission of information: a statistical theory of communications*. M.I.T. Press, Cambridge, MA, USA.

Feldman, D. P., McTague, C. S., and Crutchfield, J. P. (2008). The organization of intrinsic computation: Complexity-entropy diagrams and the diversity of natural information processing. *Chaos*, 18(4):043106.

Flecker, B., Alford, W., Beggs, J. M., Williams, P. L., and Beer, R. D. (2011). Partial information decomposition as a spatiotemporal filter. *Chaos*, 21(3):037104+.

Goh, K. I. and Barabási, A. L. (2008). Burstiness and memory in complex systems. *Europhysics Letters*, 81(4):48002.

Grassberger, P. (1983). New mechanism for deterministic diffusion. *Physical Review A*, 28(6):3666.

Grassberger, P. (1986a). Long-range effects in an elementary cellular automaton. *Journal of Statistical Physics*, 45(1-2):27–39.

Grassberger, P. (1986b). Toward a quantitative theory of self-generated complexity. *International Journal of Theoretical Physics*, 25(9):907–938.

Grassberger, P. (1989). Information content and predictability of lumped and distributed dynamical systems. *Physica Scripta*, 40(3):346.

Gray, L. (2003). A Mathematician Looks at Wolfram's New Kind of Science. *Notices of the American Mathematical Society*, 50(2):200–211.

Gutowitz, H. and Domain, C. (1997). The Topological Skeleton of Cellular Automaton Dynamics. *Physica D*, 103(1-4):155–168.

Hanson, J. E. and Crutchfield, J. P. (1992). The Attractor-Basin Portait of a Cellular Automaton. *Journal of Statistical Physics*, 66:1415–1462.

Hanson, J. E. and Crutchfield, J. P. (1997). Computational mechanics of cellular automata: An example. *Physica D*, 103(1-4):169–189.

Helvik, T., Lindgren, K., and Nordahl, M. G. (2004). Local information in one-dimensional cellular automata. In Sloot, P. M. A., Chopard, B., and Hoekstra, A. G., editors, *Proceedings of the International Conference on Cellular Automata for Research and Industry, Amsterdam*, volume 3305 of *Lecture Notes in Computer Science*, pages 121–130, Berlin/Heidelberg. Springer.

Hordijk, W., Shalizi, C. R., and Crutchfield, J. P. (2001). Upper bound on the products of particle interactions in cellular automata. *Physica D*, 154(3-4):240–258.

Jakubowski, M. H., Steiglitz, K., and Squier, R. (1997). Information transfer between solitary waves in the saturable Schrödinger equation. *Physical Review E*, 56(6):7267.

Jakubowski, M. H., Steiglitz, K., and Squier, R. K. (2001). Computing with solitons: A review and prospectus. *Multiple-Valued Logic*, 6(5-6):439–462.

Kinouchi, O. and Copelli, M. (2006). Optimal dynamical range of excitable networks at criticality. *Nature Physics*, 2(5):348–351.

Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2004). Tracking Information Flow through the Environment: Simple Cases of Stigmergy. In Pollack, J., Bedau, M., Husbands, P., Ikegami, T., and Watson, R. A., editors, *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems (ALife IX), Boston, USA*, pages 563–568, Cambridge, MA, USA. MIT Press.

Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2005). All Else Being Equal Be Empowered. In Capcarrère, M. S., Freitas, A. A., Bentley, P. J., Johnson, C. G., and Timmis, J., editors, *8th European Conference on Artificial Life (ECAL 2005)*, volume 3630 of *Lecture Notes in Computer Science*, pages 744–753, Berlin, Heidelberg. Springer Berlin / Heidelberg.

Lafusa, A. and Bossomaier, T. (2005). Hyperplane Localisation of Self-Replicating and Other Complex Cellular Automata Rules. In *Proceedings of the The 2005 IEEE Congress on Evolutionary Computation, Edinburgh*, volume 1, pages 844–849. IEEE Press.

Langton, C. G. (1990). Computation at the edge of chaos: phase transitions and emergent computation. *Physica D*, 42(1-3):12–37.

Lindgren, K. and Nordahl, M. G. (1988). Complexity Measures and Cellular Automata. *Complex Systems*, 2(4):409–440.

Lindgren, K. and Nordahl, M. G. (1990). Universal computation in simple one-dimensional cellular automata. *Complex Systems*, 4:299–318.

Lindner, M., Vicente, R., Priesemann, V., and Wibral, M. (2011). TRENTOOL: A Matlab open source toolbox to analyse information flow in time series data with transfer entropy. *BMC Neuroscience*, 12(1):119+.

Lizier, J. T. (2012). JIDT: An information-theoretic toolkit for studying the dynamics of complex systems. https://code.google.com/p/information-dynamics-toolkit/.

Lizier, J. T. (2013). *The Local Information Dynamics of Distributed Computation in Complex Systems*. Springer Theses. Springer, Berlin / Heidelberg.

Lizier, J. T., Atay, F. M., and Jost, J. (2012a). Information storage, loop motifs, and clustered structure in complex networks. *Physical Review E*, 86(2):026110+.

Lizier, J. T., Flecker, B., and Williams, P. L. (2013). Towards a synergy-based approach to measuring information modification. In *Proceedings of the 2013 IEEE Symposium on Artificial Life (ALIFE)*, pages 43–51. IEEE.

Lizier, J. T., Heinzle, J., Horstmann, A., Haynes, J.-D., and Prokopenko, M. (2011a). Multivariate information-theoretic measures reveal directed information structure and task relevant changes in fMRI connectivity. *Journal of Computational Neuroscience*, 30(1):85–107.

Lizier, J. T. and Mahoney, J. R. (2013). Moving frames of reference, relativity and invariance in transfer entropy and information dynamics. *Entropy*, 15(1):177–197.

Lizier, J. T., Pritam, S., and Prokopenko, M. (2011b). Information dynamics in small-world Boolean networks. *Artificial Life*, 17(4):293–314.

Lizier, J. T. and Prokopenko, M. (2010). Differentiating information transfer and causal effect. *European Physical Journal B*, 73(4):605–615.

Lizier, J. T., Prokopenko, M., Tanev, I., and Zomaya, A. Y. (2008a). Emergence of Glider-like Structures in a Modular Robotic System. In Bullock, S., Noble, J., Watson, R., and Bedau, M. A., editors, *Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems (ALife XI), Winchester, UK*, pages 366–373, Cambridge, MA. MIT Press.

Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2007). Detecting Non-trivial Computation in Complex Dynamics. In Almeida, Rocha, L. M., Costa, E., Harvey, I., and Coutinho, A., editors, *Proceedings of the 9th European Conference on Artifi-*

cial Life (ECAL 2007), volume 4648 of *Lecture Notes in Computer Science*, pages 895–904, Berlin / Heidelberg. Springer.

Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2008b). The information dynamics of phase transitions in random Boolean networks. In Bullock, S., Noble, J., Watson, R., and Bedau, M. A., editors, *Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems (ALife XI), Winchester, UK*, pages 374–381, Cambridge, MA. MIT Press.

Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2008c). Local information transfer as a spatiotemporal filter for complex systems. *Physical Review E*, 77(2):026110+.

Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2010). Information modification and particle collisions in distributed computation. *Chaos*, 20(3):037109+.

Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2012b). Coherent information structure in complex computation. *Theory in Biosciences*, 131(3):193–203.

Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2012c). Local measures of information storage in complex distributed computation. *Information Sciences*, 208:39–54.

Lungarella, M. and Sporns, O. (2006). Mapping information flow in sensorimotor networks. *PLoS Computational Biology*, 2(10):e144+.

MacKay, D. J. C. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, Cambridge.

Marinazzo, D., Wu, G., Pellicoro, M., Angelini, L., and Stramaglia, S. (2012). Information flow in networks and the law of diminishing marginal returns: evidence from modeling and human electroencephalographic recordings. *PloS ONE*, 7(9):e45026+.

Martinez, G. J., Adamatzky, A., and McIntosh, H. V. (2006). Phenomenology of glider collisions in cellular automaton Rule 54 and associated logical gates. *Chaos, Solitons and Fractals*, 28(1):100–111.

McIntosh, H. V. (1990). *Linear Cellular Automata*. Universidad Autónoma de Puebla, Puebla, Mexico.

Mitchell, M. (1998a). A Complex-Systems Perspective on the "Computation vs. Dynamics" Debate in Cognitive Science. In Gernsbacher, M. A. and Derry, S. J., editors, *Proceedings of the 20th Annual Conference of the Cognitive Science Society (Cogsci98), Madison, Wisconsin*, pages 710–715.

Mitchell, M. (1998b). Computation in Cellular Automata: A Selected Review. In Gramss, T., Bornholdt, S., Gross, M., Mitchell, M., and Pellizzari, T., editors, *Non-Standard Computation*, pages 95–140. VCH Verlagsgesellschaft, Weinheim.

Mitchell, M., Crutchfield, J. P., and Das, R. (1996). Evolving Cellular Automata with Genetic Algorithms: A Review of Recent Work. In Goodman, E. D., Punch, W., and Uskov, V., editors, *Proceedings of the First International Conference on Evolutionary Computation and Its Applications, Moscow*, Russia. Russian Academy of Sciences.

Mitchell, M., Crutchfield, J. P., and Hraber, P. T. (1994). Evolving Cellular Automata to Perform Computations: Mechanisms and Impediments. *Physica D*, 75:361–391.

Morgado, R., Cieśla, M., Longa, L., and Oliveira, F. A. (2007). Synchronization in the presence of memory. *Europhysics Letters*, 79(1):10002.

Obst, O., Boedecker, J., Schmidt, B., and Asada, M. (2013). On active information storage in input-driven systems. arXiv:1303.5526.

Oxford English Dictionary (2008). Accessed 8/5/2008, http://www.oed.com/.

Pahle, J., Green, A. K., Dixon, C. J., and Kummer, U. (2008). Information transfer in signaling pathways: a study using coupled simulated and experimental data. *BMC Bioinformatics*, 9:139.

Prokopenko, M., Boschietti, F., and Ryan, A. J. (2009). An Information-Theoretic Primer on Complexity, Self-Organization, and Emergence. *Complexity*, 15(1):11–28.

Prokopenko, M., Gerasimov, V., and Tanev, I. (2006). Evolving Spatiotemporal Coordination in a Modular Robotic System. In Nolfi, S., Baldassarre, G., Calabretta, R., Hallam, J., Marocco, D., Meyer, J. A., and Parisi, D., editors, *Proceedings of the Ninth International Conference on the Simulation of Adaptive Behavior (SAB'06), Rome*, volume 4095 of *Lecture Notes in Artificial Intelligence*, pages 548–559. Springer Verlag.

Prokopenko, M., Lizier, J. T., Obst, O., and Wang, X. R. (2011). Relating Fisher information to order parameters. *Physical Review E*, 84:041116+.

Prokopenko, M., Lizier, J. T., and Price, D. C. (2013). On thermodynamic interpretation of transfer entropy. *Entropy*, 15(2):524–543.

Sánchez-Montañés, M. A. and Corbacho, F. J. (2002). Towards a New Information Processing Measure for Neural Computation. In Dorronsoro, J. R., editor, *Proceedings of the International Conference on Artificial Neural Networks (ICANN 2002), Madrid, Spain*, volume 2415 of *Lecture Notes in Computer Science*, pages 637–642, Berlin/Heidelberg. Springer-Verlag.

Schreiber, T. (2000). Measuring Information Transfer. *Physical Review Letters*, 85(2):461–464.

Shalizi, C. R. (2001). *Causal Architecture, Complexity and Self-Organization in Time Series and Cellular Automata*. PhD thesis, University of Wisconsin-Madison.

Shalizi, C. R. and Crutchfield, J. P. (2001). Computational mechanics: Pattern and Prediction, Structure and Simplicity. *Journal of Statistical Physics*, 104:817–879.

Shalizi, C. R., Haslinger, R., Rouquier, J.-B., Klinkner, K. L., and Moore, C. (2006). Automatic filters for the detection of coherent structure in spatiotemporal systems. *Physical Review E*, 73(3):036104.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27.

Takens, F. (1981). Detecting strange attractors in turbulence. In Rand, D. and Young, L.-S., editors, *Dynamical Systems and Turbulence, Warwick 1980*, volume 898 of *Lecture Notes in Mathematics*, chapter 21, pages 366–381. Springer, Berlin / Heidelberg.

Von Neumann, J. (1966). *Theory of self-reproducing automata*. University of Illinois Press, Urbana.

Wang, X. R., Miller, J. M., Lizier, J. T., Prokopenko, M., and Rossi, L. F. (2012). Quantifying and Tracing Information Cascades in Swarms. *PLoS ONE*, 7(7):e40084+.

Wibral, M., Pampu, N., Priesemann, V., Siebenhühner, F., Seiwert, H., Lindner, M., Lizier, J. T., and Vicente, R. (2013). Measuring Information-Transfer delays. *PLoS ONE*, 8(2):e55809+.

Wibral, M., Rahm, B., Rieder, M., Lindner, M., Vicente, R., and Kaiser, J. (2011). Transfer entropy in magnetoencephalographic data: quantifying information flow in cortical and cerebellar networks. *Progress in Biophysics and Molecular Biology*, 105(1-2):80–97.

Williams, P. L. and Beer, R. D. (2010). Nonnegative Decomposition of Multivariate Information. arXiv:1004.2515.

Wolfram, S. (1984a). Cellular automata as models of complexity. *Nature*, 311(5985):419–424.

Wolfram, S. (1984b). Computation theory of cellular automata. *Communications in Mathematical Physics*, 96(1):15–57.

Wolfram, S. (1984c). Universality and complexity in cellular automata. *Physica D*, 10(1-2):1–35.

Wolfram, S. (2002). *A New Kind of Science*. Wolfram Media, Champaign, IL, USA.

Wuensche, A. (1999). Classifying cellular automata automatically: Finding gliders, filtering, and relating space-time patterns, attractor basins, and the Z parameter. *Complexity*, 4(3):47–66.

Yamada, T. and Aihara, K. (1994). Spatio-temporal complex dynamics and computation in chaotic neural networks. In *Proceedings of the IEEE Symposium on Emerging Technologies and Factory Automation (ETFA '94), Tokyo*, pages 239–244. IEEE.