Advances in Complex Systems

INFORMATION AND SELF-ORGANIZATION OF BEHAVIOR

DANIEL POLANI

Adaptive Systems Research Group, University of Hertfortshire
College Lane, Hatfield, AL10 9AB, United Kingdom
d.polani@herts.ac.uk

MIKHAIL PROKOPENKO*

CSIRO Information and Communication Technologies Centre, Adaptive Systems
PO Box 76, Epping, NSW 1710, Australia
mikhail.prokopenko@csiro.au

LARRY S. YAEGER

Informatics, Indiana University, 919 E. 10th St., Bloomington, IN 47408, United States
Google, 1600 Amphitheatre Parkway, Mountain View, CA 94303, United States
larryy@google.com

Received (received date)
Revised (revised date)

1. Introduction

The goal of Guided Self-Organization (GSO) is to leverage the strengths of self-organization while still being able to direct the outcome of the self-organizing process. GSO typically has the following features: (i) an increase in organization (structure and/or functionality) over some time; (ii) the local interactions are not explicitly guided by any external agent; and (iii) task-independent objectives are combined with task-dependent constraints. Over the last few years a mathematical framework has started to form around these features, promising to provide common organizational and guidance principles across multiple scales and contexts. This process is far from being complete, and every year an International GSO Workshop showcases new breakthroughs that diversify and reshape the field. Nevertheless, some themes and ideas withstand the test of time, maintaining the core of the GSO research.

One of these themes is the role of information (understood as Shannon information, i.e. "reduction in uncertainty") in guiding a self-organizing process. In particular, a lot of progress has been achieved in studying various aspects of information structure and information processing during self-organization of behavior

*Corresponding author

(molecular, neural, cognitive, social, etc.). For example, several principles based on information flows through the perception-action loops of embodied cognitive systems were recently developed [5]. These principles related GSO to the notion that adaptive behaviors emerge from interactions between brain, body, and environment while optimizing task-independent objective functions.

Having a language that describes interactions is essential for a non-reductionist science [14]. And so another common GSO trend is the use of graph theory in representing and analyzing interactions within a system, be it a cell, the brain, a social network, an ecological web, or a power grid. Several graph-theoretical measures have been devised and put to use in tracing various self-organization processes developing within networks, as well as in relating connectivity of the self-organizing systems to their function [29, 13].

This topical issue presents a selection of papers following two GSO Workshops (Bloomington, Indiana, USA, in 2010 and Hertfordshire, UK, in 2011). These papers are grouped into three sections. The first section contains three studies characterizing neural dynamics with model-free techniques. It is followed by a section on embodied (e.g., robotic) systems, consisting of five papers. These works investigate various control loops and goal-oriented behavior, shaped by specific organizational principles and constraints, such as information bottleneck, homeokinesis, empowerment, maximization of transfer entropy, etc. The concluding section broadens the scale of self-organization, extending to collective behavior of multi-agent systems. It also comprises five papers, dealing with bio-inspired algorithms, pattern formation, and social dilemmas, and offering several efficient mechanisms for guided self-organization.

2. Neural dynamics and model-free methods

The first section of the topical issue is devoted to studies of neural dynamics with model-free methods. Self-organization within a neural system is typically brought about by neuronal interactions, potentially resulting in neural plasticity and learning. Modelling and analyzing these interactions is a daunting task which remains a research challenge. For instance, Chicharro and Ledberg [10] studied the brain as a biological system consisting of multiple interacting components, showing that the influence of causal connections on the natural dynamics of the system often cannot be analyzed in terms of the causal effect of one subsystem on another.

One emerging trend bypassing this challenge suggests the use of model-free techniques that quantify directed interactions within the brain [37, 18, 36]. Some of these techniques take advantage of the generic nature of information-theoretic methods. For example, Wibral et al. [37] analyzed magnetoencephalography (MEG) source-level signals using transfer entropy, successfully detecting changes in cortical and subcortical networks between the different auditory task types.

The papers assembled in this section take the next step. The model-free techniques are not only used here to quantify the strength of interactions and/or detect specific pathways, but also form a basis for generic organizational principles that govern and guide the self-organization process at different scales.

The paper "Metabolic cost as an orginizing principle for cooperative learning" by Balduzzi et al. studies how neurons can use metabolic cost to facilitate learning at a population level. The investigation is focussed on self-organization as an internal process to the neural system, and so one may consider neurons as explicit communication channels mapping situations to actions (using spiketrains). In addition, the general inability of individual neurons to significantly manipulate their environment simplifies the optimization problems. Specifically, the paper shows that constraining reward maximization by metabolic cost forces neurons to maximize the information they encode into their spikes, aligning the information content of actions with their expected reward. The main conclusion is that "information aligns with rewards", and that metabolic cost may provide an organizing principle underlying the neural code. Moreover, one important ramification is that spikes with high information content are worth learning from, and so the suggested organizing principle may be extended to studies of neuronal learning, as well as to analysis and design of other cooperating populations.

The multi-scale relevance of generic organizational principles is a recurring theme in many papers of this topical issue. The paper by Yaeger "Identifying neural network topologies that foster dynamical complexity" links (in fact, correlates) neural complexity with behavioral complexity. The investigation is concerned with behavioral adaptation and an evolutionary selection for complexity, putting forward the question whether evolution selects for complex dynamics and specific network topologies, e.g. small-world networks. It has recently been demonstrated by Lizier et al. [19] that small-world networks are capable of maintaining and balancing comparably large information storage and information transfer. Not surprisingly, the paper quantifies the complexity of neural dynamics using information theory as well. Specifically, TSE complexity [32], is used, capturing the dual and opposing tensions towards global coordination and cooperation ("integration") and localized specialization of functionality ("segregation"). Furthermore, neural dynamics are related to several graph-theoretical measures calculated for the underlying network topologies. While the study is carried out within a simulated system of Polyworld, the resulting conclusion is far-reaching: "functional and structural evolutionary pressures cooperate to produce brains optimized for adaptation to a complex, variable world".

The paper "Combining Correlation-Based and Reward-Based Learning in Neural Control for Policy Improvement" by Manoonpong et al. concludes our first section dedicated to self-organizing neural systems. In this paper the guidance of the learning strategy is achieved by a combination of correlation-based learning and reinforcement learning. The first approach uses the correlations between external stimuli and anticipatory actions, and typically learns simple tasks quickly. The second approach uses predefined rewards / punishments as evaluations allowing an agent to optimize its expected future rewards, and in general is slower but more successful in solving complicated tasks. The combined learning model introduced here is shown to strongly improve performance of the controller. Interestingly, the learning rules do not require an explicit model of the environment, and so the presented approach is still a model-free method, in line with other techniques for guided self-organization described in this topical issue.

3. Embodiment, Robotics, and Control

Neural systems are distinguished by the fact that they operate in an environment where their counterparts communicate (and aim to communicate) with them: in other words, the self-organization in a neural system is "in the interest" of every participating entity (neurons) which cooperate in making the communication codes "jointly" understood. As one proceeds to agents operating in an external environment, this ceases to be true. In fact, the environment is not an agent and has no "intention" to communicate with the agents living in it [**?**]]gibson-ecological.

Even so, the environment influences the agent and the agent, in turn, influences the environment and thus the state of the agent and the state of the environment will be correlated. In the biological realm, evolutionary pressures will push the biological agents to modify the environment and their interaction with it to a state of selective advantage of the agents. This means that the structure of the environment will be reflected in how the agent interacts with it, and one can expect the agent's cognitive control to encompass some characteristics of the world it inhabits. The "code word" *embodiment* has been used for a while to describe how the agent's internal decisions translate into actual behavior in the world; and while in the past, rule-based and symbolic views have been adopted in which a platonic cognitive unit "hovers over the robotic reality" to select optimal behaviors [38, 21], it has become clear that this approach is doomed to fail, as it has to take into account all potential cases and exceptions, which is unrealistic except in the simplest of worlds.

In response to this, Brooks [9] proposed an inherently opposed approach which considers a non-symbolic, non-representational cognitive approach to control problems. In fact, such approaches had been explored in early cybernetics [2, 3, 34] then all but forgotten in the wake of the quick growth of the successes of the von Neumann architecture. Yet, it was well understood that even extremely simple principles could produce a variety of complex and "sensible" behaviors [7]. For the extraordinary success of these simple cognitive architectures, a core importance was attributed to embodiment, i.e. to how an agent is embedded in its physical environment. Paul [23], Pfeifer and Bongard [24], and Hoffmann and Pfeifer [15] utilized the term *morphological computation* to emphasize that the embedding into the environment participates in the cognitive process. In nature this embedding will be the product of evolutionary adaptation; in an artificial system it may be the result of a suitable design. But it makes clear that the structure of the embodiment plays an important role in producing useful behaviors. In fact, for simple examples, seemingly trivial properties of the embodiment can contribute to drastically different cognitive load on the agent itself. For example, in a simple gridworld, randomly relabeling actions individually per state instead of keeping globally consistent direction labels (i.e. global north, east, south, west labels) makes policies of equivalent performance cognitively much more expensive (measured using information theory) [26]. In terms of traditional, "platonic" Artificial Intelligence, both scenarios are completely equivalent and can be transformed into each other by a simple permutation of the actions. However, the embodiment perspective makes clear that this relabeling of actions, while in principle abstractly possible, is costly. The success of agents embodied in a real world depends on this embodiment being "naturally" supportive of the agent's task.

In typical studies, the embodiment is either a product of natural evolution (where it is taken for granted in advance) or the result of an engineering process (where it is designed into the agent on purpose). However, its acquisition and direct role in shaping cognition is comparatively rarely the direct object of study. It is to fill this gap that many of the papers of this section contribute.

The paper "Robustness of guided self-organization against sensorimotor disruptions" by Georg Martius takes one of the pioneer approaches to self-organized behavior generation for agents and robots — *homeokinesis* — to a new level. Homeokinesis was introduced by Der et al. [12, 11] as an intrinsic motivation model which operates on the assumption that an agent will try to adopt behaviors which will make its future as predictable for it as possible, while at the same time allowing for a rich set of future options. An information-theoretic generalization of the approach has used predictive information [6, 4] to the same effect. Homeokinesis is a purely self-organized approach, but has seen extension towards incorporating goal-directed behaviors [20] by incorporating preferences for preferred modes of behavior of the self-organizing system.

The new paper by Martius now extends the existing homeokinetic approach by taking into account not just proximity sensors (such as contact sensors), but now also distal sensors, namely in the form of vision sensors. The basic coordination and task (ball pushing), including vision sensors, are learnt in the self-organizing homeokinesis process in a short time. Their work now studies the advantages of the guided self-organizing system in the given configuration with respect to a purely self-organized or purely guided system. The combined system is able to handle massive sensoric reconfigurations rapidly and with good reliability. Turning off the self-organization part while keeping the "guided" component active, an agent loses its ability to explore and is less flexible when it comes to extricating itself from difficult situations, such as escaping from corners. Pure self-organization without control, on the other hand, tends to lose focus and sight of the goal behavior. The combination of guidance and self-organization provides a prime example for the strengths of each of the contributions and the power that the combined approach can offer.

The paper "A Goal-Orientation Framework for Self-Organizing Control" by Hesse and Wörgötter consider again homeokinesis as a generator for low level controls. However, their approach consists of separating and layering the low-level self-organizational behaviors and the high-level goal-directed behaviors. Thus the high-level behaviors make use of the adapted homeokinetic low-level behaviors for their optimization. The resulting high-level behaviors will thus not be perfectly optimal, as the robustness of the low-level behaviors comes at a price. However, the combination provides a significant increase in flexibility. The high-level controller has the choice of presetting the scenarios (such as a fixed hand opening which is externally imposed at the beginning of a ball-gripping experiment), or defining the basic movement patterns. The paper emphasizes a strict separation of goals and self-organization for the purpose of clarifying which level is responsible for which part of the behavior. This also provides an interface for potential extensions of one level independently of the other.

The relation between goal-directed learning and the sometimes surprising success of self-organizing behavior generation has been quite in the dark. While traditional AI considers mostly goal-directed learning, much of its "intelligent" achievements can already be addressed by self-organizing approaches. The paper "Informational Constraints-Driven Organization in Goal-Directed Behavior" by van Dijk and Polani takes some steps to understand why that could be the case. It considers an informational picture of cognitive processing in which constraints on Shannon information determine what preferred behaviors will look like in a given environment. For this purpose, it is assumed that agents have a limited goal-related working memory. Transition points in the task (such as doors) will lead to a "re-caching" of goal information; these transition points are prime candidates to be used as subgoals, but emerge entirely from the necessity to limit working memory intake. The goal-information perspective also provides an approach to producing a natural classification of the environment into different rooms, without having an explicit concept of "rooms". Another variant produced clusters of locations in the world which would be neighbouring if not for the walls between them — a "spillover" map that ignores the incidental local walls and rather respects the global neighbourhoods of the world. These outcomes indicate that there is already significant prior structure in the world which well-designed self-organizing methods can successfully exploit to great effect.

A special case of environment structure is the dynamics of the perception-action loop of an agent. Many approaches, including homeokinesis, use this loop. Another intrinsic motivation model is empowerment, which considers the channel capacity of the external perception-action loop; i.e., the potential perceivable change that an agent can reliably effect on its environment, measured in terms of information theory. Empowerment has been studied in a number of scenarios [16], but its computation beyond small discrete cases is computationally expensive. While for discrete worlds useful approximations and extensions have been developed [1], the extension to the continuous domain has remained difficult. The paper "Approximation of Empowerment in the Continuous Domain" by Salge et al. introduces a quasi-linear Gaussian framework which takes advantage of the fact that often perception-action dynamics can be considered locally Gaussian. The authors apply the framework to an inverse pendulum problem where, although the approximated empowerment landscape differs from the more accurate one, the induced dynamics is very similar. Thus, the quasilinear Gaussian approximation promises to be a practically viable approach to compute empowerment in an efficient way.

"Bootstrapping Perception Using Information Theory: Case Studies in a Quadruped Robot Running on Different Grounds" by Schmidt et al. explores the sensorimotor structure of an embodied hardware robot which exhibits a number of self-organized gaits driven by the interaction with the ground, the "Puppy" robot. The authors use transfer entropy-based measures to characterize the sensoric and motoric variables of the Puppy robot. The transfer entropy is being used to capture the temporal (however, not causal) and interconnectedness of the various sensors and actuators in a time-relevant manner. The analysis is able to identify "easier" and "harder" ways to accomplish certain desired behaviors by quantifying how certain desired aspects of the movement (acceleration or turning) are affected by other dynamics of the system. Also externalities, such as a change of the ground and its properties, can be identified. Importantly, only the agent's own sensors are used, and thus all results of the analysis can be directly considered as knowledge available to the agent. In a similar vein to Olsson et al. [22], sensoritopic maps are constructed, this time, however, based on the transfer-entropy characteristics of the system. The information-theoretic analysis reveals a number of salient points: not only are many aspects of the robot dynamics reflected in the informational signature of the system, but it also shows a blurred boundary between proprio- and exteroreceptive sensors. For instance, whereas in traditional views, the knee angular receptors would be classified as pure proprioreceptive sensors, in the present work the dependence of their status on externalities (such as orientation and ground properties), as revealed by the informational analysis, places the classification of this sensor into a gray zone between proprio- and exteroreceptive sensors. Thus, the informational analysis provides a deeper insight into the dynamics of the agent.

4. Bio-inspired collective systems

As we increase the scale of self-organizing systems, we turn our attention to collective phenomena brought about by interactions of individual components of the systems. Again, the theme of information often takes the central stage, underlying different phases of dynamics observed during self-organizing processes, and suggesting generic organizational principles behind complex collective behavior. It is worth pointing out that the role of information-theoretic characterizations is not just in providing a common unifying language for an analysis of various systems across scales, but also in pinpointing and/or predicting specific features of collective dynamics. For instance, collective memory and information cascades in swarms were recently quantified by Wang et al. [35] via information theoretic measures, relating these features to self-organization within the swarm. Another example is given by the study of the collective behaviour and its dynamics in the context of self-organized path formation in a swarm of robots [31], which introduced and verified a measure based on information entropy.

Often self-organizing phenomena are investigated with network models, and in this case one may need to inter-relate information-theoretic and graph-theoretic analyses. For example, the prevalence of particular motifs (directed feedback and feedforward loops) in biological and artificial networks was recently explained via their role in contributing to local information storage, measured information-theoretically and correlated with clustering coefficient(s) [17]. This inter-relationship between graph theory and information theory goes beyond a purely methodological gain: as noted by Rosvall and Bergstrom [28], "if we want to understand how network structure relates to system behavior, we need to understand the flow of information on the network".

Once generic organizational principles, supported by reliable measures, are in place, one may contemplate the task of designing or synthesizing a desirable collective behavior, while balancing the expected degrees of robustness, flexibility and scalability. This task may, for example, be carried out with evolutionary robotics techniques [33].

The studies presented in this section are also strongly motivated by their biological analogies, where collective communications and information processing give rise to intricate patterns and functions at a global level.

The paper "Self-organizing particle systems" by Harder and Polani attempts to discover general information-processing principles underlying one of the most entangled and perplexing processes: self-organization of cells into a living organism. In doing so, the study considers a system of interacting particles that roughly mimic biological cells by exhibiting differential adhesion behavior. Curiously, it shows that particles can self-organize without the emergence of pattern-like structures, highlighting once more the difference between self-organization and emergence [30, 27]. Nevertheless, the main conclusion is that regular pattern-like structures do help to overcome limitations of self-organization that are imposed by spatial structure of interactions. In this paper, self-organization is measured via multi-information which is capable of detecting long-range correlations among interacting particles, and distinguish between qualitatively different phases of the self-organization process.

The paper "Bio-development of motorway networks in the Netherlands: A slime mould approach" by Adamatzky et al. stands out in its bold attempt to utilize the foraging behavior of a real biological organism, *Physarum polycephalum*, in constructing networks that resemble dense motorway networks in Europe. This study not only provides a striking example of similarities between vastly different scales (bio-development is carried out by a plasmodium of *acellular* slime mould, while transport networks emerge as a result of diverse social, economic, political and other factors), it also highlights the use of generic graph-theoretic measures, such as Harare index and Randić index, in the analysis of complex self-organizing topologies, e.g. relative neighborhood graphs, β-skeletons, minimum spanning trees, etc. Using these methods the authors are able to demonstrate the discovery of different graph classes by the slime mould growth dynamics.

This theme is continued in the paper "An Ant-Based Algorithm with Local Optimization for Community Detection in Large-Scale Networks" by He et al. Community structure appears when some network nodes cluster into groups with a high density of intra-cluster edges and a relatively lower density of inter-cluster edges. The method proposed to detect such structures is based on local optimization of modularity using ant-based algorithms. In general, the search for the partition with maximal modularity has been shown to be an NP-complete problem [8]. The bio-inspired algorithm described in this paper has a low computational complexity but still identifies high-quality community division. Interestingly, as network heterogeneity is increased, the method needs adjustments in order to keep up with performance of Infomap, a method that traces information flows in discovering a modular organization [28]. Infomap identifies a well connected module as a group of nodes among which information flows quickly and easily. The adjustment discussed in the paper (a weighting scheme that better discriminates between inter-cluster and intra-cluster edges) highlights an important difference between the specific topological approximation and the generic organizational principle underlying Infomap (that is, modularity is maximized when information flow is maximized). The more comprehensive performance of Infomap for heterogeneous networks may, of course, be traded off against computational complexity of distributed algorithms.

Kitto and Boschetti consider even a higher scale of collective phenomena in their paper "Attitudes, Ideologies and Self-Organization: Information Load Minimization in Multi-Agent Decision Making": social behavior and formation of ideologies. They propose a model of human decision making in which attitudes of individual agents self-organize into ideologies. Moreover, the formed ideologies further guide agent-based attitude changes. The paper reports a number of interesting observations, including a tendency to minimize the entropy in the system that quantifies and aggregates agents' levels of cognitive dissonance. The study argues that a considerable re-organization in the system can be attained by a single external act of ideological guidance: "rather than having to convince several tens of agents to shift their local framing of an issue, all that was needed was to change the ideology that they subscribe to". Importantly, this action requires a smaller information processing effort whenever the system is organized. Thus, we again see a dependency between levels of a system's (self-)organization and the extent of its information dynamics.

This topical issue is concluded with the paper "Preferential opponent selection in public goods games" by Brede, in which the subject of social behavior is revisited. The paper begins with the definition: "Altruism describes individual behaviour that benefits the group, but comes at a cost to the individual", and proceeds to rigorously model the emergence of cooperation in evolutionary social dilemmas. The main result is the observation that a preference to play with successful opponents strongly enhances the prevalence of cooperation. This result is complemented with a mechanism for the emergence of opponent selection biases. As with many other papers in this section, the analysis is carried out on networks, revealing some thought-provoking effects. For example, network heterogeneity is shown to boost cooperation. This effect "results from the superior ability of hub nodes to generate payoff and spread their strategy to adjacent nodes", while "a defector hub would undermine its position by surrounding itself by defectors". Assortative mixing of network nodes (i.e., their preference to connect with similar nodes) has been previously shown to be strongly related to information content [25], and so the effect of network heterogeneity reported by Brede can also be interpreted as an effect of information structure on the mechanism of (social) interactions.

5. Conclusion

This topical issue presents a range of studies, at increasing scales: starting with self-organization of neural dynamics within the brain, moving to generation of coordinated and goal-oriented behavior of embodied agents, and culminating in collective, multi-agent, behavior inspired by distributed biological and social systems. Almost all these works were unified by the use of model-free and task-independent objectives, motivated and supported by organizational / optimization principles. The principles in turn are unified by a common language provided by information theory and graph theory. In particular, it is interesting to observe how information dynamics were utilized to capture important dependencies and processes in seemingly unrelated areas: from neuronal interactions to cellular aggregation to perception-action loops to formation of ideologies. In addition, it was revealing to see a significant overlap in topological features that were shown to be critical for self-organization of behavior in small-world, modular, heterogeneous and other types of networks.

The emergence of unifying generic principles is the main goal and achievement of GSO research over the last several years. We hope that the papers presented in this topical issue will inspire further progress in this field.

6. Acknowledgments

We would especially like to thank the anonymous reviewers for this topical issue for their timely responses and useful comments. We also appreciate the professional effort of all contributing authors. Finally, we are grateful for the support provided by the organizers of GSO-2010 at Indiana University and GSO-2011 at the University of Hertfordshire and CSIRO Complex Systems Science team. GSO 2011 was partly supported by the European Commission as part of the CORBYS (Cognitive Control Framework for Robotic Systems) project under contract FP7 ICT-270219. The views expressed in this Topical Issue are those of its editors and authors, and not necessarily those of the consortium.

References

[1] Anthony, T., Polani, D., and Nehaniv, C., Impoverished empowerment: 'meaningful' action sequence generation through bandwidth limitation, in *Advances in Artificial Life. Darwin Meets von Neumann - 10th European Conference, ECAL 2009, Budapest, Hungary, September 13-16, 2009, Revised Selected Papers, Part II*, eds. Kampis, G., Karsai, I., and Szathmáry, E., *Lecture Notes in Computer Science*, Vol. 5778 (Springer, 2009), pp. 294–301.

[2] Ashby, W. R., Principles of the self-organizing dynamic system, *J. Gen. Psychol.* **37** (1947) 125–128.

[3] Ashby, W. R., *Design for a Brain* (Wiley & Sons, New York, 1952).

[4] Ay, N., Bertschinger, N., Der, R., Güttler, F., and Olbrich, E., Predictive information and explorative behavior of autonomous robots, *European Journal of Physics B* **63** (2008) 329–339.

[5] Ay, N., Der, R., and Prokopenko, M., Guided self-organization: perception-action loops of embodied systems, *Theory in Biosciences* **131** (2012) 125–127.

[6] Bialek, W., Nemenman, I., and Tishby, N., Predictability, complexity and learning, *Neural Computation* **13** (2001) 2409–2463.

[7] Braitenberg, V., *Vehicles: Experiments in Synthetic Psychology* (MIT Press, Cambridge, 1984).

[8] Brandes, U., Delling, D., Gaertler, M., Goerke, R., Hoefer, M., Nikoloski, Z., and Wagner, D., Maximizing modularity is hard (2006).

[9] Brooks, R. A., Intelligence without representation, *Artificial Intelligence* **47** (1991) 139–159.

[10] Chicharro, D. and Ledberg, A., When two become one: The limits of causality analysis of brain dynamics, *PLoS ONE* **7** (2012) e32466.

[11] Der, R., Self-organized acqusition of situated behavior, *Theory Biosci.* **120** (2001) 1–9.

[12] Der, R., Steinmetz, U., and Pasemann, F., Homeokinesis – a new principle to back up evolution with learning, in *Computational Intelligence for Modelling, Control, and Automation*, ed. Mohammadian, M., *Concurrent Systems Engineering Series*, Vol. 55 (IOS Press, 1999), pp. 43–47.

[13] Gershenson, C., Guiding the self-organization of random boolean networks, *Theory in Biosciences* **131** (2012) 181–191.

[14] Gershenson, C. and Heylighen, F., How can we think the complex?, in *Managing Organizational Complexity: Philosophy, Theory and Application*, ed. Richardson, K. (Information Age Publishing, 2005), pp. 47–61.

[15] Hoffmann, M. and Pfeifer, R., The implications of embodiment for behavior and cognition: animal and robotic case studies, in *'The Implications of Embodiment: Cognition and Communication'*, eds. Tschacher, W. and Bergomi, C. (Exeter: Imprint Academic, 2011), to appear.

[16] Klyubin, A. S., Polani, D., and Nehaniv, C. L., Keep your options open: An information-based driving principle for sensorimotor systems, *PLoS ONE* **3** (2008) e4018.

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

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

[19] Lizier, J. T., Pritam, S., and Prokopenko, M., Information dynamics in small-world

Boolean networks, *Artificial Life* **17** (2011) 293–314.
[20] Martius, G. and Herrmann, J. M., Variants of guided self-organization for robot control, *Theory in Biosciences* **131** (2012) 129–137.
[21] Nilsson, N. J., Shakey the robot (1984).
[22] Olsson, L., Nehaniv, C. L., and Polani, D., From unknown sensors and actuators to actions grounded in sensorimotor perceptions, *Connection Science* **18** (2006) 121–144.
[23] Paul, C., Morphological computation: A basis for the analysis of morphology and control requirements, *Robotics and Autonomous Systems* **54** (2006) 619–630.
[24] Pfeifer, R. and Bongard, J., *How the Body Shapes the Way We think: A New View of Intelligence* (Bradford Books, 2007).
[25] Piraveenan, M., Prokopenko, M., and Zomaya, A., Assortativeness and information in scale-free networks, *The European Physical Journal B* **67** (2009) 291–300.
[26] Polani, D., An informational perspective on how the embodiment can relieve cognitive burden, in *Proc. IEEE Symposium Series in Computational Intelligence 2011 — Symposium on Artificial Life* (IEEE, 2011), pp. 78–85.
[27] Prokopenko, M., Boschetti, F., and Ryan, A. J., An information-theoretic primer on complexity, self-organization, and emergence, *Complexity* **15** (2009) 11–28.
[28] Rosvall, M. and Bergstrom, C. T., Maps of random walks on complex networks reveal community structure, *Proceedings of the National Academy of Sciences* **105** (2008) 1118–1123.
[29] Rubinov, M. and Sporns, O., Complex network measures of brain connectivity: Uses and interpretations, *NeuroImage* **52** (2010) 1059–1069.
[30] Shalizi, C. R., *Causal architecture, complexity and self-organization in time series and cellular automata*, Ph.D. thesis, The University of Wisconsin – Madison (2001).
[31] Sperati, V., Trianni, V., and Nolfi, S., Self-organised path formation in a swarm of robots, *Swarm Intelligence* **5** (2011) 97–119.
[32] Tononi, G., Sporns, O., and Edelman, G. M., A measure for brain complexity: relating functional segregation and integration in the nervous system, *Proceedings of the National Academy of Sciences of the United States of America* **91** (1994) 5033–5037.
[33] Trianni, V. and Nolfi, S., Engineering the evolution of self-organizing behaviors in swarm robotics: A case study, *Artificial Life* **17** (2011) 183–202.
[34] Walter, W. G., A machine that learns, *Scientific American* **185** (1951) 60–63.
[35] Wang, X., Miller, J., Lizier, J., Prokopenko, M., and Rossi, L., Quantifying and tracing information cascades in swarms., *PLoS One* **7** (2012) e40084.
[36] Wibral, M., Pampu, N., Priesemann, V., Siebenhhner, F., Seiwert, H., Lindner, M., Lizier, J. T., and Vicente, R., Measuring information-transfer delays, *PLoS ONE* **8** (2013) e55809.
[37] Wibral, M., Rahm, B., Rieder, M., Lindner, M., Vicente, R., and Kaiser, J., Transfer entropy in magnetoencephalographic data: quantifying information flow in cortical and cerebellar networks., *Progress in Biophysics and Molecular Biology* **105** (2011) 80–97.
[38] Winograd, T., Procedures as a representation for data in a computer program for understanding natural language, AI Technical Report 235, MIT (1971).