

Emergence of Glider-like Structures in a Modular Robotic System

Joseph T. Lizier^{1,2}, Mikhail Prokopenko¹, Ivan Tanev³ and Albert Y. Zomaya²

¹CSIRO Information and Communications Technology Centre, Locked Bag 17, North Ryde, NSW 1670, Australia

²School of Information Technologies, The University of Sydney, NSW 2006, Australia

³Department of Information Systems Design, Doshisha University, 1-3 Miyakodani, Tatara, Kyotanabe, Kyoto 610-0321, Japan
jlizier@it.usyd.edu.au

Abstract

Information-driven evolutionary design has been proposed as an efficient method for designing self-organized multi-agent systems. *Information transfer* is known to be an important component of distributed computation in many complex systems, and indeed it has been suggested that maximization of information transfer can give rise to interesting behavior and induce necessary structure in a system. In this paper, we report the first known application of a direct measure of information transfer, *transfer entropy*, as a fitness function to evolve a self-organized multi-agent system. The system evolved here is a simulated snake-like modular robot. In the most fit snakebot in the final generation, we observe coherent traveling information transfer structures. These are analogous to gliders in cellular automata, which have been demonstrated to represent the coherent transfer of information across space and time, and play an important role in facilitating distributed computation. These observations provide evidence that using information transfer to drive evolutionary design can produce useful structure in the underlying system.

Introduction

The principle of self-organization is well known to offer the advantages of flexibility, robustness and scalability over centralized system designs (Prokopenko et al., 2006a). Most self-organized solutions are currently designed using a genetic algorithm of some form, with fitness functions measuring achievement of the task required of the system (*task-based evolution*). Several authors have recently been investigating the potential for *information-driven evolutionary design* to push the advantages of self-organization even further, e.g. (Prokopenko et al., 2006a; Polani et al., 2007; Klyubin et al., 2005; Sporns and Lungarella, 2006). This concept proposes the use of information-theoretical measures of the information processing carried out by the system as generic fitness functions in evolutionary design. From an engineering perspective, template-based evolution for generic information processing skills could be simpler and afford a framework based approach to such design of self-organized systems. It also provides to us the potential to better understand the evolved solutions, and more importantly the opportunity to study and understand the emergence rather than engineering of intelligence (Polani et al., 2007).

We believe information-driven self-organization is best facilitated using measures of the information dynamics of distributed computation (Lizier et al., 2007). Any task we wish to evolve the system to solve involves a distributed computation, so evolving for the fundamental building blocks of the computation is a direct way to allow that computation to emerge. We could evolve directly for a particular computational property (e.g. information storage as opposed to transfer), or for a mix of those properties.

Information transfer has been suggested to be a particularly important fitness function here. It has been conjectured that information transfer can give rise to interesting behavior and induce necessary structure in a multi-agent system (Prokopenko et al., 2006a). One inspiration of this viewpoint is the concept of *empowerment* (Klyubin et al., 2005), which refers to an agent's self-perception of its influence over its environment. Alluding to (but not directly measuring) information transfer, it is quantified as the channel capacity between an agent's actuators and sensors through the environment. Maximization of empowerment has been suggested to be an intrinsic selection pressure¹. With or without the presence of explicit actuator-sensor channels, we expect information transfer to be a useful fitness function because of its important role in distributed computation.

Here, we present the first experiment of the use of a direct measure of information transfer, transfer entropy (Schreiber, 2000), as the sole fitness function in an evolutionary design task. An initial aim of the experiment is to check whether information transfer underpins co-ordinated motion, as was suggested in previous work (Prokopenko et al., 2006a). More importantly, we aim to investigate what type of behavior emerges when a system is evolved to maximize information transfer. Much previous work on information-driven evolution has sought to confirm whether it can approximate direct evolution for a given task. Here, we simply seek to investigate what type of solution or computation

¹The justification or otherwise of the suggestion that *natural* evolution is driven by the intrinsic forces of information processing is irrelevant to whether information-driven evolutionary design can be used as a successful tool for *artificial* systems.

is generated by evolution for information transfer, and hypothesize that it will induce useful computation in the system. Our findings will help us to understand the role that information transfer can play in a unified framework for information-driven evolutionary design, focusing on the information dynamics of distributed computation.

We use a snake-like modular robot (the *snakebot*) for experimentation: information structure has been observed to emerge previously with a fitness function for fastest motion (Prokopenko et al., 2006b), and conversely fast motion has emerged from evolution with a measure of co-ordination as the fitness function (Prokopenko et al., 2006a). We measure information transfer using the transfer entropy (Schreiber, 2000) between neighboring modules of the snakebot, and evolve the snakebot to maximize this quantity. Information transfer in this fashion could be utilized by the snake in leading to co-ordinated motion between the modules, communicating information about obstacles, or driving new behaviors in a given direction along the snake.

We report that *coherent* traveling information transfer structures were observed to emerge (using local transfer entropy (Lizier et al., 2008a)) in the evolved snakebot. We say “emerged” because while high information transfer was selected for, local coherent structures were not part of the specification. This is an important finding, because these structures are analogous to *glider* structures in cellular automata (CAs). Gliders are known to be the information transfer agents in CAs, providing for long-range correlations across space and time and playing a fundamental role in the distributed computation carried out in the CA (Lizier et al., 2008a). As such, we have provided evidence that using a direct measure of information transfer as a fitness function in information-driven evolutionary design can indeed produce useful structure in the system.

Information-driven evolution

Task-based evolution, the incumbent method of designing self-organized systems, can be impractical. Hand-crafting fitness functions for every task can be time-consuming and tedious, and requires specialized human understanding of the task. It has the potential to under-specify the problem (thereby solving a different task) or perhaps over-specify it (leading to an inflexible design). Also, the intelligent designer may not be completely sure of how to measure performance of the required task, or this may be difficult (e.g. measuring speed may require extra sensors). Furthermore, if the initial task-based fitness landscape is flat and features no gradients, task-based evolution has no foothold around which to begin designing a solution. Finally, evolution often delivers intricate solutions for which (human) system managers cannot understand the inner workings: this is particularly undesirable for critical systems where maintenance or prediction of behavior is required.

As an alternative, information-driven evolutionary design

proposes the use of information-theoretic measures to design the required information processing structure in self-organized systems. This has been prompted by observations of complexity to grow or necessary information-theoretic structure to emerge during task-based evolution. Growth of complexity during evolution has been observed by Adami (2002) (measuring “physical complexity” in the Avida simulation system) and Yaeger and Sporns (2006) (measuring neural complexity of evolved agents in the PolyWorld simulation system). Looking at evolution for particular tasks, Prokopenko et al. (2006b) observed co-ordination (measured as excess entropy (Crutchfield and Feldman, 2003)) to increase in snakebots evolved for maximum velocity, and Baldassare et al. (2004) observed a decrease in entropy in a swarm evolved for co-ordinated motion.

These observations suggest that such information-theoretic metrics could be used themselves in information-driven evolutionary design. This idea is fundamentally based on the theory that information structure is vital to the emergence of self-organized intelligence (Polani et al., 2007). The concept could provide a consistent framework for the evolutionary design of self-organized systems, using template-based evolution for required computational tasks. This framework would be able to produce useful structure where task-based evolution faces initially flat task-based fitness landscapes, perhaps serving as a platform from which to launch better-equipped task-based evolution. Furthermore, it may provide solutions which are simpler for humans to understand in terms of the underlying information dynamics. Perhaps most important is the potential for this approach to provide insight into the emergence rather than engineering of intelligence (Polani et al., 2007), and thereby facilitate unsupervised learning.

Several examples of successful information-driven evolutionary design exist in the literature. Maximization of empowerment has been shown to induce a necessary structure in agent’s behavior by Klyubin et al. (2005). Sporns and Lungarella (2006) have evolved hand-eye co-ordination to grab a moving object using maximization of neural complexity, and demonstrated that this solution contained more intrinsic diversity than solutions from task-driven evolution; the increased diversity may afford greater flexibility to the system. Prokopenko et al. (2006a) were able to evolve fast-moving snakebots using maximization of an information-theoretic measure of co-ordination. Also, Sperati et al. (2007) have observed interesting periodic behavior and complex structure in groups of robots which were evolved to maximize their mutual information.

We suggest that the information dynamics of distributed computation (Lizier et al., 2007, 2008a) provide the most intuitive basis for information-driven evolution. These information dynamics are the primitive functions of Turing universal computation, i.e. *information storage, transfer* and *modification*. Any task we wish the system to achieve in-

volves some form of computation. As such, using a framework for distributed computation allows us to target the evolution toward the computational requirements of the task at hand, i.e. selecting either the most relevant computational function as the fitness function, or balancing the functions in a more complex manner. Importantly, using such a framework provides a basis through which to understand the computation carried out by the evolved solution. Also, guiding a system toward the building blocks of distributed computation is perhaps the most intuitive way to facilitate the emergence of collective intelligence.

Information transfer is an important candidate fitness function here. It has been observed to be a critical part of the dynamics of many complex systems, for example being manifested in dipole-dipole interactions in microtubules which give rise to self-organization there (Brown and Tuszynski, 1999). Another important example are particles or gliders in CAs (e.g. see Fig. 1), which are coherent traveling information transfer structures in those systems (Lizier et al., 2008c). Much importance has been placed on the role of gliders in CA dynamics; in fact, they have been demonstrated to transport information for the distributed computation carried out in CAs (Lizier et al., 2008c). For example in a density-classification task, gliders appear to transport information about the density in the region of the CA where they originated, with glider collisions processing this information to make a decision about the overall density (Mitchell et al., 1994). Information transfer is also related to the concept of empowerment (Klyubin et al., 2005), with much importance placed on the maximization of the capacity of the information channel between an agent’s actuators and sensors here. Importantly also, it has long been conjectured that information transfer is maximized in the vicinity of an order-chaos phase transition (Langton, 1990), where critical dynamics are said to facilitate the emergence of complex computation. Several authors have since inferred this conclusion from related measures (Solé and Valverde, 2001), however evidence from a directed, dynamic measure of information transfer has only recently been provided (Lizier et al., 2008b). In the following section, we describe this measure of information transfer.

Information transfer

Our measure of information transfer is of course found in the domain of information theory (MacKay, 2003), which is proving to be a useful framework for the analysis and design of complex systems, e.g. (Prokopenko et al., 2006a). The fundamental quantity in this domain is the (Shannon) *entropy*, which represents the uncertainty in a sample x of a random variable X : $H_X = -\sum_x p(x) \log_2 p(x)$ (all with units in bits). The *joint entropy* of two random variables X and Y is a generalization to quantify the uncertainty of their joint distribution: $H_{X,Y} = -\sum_{x,y} p(x,y) \log_2 p(x,y)$. The *conditional entropy* of X given Y is the average un-

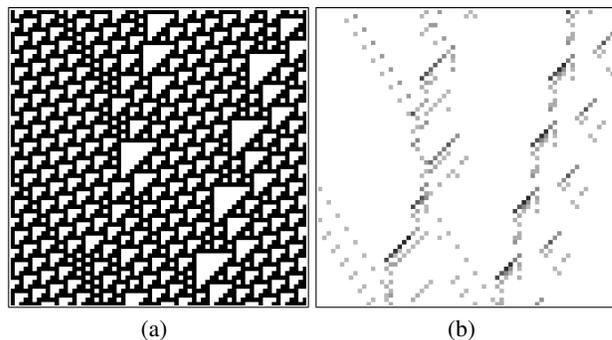


Figure 1: Elementary CA rule 110. (a) Raw states. (b) Local transfer entropy with $k = 16$ (for transfer one step to the left per time step) highlights glider structures.

certainty that remains about x when y is known: $H_{X|Y} = -\sum_{x,y} p(x,y) \log_2 p(x|y)$. The *mutual information* between X and Y measures the average reduction in uncertainty about x that results from learning the value of y , or vice versa: $I_{X;Y} = H_X - H_{X|Y}$. The *conditional mutual information* between X and Y given Z is the mutual information between X and Y when Z is known: $I_{X;Y|Z} = H_{X|Z} - H_{X|Y,Z}$.

The mutual information has previously been used as a de facto measure for information transfer (e.g. by Solé and Valverde (2001)), however this approach is criticized by Schreiber (2000) as a symmetric measure of statically shared information. To address these concerns, Schreiber introduced the transfer entropy to quantify the *information transfer* between a source and a destination agent as the average information provided by the source about the destination’s next state that was not contained in the past of the destination. This formulation provides a properly directional and dynamic measure of information transfer. The transfer entropy is the average mutual information between the previous state of the source² y_n and the next state of the destination x_{n+1} , *conditioned* on the past of the destination $x_n^{(k)}$:

$$T_{Y \rightarrow X}(k) = \sum_{u_n} p(u_n) \log_2 \frac{p(x_{n+1}|x_n^{(k)}, y_n)}{p(x_{n+1}|x_n^{(k)})}. \quad (1)$$

This average is over all state transition tuples $u_n = (x_{n+1}, x_n^{(k)}, y_n)$. From another perspective, it is also an average over a *local transfer entropy* (Lizier et al., 2008c) at all observed time points:

$$t_{Y \rightarrow X}(n+1, k) = \log_2 \frac{p(x_{n+1}|x_n^{(k)}, y_n)}{p(x_{n+1}|x_n^{(k)})}, \quad (2)$$

$$T_{Y \rightarrow X}(k) = \langle t_{Y \rightarrow X}(n, k) \rangle \quad (3)$$

²The transfer entropy can be formulated using the l previous states of the source. However, where only the previous state is a causal information contributor, we set $l = 1$ to measure direct transfer only at step n .

In general, these measures are only completely accurate in the limit $k \rightarrow \infty$ (Lizier et al., 2008c), since this removes all information that was already in the history of the destination from being mistaken as transferred. This is computationally infeasible however, so we use as large a history k as is facilitated by our observation set.

The transfer entropy can also be formulated to condition on the states of all other *causal information contributors* to the destination, so as to completely account for the contribution of the source Y . This form is known as the *complete* transfer entropy (see Lizier et al. (2008c)). The formulation in Eq. (2) is then labeled the *apparent* transfer entropy (note: in this paper, we refer to this form unless otherwise stated).

The transfer entropy has been studied in a number of interesting applications, for example in characterizing information flow in sensorimotor networks by Lungarella and Sporns (2006). Bertschinger et al. (2006) used the transfer entropy to investigate the distinction of a system from its environment, and the autonomy of the system. Studies of the local transfer entropy in CAs provided the first quantitative evidence for the long-held conjecture that gliders are the information transfer agents therein (Lizier et al., 2008c) (see Fig. 1). Application to random boolean networks (RBNs) suggests that the apparent transfer entropy is maximized in the vicinity of a phase transition from ordered to chaotic behavior, while the complete transfer entropy continues increasing into the chaotic regime (Lizier et al., 2008b). Large apparent transfer entropy appears to indicate that the dynamics support coherent information transfer (in the form of gliders in CAs) as an important component of complex distributed computation (Lizier et al., 2008a).

To compute the transfer entropy for continuous variables, a simple approach is to discretize the continuous variables and apply Eq. (1), however with a slight increase in effort, one can remain in the continuous regime. In doing so, Schreiber (2000) recommends using the method of *kernel estimation* to estimate the required probabilities, rather than an approach based on correlation integrals. (The same technique is used under different guises in computing the “pattern entropy” by Dettmann and Cohen (2000) and the “approximate entropy” by Pincus and Singer (1996)). This method has been used, for example, to compute transfer entropy in signal transduction by calcium ions by Pahle et al. (2008). With the kernel estimation method, the joint probability of the state transition tuple $u_n = (x_{n+1}, x_n^{(k)}, y_n)$ for example is estimated by counting similar tuples:

$$\hat{p}_r(u_n) = \frac{1}{N} \sum_{n'} \Theta \left(\left\| \begin{pmatrix} x_{n+1} - x_{n'+1} \\ x_n^{(k)} - x_{n'}^{(k)} \\ y_n - y_{n'} \end{pmatrix} \right\| - r \right), \quad (4)$$

where by default Θ is the step kernel ($\Theta(x > 0) = 1$, $\Theta(x \leq 0) = 0$) using the precision r , and the norm $\|\cdot\|$ is the maximum distance, though other choices are possible. The average transfer entropy $T_{Y \rightarrow X}(k)$ is then computed

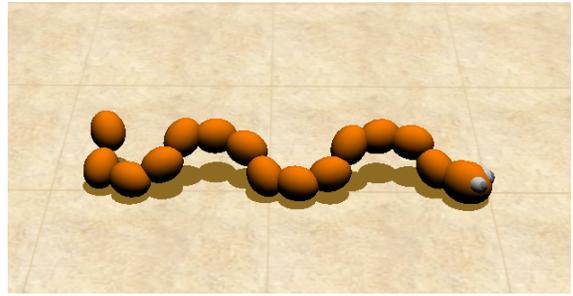


Figure 2: Snakebot

as the average of local transfer entropies (see Eq. (3) and Eq. (2)), where each local transfer entropy uses these kernel estimations to compute the relevant probability distribution functions. That is, computation of the average transfer entropy for continuous variables is *necessarily* a computation over each local point in time rather than over all possible state transition tuples. Here, we will present the first use of the local transfer entropy values for continuous variables.

Evolving the snakebot for maximum information transfer

The snakebot is a snake-like modular robot, introduced in (Tanev et al., 2005), which is simulated in the Open Dynamics Engine (ODE). As shown in Fig. 2, it consists of a set of identical spherical morphological segments which are linked by universal joints. The joints each have two actuators for joint rotation, which are oriented vertically and horizontally in the initial standstill position of the snakebot, and all have identical angle limits. No anisotropic friction between the morphological segments and the surface is considered. The genome for the snakebot is an algebraic expression for the *desired* turning angles of its horizontal and vertical actuators as a function of time and actuator index. The periodic functions \sin and \cos are included in the function set, providing support for periodic gaits. The turning angles however are constrained by interactions between the segments and with the terrain; as such the *actual* actuator angles represent the emergent dynamics. Here, $\alpha_{i,n}$ and $\beta_{i,n}$ represent the actual horizontal and vertical turning angles respectively at time step n , where i is the actuator index (so $1 \leq i \leq S$ where $S = 14$ is the number of joints), and $1 \leq n \leq N$ for $N = 1800$ time steps in the simulation run.

Initial experiments to evolve fastest motion in any direction indicated that side-winding motion (i.e. locomotion predominantly perpendicular to the long axis of the snakebot) provided superior speed characteristics (Tanev et al., 2005). As previously mentioned, subsequent experiments observed the increase in co-ordination (as excess entropy) with this evolution (Prokopenko et al., 2006b), and then evolved similar fast moving side-winding locomotion using this mea-

sure of co-ordination as a fitness function (Prokopenko et al., 2006a). In capturing correlation across space and time, the (two-dimensional) excess entropy is something of an overall measure of distributed computation which balances the underlying components of information storage and transfer. Here, we evolve the snakebot using transfer entropy, in order to maximize the information transfer component of distributed computation. It was suggested in (Prokopenko et al., 2006a) that information transfer underpinned co-ordinated motion. An information transfer is certainly required in a transient sense to achieve co-ordinated motion, but the level of information transfer in this initial phase may not be very significant compared to the information transfer averaged over longer experimental periods for other behaviors. The evolution of the snakebot here will take place in a flat environment. We will observe what types of behavior emerge as a result of selecting for information transfer.

In evaluating the fitness of each snakebot after it is simulated for N time steps, we compute the average transfer entropy $T_{i+1 \rightarrow i}(k)$ between each pair of consecutive modules $i + 1$ and i , in the direction from the tail toward the head (i.e. decreasing module number i). The transfer entropy is computed using the time series of actual horizontal turning angles $\alpha_{i,n}$. Kernel estimation is used with these continuous values, with r set to one quarter of the standard deviation of the turning angles. Also, we use the default step kernel and maximum distance norm, ignoring matched pairs within 20 time steps *and* neighboring modules to avoid spurious dynamic correlations (as recommended by Schreiber (2000)). The direction of tail toward head is selected because each module only applies desired turning angles to the actuators in front of it (i.e. in the direction of the head), thereby giving preferential treatment to information traveling in this direction. Although it is possible for information to be transferred across more than one joint per time step, we consider only consecutive pairs since this is likely to be the dominant transfer mode. Also, as per footnote 2, we only consider transfer from a single previous state of the source variable, so as to consider information transferred directly at the given time step. We use a past history length $k = 30$ (as for the correlation entropy calculations in Prokopenko et al. (2006a)). This is large enough to eliminate information storage from the calculation (see Results), while allowing adequate sampling of the underlying distributions (because the presence of sin and cos functions mean that the emergent turning angle sequences are generally quasi-periodic and therefore much of the state space of $\alpha_{i,n}^{(k)}$ remains unexplored). Our *fitness function* is then the average of these transfer entropies over all $S - 1$ consecutive module pairs for the given snakebot:

$$T_{tail \rightarrow head}(k) = \frac{1}{S-1} \sum_{i=1}^{S-1} T_{i+1 \rightarrow i}(k). \quad (5)$$

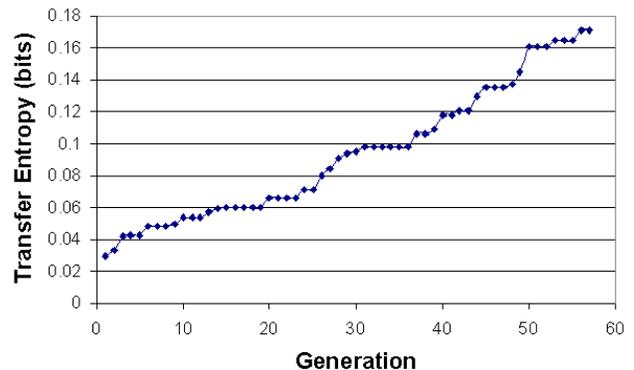


Figure 3: Snakebot fitness (average transfer entropy $T_{tail \rightarrow head}(k = 30)$) per generation, plotted for the best performer in each generation.

The Genetic Programming (GP) techniques used for snakebot evolution are described by Tanev et al. (2005). The snakebots evolve within a population of 200 individuals, with the best performers selecting using the fitness function described above. No minimum limit is placed on how far the snakebot moves, since we are not evolving for fast locomotion. The selection is based on a binary tournament with selection ratio of 0.1 and reproduction ratio of 0.9. Random subtree mutation is used with a ratio of 0.01.

Results and discussion

First, we note that snakebots exhibiting a high degree of co-ordinated motion (as exemplified by most fit individual from (Prokopenko et al., 2006a)) were found to have significantly lower transfer entropy than individuals specifically evolved to maximize transfer entropy (e.g. 0.007 bits versus 0.175 bits for the most fit snakebot here). Highly co-ordinated snakebots exhibited very short transients before becoming co-ordinated, and minimal transfer entropy in their ongoing behavior. Co-ordinated motion is certainly more strongly associated with memory (in fact is a distributed memory (Lizier et al., 2008a)) than information transfer. When neighboring modules achieve perfect co-ordination, they have effectively reached a periodic attractor: their next states are completely predictable from their individual pasts, and so no additional information from the neighbor is measured as transfer entropy. It is possible that transfer entropy might be measured to be higher for snakebots attempting co-ordinated motion in a challenging environment, where information transfer in the longer and more significant transient toward co-ordination may play an important role in the dynamics.

In our evolution of snakebots for transfer entropy, the growth in the average transfer entropy $T_{tail \rightarrow head}(k = 30)$ of the most fit snakebot in each generation is shown in Fig. 3.

We will focus on the most fit individual in the final (57th)

generation as the result of this evolution, which had an average transfer entropy of 0.175 bits between neighboring modules toward the head per time step. This snakebot did not display a fast, well co-ordinated side-winding locomotion. Instead, it displayed a complex form of wriggling behavior, where thrashing of the tail appeared to drive new behavior along the body of the snake, achieving a slow movement to the side³. The dynamics of this behavior are clearer when examining the time-series of the actual horizontal turning angles $\alpha_{i,n}$, as displayed in Fig. 4(a). Here, we see that coherent waves of behavior are consistently traveling along the snakebot, from the tail toward the head. Each wave involves the modules turning in alternating directions along the snake (visible in color image online), reaching a maximum angle then coming back to a rest position. The modules then swap their turning angles in the next wave. Importantly, these waves are not completely periodic, allowing scope for information transfer effects.

Already, we note a fairly clear correspondence to emergent traveling structures in microtubules and gliders in CAs, however to confirm the information transfer properties, we examine the local transfer entropy profile in Fig. 4(b). The local transfer entropy profile here tells us much more about the snakebot dynamics than the average transfer entropy does (as was observed for CAs in (Lizier et al., 2008c)). As expected, we confirm that we have coherent traveling waves of information transfer moving along the snakebot from the tail toward the head, which coincide in direction and approximately in time with the time-series waves previously observed. As an example, note the images of the snakebot in Fig. 5 with modules colored to indicate local transfer entropy (also, videos with the modules of the snake highlighted according to their local transfer entropy are available online, see footnote 3). We can be confident that the information transfer measured is not misattributed information storage, because our use of $k = 30$ considers a longer past history than the length of the time-series waves here. Note that these coherent transfer structures were not observed in fully-coordinated or random snakebots.

There is a wide variation in the types of such information transfer structures observed here: some move faster than others (indicated by a flatter structure), some are more highly localized in time (thinner structures), some contain higher local transfer entropies (darker coloring), and some do not coherently travel the whole way along the body of the snakebot. Importantly, none of these differences are detectable by superficial examination of the time-series of the actual actuator angles. Indeed, apart from their coincidence in direction and approximately in time, there is little correspondence between the time-series waves and the informa-

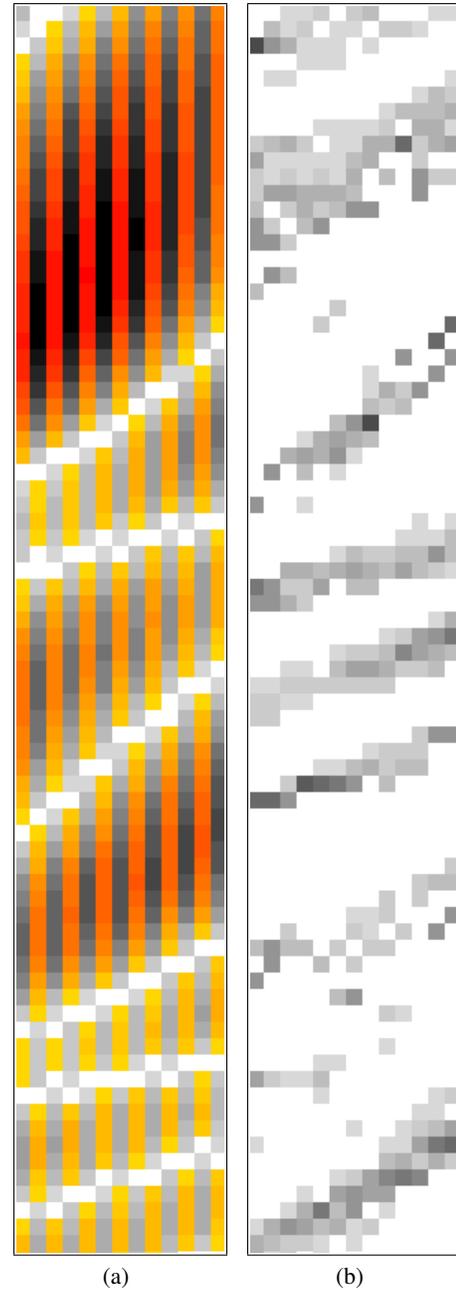


Figure 4: Local apparent transfer entropy highlights “gliders” in the evolved snakebot. (a) Raw actuator turning angles for each of the 13 destination modules (head at left, tail at right) of the snakebot for 76 consecutive time steps (time increases down the page): grayscale represents a positive turning angle, yellow-red (color online) represents a negative turning angle; range is -50 to 50 degrees. (b) Local transfer entropy $t_{i+1-i}(n, k = 30)$ into each of the 13 information destination modules of the snakebot, between consecutive modules in the tail \rightarrow head direction: grayscale, range 0.0 bits (white) to 2.8 bits (black).

³Videos of the snakebot, showing raw motion and local transfer entropy are available at <http://www.it.usyd.edu.au/~jlizier/publications/08ALifeSnakebotTe> or http://www.prokopenko.net/modular_robotics.html

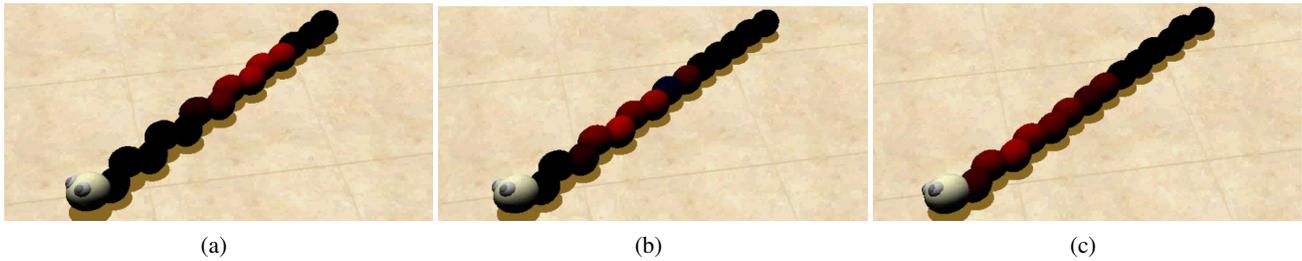


Figure 5: Snakebot modules colored to indicate incoming local transfer entropy (black is 0.0 bits, red is 2.8 bits; color online) from neighboring module toward the tail, for three consecutive time steps. The information transfer from the tail appears to communicate a straightening behavior here.

tion structure that is obvious to the observer. Certainly, there is no simple method of using the time-series waves to infer the location in time of the local information transfer structures: these are observed to begin and end at various time points within the time-series waves. Local transfer entropy reveals the *precise* space-time dynamics of the manner in which the tail drives new behavior in the snakebot in a way not possible by examining the time-series alone.

As coherent traveling local information transfer, these structures are clearly analogous to gliders in CAs (see Fig. 1). This finding is significant because of the important role that gliders play in CA dynamics, where they coherently transfer information relevant to the collective computation of the CA. We previously noted that the coincidence of gliders and coherent information transfer with a maximization of (apparent) transfer entropy (Lizier et al., 2008a). Here, we have demonstrated the emergence of glider-like structures when (apparent) transfer entropy is optimized, *without* explicitly selecting for such local coherence. This suggests that coherent glider-like structures are the most efficient mode of (apparent) information transfer. This has significant implications for glider-like structures observed in natural systems, e.g. dipole-dipole interactions in microtubules (Brown and Tuszynski, 1999), which could have evolved to exploit this efficient mode of information transfer where coherent communication or effect over some distance is beneficial.

The coherence of glider structures is of particular importance to the computation in CAs; without coherence of information transfer, complex computation does not appear to take place (Lizier et al., 2008a,b). A second requirement for such truly distributed computation though is *bidirectional* information transfer. Here, with strong information transfer encouraged in one direction only, although we have demonstrated the emergence of an important building block for non-trivial computation, we have evolved only a trivial type of computation. (This is effectively the reason that there are very few points of negative local transfer entropy measured in the snakebot here). In future work, we will build on our results here to evolve bidirectional information transfer for true distributed computation.

Conclusion

We have presented the first experiment of the use of transfer entropy as a generic fitness function for information-driven evolutionary design. We have demonstrated that maximizing information transfer in this manner can lead to the emergence of *coherent* transfer structures which, as manifested by gliders, are known to underpin distributed computation in CAs. Here, this useful generic skill was not fully capitalized on by the snakebot, but the important finding is that the use of information transfer as a fitness function led to the emergence of this computational capability. Also, our experiment implies that glider-like structures are the most efficient mode of coherent information transfer, which is itself significant insight into the nature of information transfer.

All agent-based systems compute; indeed it is their computation that makes them useful to us. Here, the snake computes where to move. While information transfer does not appear to be important for co-ordinated motion in flat environments, it could underpin computation for tasks such as successful navigation in challenging environments, where different parts of the body could sample many sections of the environment in parallel, and communicate information about the environment along the structure. Information transfer could be used to develop the required computational capability for tasks such as these in future work.

We intend to explore the use of information transfer in information-driven evolutionary design in other settings where bidirectional information transfer may be required for distributed computation. We also intend to investigate the use of the other information dynamics of computation (information storage and modification) (Lizier et al., 2007) in such design, and explore the circumstances under which each should be used and indeed how they can be used together.

Acknowledgements

JL thanks Jürgen Pahle for discussions on local transfer entropy for continuous-valued variables. IT was supported in part by the National Institute of Information and Communications Technology of Japan.

References

- Adami, C. (2002). What is complexity? *Bioessays*, 24(12):1085–1094.
- Baldassare, G., Parisi, D., and Nolfi, S. (2004). Measuring coordination as entropy decrease in groups on linked simulated robots. In *Proceedings of the International Conference on Complex Systems (ICCS2004)*, Boston. To be published.
- Bertschinger, N., Olbrich, E., Ay, N., and Jost, J. (2006). Information and closure in systems theory. In Artmann, S. and Dittrich, P., editors, *Proceedings of the 7th German Workshop on Artificial Life (GWAL-7)*, Jena, Germany, Amsterdam. IOS Press.
- Brown, J. A. and Tuszynski, J. A. (1999). A review of the ferroelectric model of microtubules. *Ferroelectrics*, 220:141156.
- Crutchfield, J. P. and Feldman, D. P. (2003). Regularities unseen, randomness observed: Levels of entropy convergence. *Chaos*, 13(1):25–54.
- Dettmann, C. P. and Cohen, E. G. D. (2000). Microscopic chaos and diffusion. *J. Stat. Phys.*, 101(3):775–817.
- Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2005). All else being equal be empowered. In Capcarrere, M. S., Freitas, A. A., Bentley, P. J., Johnson, C. G., and Timmis, J., editors, *Proceedings of the 8th European Conference on Artificial Life (ECAL 2005)*, Kent, UK, volume 3630 of *Lecture Notes in Computer Science*, pages 744–753, Berlin / Heidelberg. Springer.
- Langton, C. G. (1990). Computation at the edge of chaos: phase transitions and emergent computation. *Physica D*, 42(1-3):12–37.
- Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2007). Detecting non-trivial computation in complex dynamics. In Almeida e Costa, F., Rocha, L. M., Costa, E., Harvey, I., and Coutinho, A., editors, *Proceedings of the 9th European Conference on Artificial Life (ECAL 2007)*, Lisbon, Portugal, volume 4648 of *Lecture Notes in Artificial Intelligence*, pages 895–904, Berlin / Heidelberg. Springer.
- Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2008a). A framework for the local information dynamics of distributed computation in complex systems. Submitted to *Physica D*.
- Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2008b). The information dynamics of phase transitions in random boolean networks. In *Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems (ALife XI)*, Winchester, UK. MIT Press. to appear.
- 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.
- Lungarella, M. and Sporns, O. (2006). Mapping information flow in sensorimotor networks. *PLoS Computational Biology*, 2(10):e144.
- MacKay, D. J. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, Cambridge.
- Mitchell, M., Crutchfield, J. P., and Hraber, P. T. (1994). Evolving cellular automata to perform computations: Mechanisms and impediments. *Physica D*, 75:361–391.
- 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.
- Pincus, S. and Singer, B. H. (1996). Randomness and degrees of irregularity. *Proceedings of the National Academy of Sciences*, 93(5):2083–2088.
- Polani, D., Sporns, O., and Lungarella, M. (2007). How information and embodiment shape intelligent information processing. In Lungarella, M., Iida, F., Bongard, J., and Pfeifer, R., editors, *Proceedings of the 50th Anniversary Summit of Artificial Intelligence*, New York, volume 4850 of *Lecture Notes in Computer Science*, pages 99–111, Berlin / Heidelberg. Springer.
- Prokopenko, M., Gerasimov, V., and Tanev, I. (2006a). 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., Gerasimov, V., and Tanev, I. (2006b). Measuring spatiotemporal coordination in a modular robotic system. In *Proceedings of the 10th International Conference on the Simulation and Synthesis of Living Systems (ALifeX)*, Bloomington, Indiana, USA. MIT Press.
- Schreiber, T. (2000). Measuring information transfer. *Physical Review Letters*, 85(2):461–464.
- Solé, R. V. and Valverde, S. (2001). Information transfer and phase transitions in a model of internet traffic. *Physica A*, 289(3-4):595–605.
- Sperati, V., Trianni, V., and Nolfi, S. (2007). Evolving coordinated group behaviors through maximization of mean mutual information. Submitted to *Swarm Intelligence*.
- Sporns, O. and Lungarella, M. (2006). Evolving coordinated behavior by maximizing information structure. In Rocha, L. M., Yaeger, L. S., Bedau, M. A., Floeano, D., Goldstone, R. L., and Vespignani, A., editors, *Proceedings of the Tenth International Conference on Simulation and Synthesis of Living Systems (ALifeX)*, Bloomington, Indiana, USA, pages 323–329. MIT Press.
- Tanev, I., Ray, T., and Buller, A. (2005). Automated evolutionary design, robustness, and adaptation of sidewinding locomotion of a simulated snake-like robot. *IEEE Transactions on Robotics*, 21(4):632–645.
- Yaeger, L. and Sporns, O. (2006). Evolution of neural structure and complexity in a computational ecology. In Rocha, L. M., Yaeger, L. S., Bedau, M. A., Floeano, D., Goldstone, R. L., and Vespignani, A., editors, *Proceedings of the Tenth International Conference on Simulation and Synthesis of Living Systems (ALifeX)*, Bloomington, Indiana, USA, pages 330–336. MIT Press.