Dynamical Properties of Artificially Evolved Boolean Network Robots

Andrea Roli[1(✉)], Marco Villani[2], Roberto Serra[2], Stefano Benedettini[1], Carlo Pinciroli[3], and Mauro Birattari[4]

[1] Department of Computer Science and Engineering, Alma Mater Studiorum Università di Bologna, Bologna, Italy
andrea.roli@unibo.it
[2] Department of Physics, Informatics and Mathematics, Università di Modena e Reggio Emilia & European Centre for Living Technology, Venice, Italy
[3] MIST, École Polytechnique de Montreal, Montreal, Canada
[4] IRIDIA-CoDE, Université libre de Bruxelles, Brussel, Belgium

Abstract. In this work we investigate the dynamical properties of the Boolean networks (BN) that control a robot performing a composite task. Initially, the robot must perform phototaxis, i.e. move towards a light source located in the environment; upon perceiving a sharp sound, the robot must switch to antiphototaxis, i.e. move away from the light source. The network controlling the robot is subject to an adaptive walk and the process is subdivided in two sequential phases: in the first phase, the learning feedback is an evaluation of the robot's performance in achieving only phototaxis; in the second phase, the learning feedback is composed of a performance measure accounting for both phototaxis and antiphototaxis. In this way, it is possible to study the properties of the evolution of the robot when its behaviour is adapted to a new operational requirement. We analyse the trajectories followed by the BNs in the state space and find that the best performing BNs (i.e. those able to maintaining the previous learned behaviour while adapting to the new task) are characterised by generalisation capabilities and the emergence of simple behaviours that are dynamically combined to attain the global task. In addition, we also observe a further remarkable property: the complexity of the best performing BNs increases during evolution. This result may provide useful indications for improving the automatic design of robot controllers and it may also help shed light on the relation and interplay among robustness, evolvability and complexity in evolving systems.

1 Introduction

Genetic regulatory networks (GRNs) model the interaction and dynamics among genes. From an engineering and computer science perspective, GRNs are extremely interesting because they are capable of producing complex behaviours, notwithstanding the compactness of their description. Cellular systems are also

M. Gavanelli et al. (Eds.): AI*IA 2015, LNAI 9336, pp. 45–57, 2015.
DOI: 10.1007/978-3-319-24309-2_4

both robust and adaptive, i.e. they can maintain their basic functions in spite of damages and noise, and they are able to adapt to new environmental conditions. Such a complex behaviour can be interpreted from an artificial system design's viewpoint, suggesting the possibility of achieving robust and adaptive behaviours in agents, robots, and group of robots, by exploiting the properties of GRN models.

Among the most studied models for GRNs, are Boolean networks (BNs), first introduced by Kauffman [11]. A BN is a discrete-time discrete-state dynamical system whose state is a N-tuple in $\{0, 1\}^N$, $(x_1, \ldots, x_N)$. The state is updated according to the composition of N Boolean functions $f_i(x_{i_1}, \ldots, x_{i_{K_i}})$, where K_i is the number of inputs of node i, which is associated to Boolean variable x_i. Each function f_i governs the update of variable x_i and depends upon the values of variables $x_{i_1}, \ldots, x_{i_{K_i}}$. Most works on BNs deal with so-called *autonomous* networks, i.e. systems that evolve in time without input from the external—at most, only the initial state may be extrenally imposed. Usually, BNs are subject to a deterministic, synchronous and parallel node update, even if other update schemes are possible [28]. In the synchronous and deterministic update scheme, every state has a unique successor and the trajectory is composed of a transient and a state cycle (possibly a fixed point, i.e. a cycle of length 1).

BNs have received considerable attention in the community of complex system science. Works in complex systems biology show that BNs provide powerful model for cellular dynamics [26,29], cellular differentiation [25,31] and interactions among cells and environment [24]. A specific dynamical regime at the boundaries between order and chaos, called the *critical* regime, is of notable interest. Critical networks enjoy important properties, such as the capability of optimally balancing evolvability and robustness [1] and maximising the average mutual information among nodes [21]. Hence the conjecture that living cells, and living systems in general, are critical [17].

In recent works, it has been shown that such kind of BNs can be used to control robots [6,22,23]. In this case, the BN evolution in time also depends on the values of some "input" nodes which are set depending on the robot's sensor readings. The BN is trained by means of a learning algorithm that manipulates the Boolean functions. The algorithm employs as learning feedback a measure of the performance of the BN-controlled robot (in the following, BN-robot) on the task to perform. The effectiveness of this approach was demonstrated through experiments on both simulated and real robots.

In this contribution, we outline some results on the analysis of the BN-robot's dynamics along the learning process. We analyse the trajectories followed by the BN-robot in the space of BN states and compute significant features, such as state number and frequency of state occurrence in sample trajectories. In addition, we compute the number of *fixed points*, i.e. BN states repeated as long as the BN inputs do not change. The number of fixed points is an indicator of the generalisation capabilities of the system, as they represent simple functional building blocks of the type `while <condition> do <action>`, which compose the overall system dynamics. Moreover, we estimate the statistical complexity of the system by means of a complexity measure called the *LMC complexity* [15]. The dynamics of a complex system is neither totally disordered (as an ideal gas at equilibrium), nor perfectly ordered (as a crystal); therefore we expect that a measure of the distance of a system from these two conditions should have very high values when the system exhibits complex behaviours. While we are of course aware of the fact that there is no general agreement on an all-encompassing definition of a measure of complexity, LMC seems particularly interesting in this case, as it will be discussed in Section 3.

We found that the successful performing BN-robots, which show the capability of attaining the learned behaviours also in spite of noise and perturbations (*robustness*) while adapting to new tasks to perform (*evolvability*), are characterised by both number of fixed points and LMC complexity higher than those of unsuccessful ones. These preliminary results may provide useful indications for improving the automatic design of robot controllers and may help shed light on the relation and interplay among robustness, evolvability and complexity in evolving systems.

The structure of the paper is as follows. After a summary of the experimental setting in Section 2, we discuss the main results of the analysis of the dynamics of the BNs controlling the robot in Section 3 and we conclude with a discussion and an outlook to future work in Section 4.

2 Experimental Setting

In this experiment, we control an *e-puck* robot [16] by means of a BN. The values of a set of network nodes (BN input nodes) are imposed by the robot's sensor readings, and the values of another set of nodes (BN output nodes) are observed and used to encode the signals for maneuvering the robot's actuators. The BN controlling the robot is subject to synchronous and parallel update. As described in the following, the Boolean functions are set by a search process, whilst the topology of networks is set at random.[1] The sensors consists of four light sensors and one sound sensor, while the actuators correspond to right and left wheel speed controllers. The Boolean values of the output nodes are sent to wheel actuators after a preprocessing consisting in a moving average, so as to feed the motors with signals in the range [0,1]. The robot is placed in a random position and with random orientation in a squared arena, with one light source in a corner. The BN-robot must accomplish the following task: initially, it must perform phototaxis, that is, move towards the light source; upon perceiving a sharp sound, the BN-robot must switch to antiphototaxis, that is, move away from the light source.[2] The robot is trained in simulation by means of an *adaptive walk*: the process starts from a randomly generated BN, it iteratively mutates its functions and keeps only the changes that either improve the BN-robot's performance or do not decrease it. Mutation is implemented by

[1] The choice for a random topology is not a limitation, as discussed in [22].

[2] A video of a typical run of a best performing BN-robot is available at https://www.youtube.com/watch?v=6ZF9Ijpwkd8.

randomly choosing a node and an entry in its Boolean function truth table and flipping it. The algorithm is therefore a stochastic descent in the space of Boolean functions. [3] Advanced search strategies can of course be devised so as to attain a higher performance; nevertheless, this subject is beyond the scope of this paper.

The BN-robot is trained in two sequential phases. In the first phase, the learning feedback is an evaluation of the robot's performance in achieving only phototaxis. In the second phase, the learning feedback is composed of a performance measure accounting for both phototaxis and antiphototaxis. In this way, we can study the properties of the evolution of the BN-robot when its behaviour has to be adapted to a new operational requirement. We define the performance of a BN-robot as a function of an error $E \in [0, 1]$. The smaller is the error, the better is the robot performance. The error function is given by a weighted sum of phototaxis and antiphototaxis errors: at each time step $t \in \{1, \ldots, T\}$, the robot is rewarded if it is moving in the correct direction with respect to the light. Let t_c be the time instant at which the clap is performed. The error function E is defined as follows:

$$E = \alpha\Big(1 - \frac{\sum_{i=1}^{t_c} s_i}{t_c}\Big) + \Big(1 - \alpha\Big)\Big(1 - \frac{\sum_{i=t_c+1}^{T} s_i}{T - t_c}\Big),$$

where:

$$\forall i \in \{1, \ldots, t_c\},\ s_i = \begin{cases} 1 & \text{if the robot goes towards to the light at step } i \\ 0 & \text{otherwise} \end{cases}$$

$$\forall i \in \{t_c{+}1, \ldots, T\},\ s_i = \begin{cases} 1 & \text{if the robot moves away from the light at step } i \\ 0 & \text{otherwise} \end{cases}$$

In the first phase of the training is $\alpha = 1$, whilst in the second phase α is set to 0.5 so as to take into account both phototaxis and antiphototaxis.

One hundred independent runs of the entire training process were executed,[4] starting from 100 initial BNs generated at random (with 20 nodes, 3 inputs per node and no self-connections).

During the training process BN-robots are subject to random perturbations, so as to train them also for operating in noisy environments. Along the training process we tested the BN-robot and collected statistics on the BN states traversed.

The experiments in simulation have been run by means of the open source simulator ARGoS [19].

3 Analysis of BN Dynamics

A significant fraction of the training experiments—about 30%—leads to a successful BN-robot, i.e. a robot able to perform both phototaxis and antiphototaxis

[3] Details can be found in [22].

[4] The experiments reported in [22] were re-run so as to have a greater number of replicas.

and to switch between the first and the second behaviour when it perceives a sharp sound signal. The unsuccessful BN-robots are either able to perform phototaxis only or not even that task. In the successful cases, the phototaxis capability acquired by the BN-robot in the first training phase is maintained while also the antiphototaxis behaviour is learned. Whence these systems have the possibility of successfully balancing robustness and evolvability.[5]

A question may rise at to what extent topology affects the results: after visual inspection of a sample of the BNs we discover that topology has an impact only in pathological cases, such as complete disconnection of all sensors or actuators. Notably, one of the best performing BNs has a topology in which the South light sensor is disconnected, which means that the network was anyway able to integrate this piece of information.

We studied the properties of the BN trajectories as they control the robot during its actions. The BN that controls a robot is coupled with the environment, as some of its nodes are forced to values imposed by the sensors and some of its outputs control the robot actuators (the wheels in this case). As a consequence, the network itself is *embodied* and its dynamics must be studied in the operational setting in which the robot is functioning, characterised by a specific sensors–actuators loop mediated by the environment. Therefore, we studied the dynamics of such BNs by means of the properties of their trajectories in the state space collected during robot runs. More precisely, for each BN-robot we run the robot starting from 1000 different initial conditions and recorded the sequence of BN states the network traverses during the run. This collection of state sequences is then merged into a graph whose vertices are the network states traversed by the BN and the edges the observed transitions between two states (see a typical trajectory graph in Figure 1). Moreover, the frequency of occurrence of states in the trajectories is recorded. This information is used to compute several features of the BN dynamics, which will be detailed in the following. The statistics that will be shown are computed by subdividing the BN-robots into three classes: BN-robots able to attain correctly the task (*both* class, about 30/100), BN-robots able to perform phototaxis only (*pt* class, about 60/100) and totally failing robots (*none* class, about 10/100).

3.1 Number of States

The number of unique states in the collection of trajectories—i.e. the number of different states in the set collecting all the states in the sampled trajectories—is an indicator of the size of the state space the BN dynamics occupies, as it represents the portion of state space actually explored by the BN. The smaller this size, the greater the generalisation capability of the network. Indeed, a large number of unique states denotes BN trajectories that do not overlap, which in turn means that the network has simply learned collections of successful examples. Conversely, a small number of unique states denotes trajectories that share

[5] We use the terms robustness and evolvability with the same meaning as in the work by Aldana et al. [1]

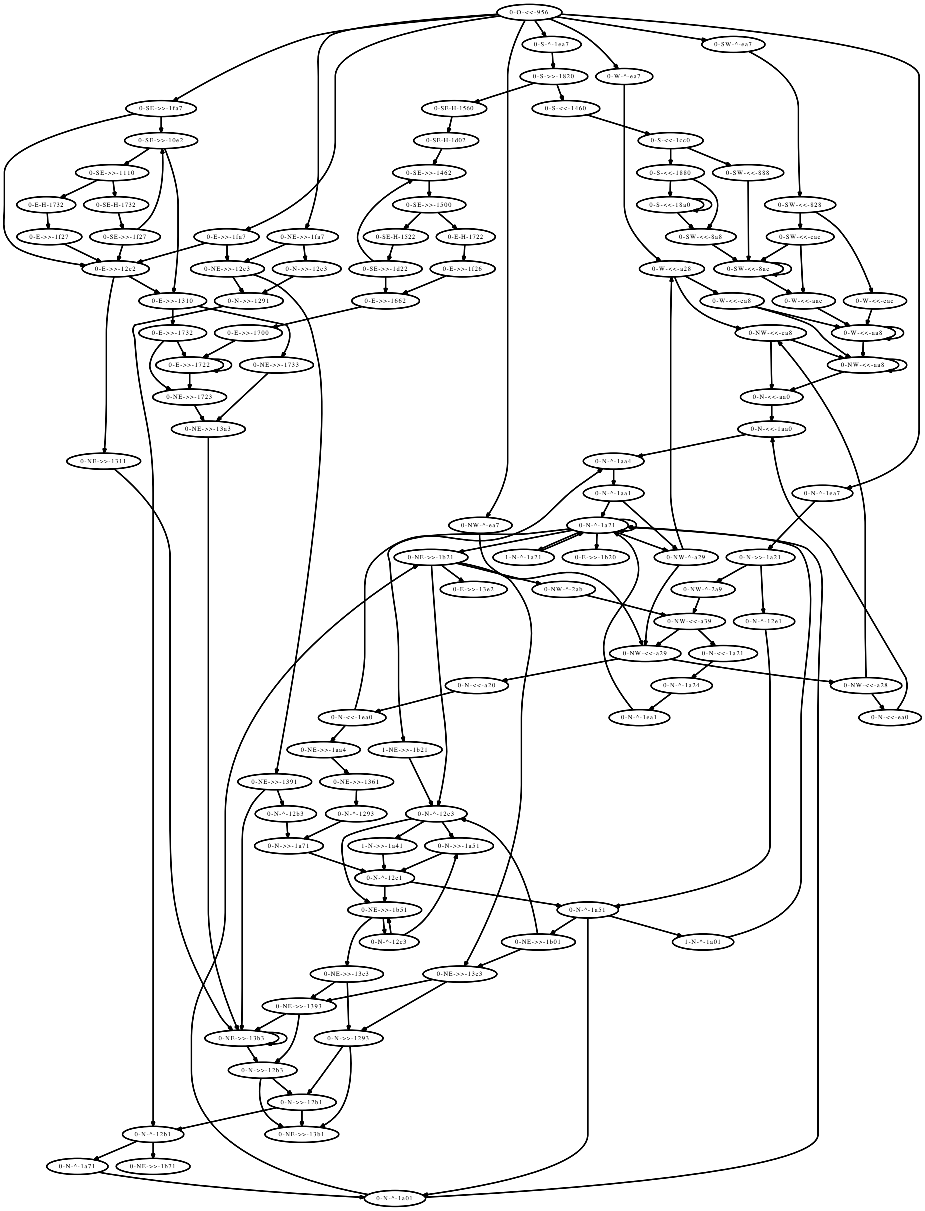

Fig. 1. Typical trajectory graph of a BN-robot. Transitions between nodes occur either for internal network state update or caused by input change. Node labels—not relevant for this context—denote the encoded binary state of the network.

a large fraction of transitions, which is a property of a system that was able to generate a compact model of the world.

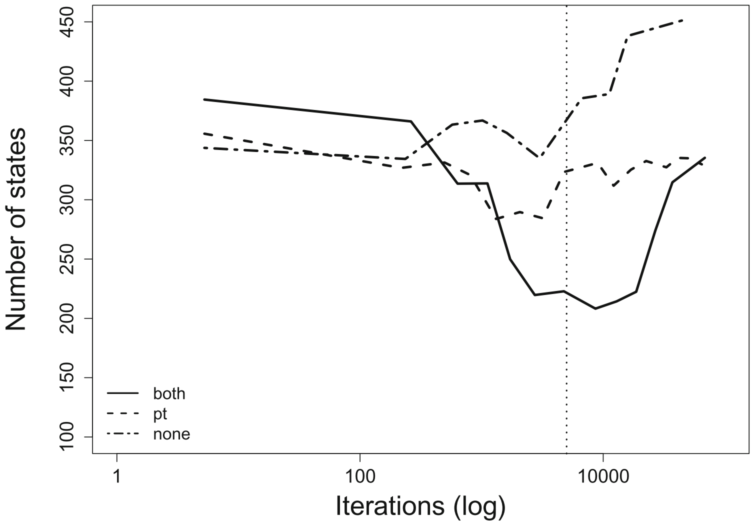

Fig. 2. Average number of different states in the trajectory collections as a function of learning algorithm's iteration. Averages are taken across three different behaviour classes: *both* class ↔ successful BN-robots, *pt* class ↔ BN-robots able to perform phototaxis only, *none* class ↔ failing BN-robots.

In Figure 2 the average number of states in the trajectory collection for each class of robots is plotted along the training phase. The dashed vertical line denotes iteration 5000 at which the objective function was changed so as to include also the evaluation on the antiphototaxis behaviour. We observe that the successful BN-robots are characterised by a decreasing number of unique states up to iteration 5000, when the BN is forced to accomplish a more complex behaviour and the number of states starts to increase, meaning that the training process is still acting so as to adapt the BN-robot to achieve the compound task. BN-robots able to perform phototaxis only show a similar but far less marked pattern, whilst—as expected—worse BN-robots show no tendency to generalise.

3.2 Number of Fixed Points

Some states in BN-robot trajectories are repeated until a change occurs in the input. With slight abuse of term, we call these states *fixed points*. These states represent simple functional building blocks of the type `while <condition> do <action>` (e.g. “turn right until the light input changes”) which are combined to achieve a global behaviour. The emergence of fixed points reveals that the BN is able to extract regularities in the environment and to classify them.

The curves in Figure 3 show that the average number of fixed points in the successful BN-robots increases with training and it consistently increases when the more complex task has to be learned. Instead, the BN-robots of the other

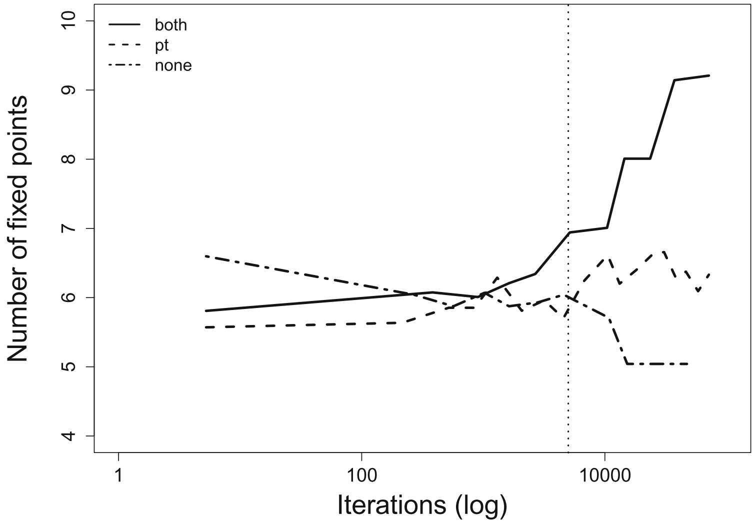

Fig. 3. Average number of fixed points as a function of learning algorithm's iteration. Averages are taken across three different behaviour classes: *both* class ↔ successful BN-robots, *pt* class ↔ BN-robots able to perform phototaxis only, *none* class ↔ failing BN-robots.

two classes maintain approximately the same number of fixed points along the training.

3.3 Statistical Complexity

An analysis of the trajectories of a system may also be focused to capture a further notable dynamical property, which is usually called *statistical complexity* [3,7,8,15,20,27]. This quantity is aimed at estimating to what extent a system works at the edge of order and disorder, i.e. in critical regime. Critical regimes may provide an optimal trade-off between reliability and flexibility, i.e. they make the system able to react consistently with the inputs and, at the same time, capable to provide a sufficiently large number of possible outcomes. This conjecture has been introduced with the expression "computation at the edge of chaos" [4,13,18] and it is supported by results on different computational models such as ϵ-machines [30], cellular automata [9], and neural networks of different kinds [2,12,14].

A system that does not change in time (i.e. in the ordered regime), as well as a system characterised by random behaviour (i.e. in the disordered regime) should be evaluated with low complexity. High complexity is expected to characterise systems in the critical regime accomplishing non trivial tasks. Several measures have been proposed [20] to account for *statistical complexity* (SC), i.e., the algorithmic complexity of a program that reproduces the statistical proper-

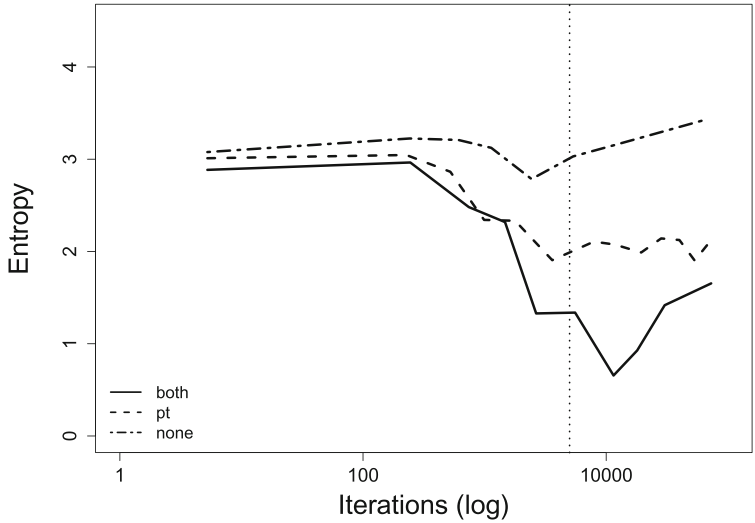

Fig. 4. Entropy of the BN controller as a function of the learning algorithm's iteration. Averages are taken across three different behaviour classes: *both* class ↔ successful BN-robots, *pt* class ↔ BN-robots able to perform phototaxis only, *none* class ↔ failing BN-robots.

ties of a system. In this light, the SC of both a constant and a random sequence is low.

Among various measures of SC, we have chosen a simple yet effective one, which is called LMC complexity, by the name of its inventors [15]. The idea is rather simple: if we want the SC of a system to be high in intermediate regions between order and disorder, we can define it as the product of a measure that increases with disorder and another which decreases with it. The first measure is the Shannon *entropy*, computed over the frequency of the states traversed by the BN-robot. If the BN-robot traverses states $x \in X$ with probabilities $P(x)$ estimated by means of their frequencies, the entropy is defined as:

$$H(X) = -\sum_{x \in X} P(x)\; logP(x)$$

In the definition of $H(X)$ we assume $0\; log0 = 0$.

The second measure contributing to SC is *disequilibrium*:

$$D(X) = \sum_{x \in X} \left(P(x) - \frac{1}{|X|}\right)^2$$

The disequilibrium estimates the extent to which a system exhibits patterns far from equidistribution. For example, if the trajectory of a system is composed of only few of the possible states (e.g., a short cyclic attractor), then it has a high disequilibrium.

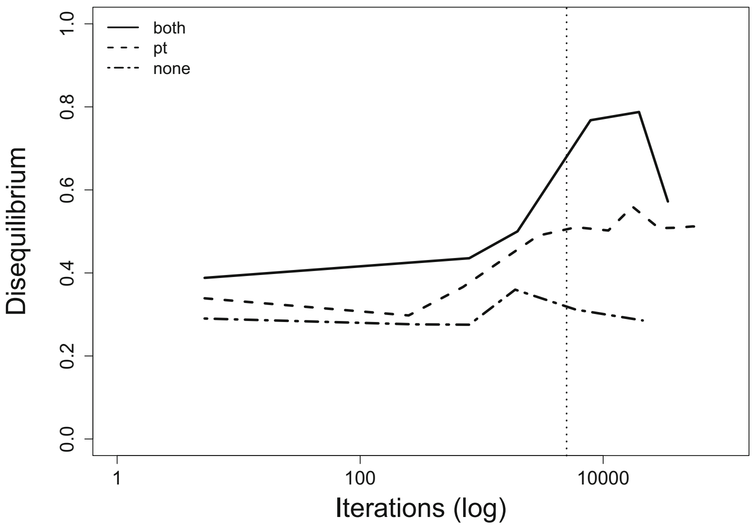

Fig. 5. Disequilibrium of the BN controller as a function of the learning algorithm's iteration. Averages are taken across three different behaviour classes: *both* class ↔ successful BN-robots, *pt* class ↔ BN-robots able to perform phototaxis only, *none* class ↔ failing BN-robots.

Finally, the LMC complexity is defined as:

$$C(X) = H(X) \cdot D(X)$$

A high entropy means that the sequences of states in the BN trajectories are highly diversified. Conversely, a high disequilibrium among the states characterises trajectories mostly composed of the repetition of few states. It is conjectured that a complex system operates in a dynamical regime such that a balance between these two quantities is achieved [15].

It is quite informing to separately observe the three measures, namely entropy, disequilibrium and complexity. In Figure 4 the entropy of BN controllers is shown along the adaptive process. As in previous graphs, the average value for the three performance classes is plotted. Notably, the entropy of well performing BN-robots decreases up to the fitness function change, providing evidence that the adaptive process is successfully achieving generalisation of the task. At iteration 5000, when the fitness function is change so as to include also antiphototaxis, the entropy starts to increase as the BN is adapting to the new task to be accomplished. The reason for this increase has to be ascribed to the adaptive process which does not seem to be completed for all the best performing BNs at the 10000th iteration. The entropy of BN robots that do not perform the complete task shows instead a different behaviour, as it just slightly decreases in the case of BN-robots performing phototaxis only, while it even increases for the worst performing BN-robots.

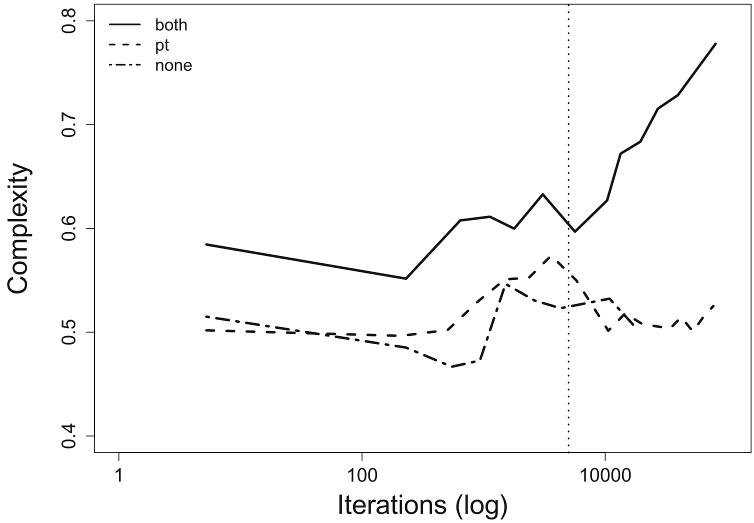

Fig. 6. Complexity of the BN controller as a function of the learning algorithm's iteration. Averages are taken across three different behaviour classes: *both* class ↔ successful BN-robots, *pt* class ↔ BN-robots able to perform phototaxis only, *none* class ↔ failing BN-robots.

Disequilibrium shows a complementary behaviour with respect to entropy, as illustrated by Figure 5. Finally, as shown in Figure 6, the complexity of the successful BN-robots increases steadily during the training process, whilst it is almost constant for the unsuccessful ones. This result supports the conjecture that complexity characterises systems that perform non-trivial tasks [4,15]. Nevertheless, this point deserves further investigations, especially to be compared with previous work on similar subjects [5,10] in which the relation between fitness and complexity is addressed.

4 Conclusion and Future Work

The main finding of the analysis of the trajectories of BN-robots is that the networks that optimally balance robustness and evolvability are characterised by generalisation capability and high statistical complexity of their trajectories. Even if preliminary, these results suggest that also artificial systems that has to cope with changing environments may have an advantage in enjoying the same properties. In the settings in which this hypothesis turned out to hold, additional information for both training and analysing these systems would be available. In particular, the evaluation of features such as fixed points and complexity may be profitably incorporated into the objective function of the adaptive process, with the goal of guiding it towards high performing networks.

Conversely, experiments on simple artificial systems provide a controlled environment for studying general properties of living systems. The use of BNs and their trajectories make it possible to link results in digital worlds with biological ones, as these models have been proven to capture relevant biological phenomena.

The results presented in this paper concern preliminary experiments on the subject, which may be further investigated in several directions. The robustness of the results against changes in the search strategy and input and output encoding should be assessed.

In the next future, we plan to investigate the relation between complexity measures and performance of BN-robots in noisy and varying environments. First of all, this is expected to provide guidelines for the automatic design of truly adaptive robotic systems; furthermore, we aim at contributing elucidate the elusive interplay among complexity, robustness and evolvability.

Acknowledgments. We thank the anonymous referees who carefully read the paper and provided pertinent and valuable suggestions for preparing this final version.

References

1. Aldana, M., Balleza, E., Kauffman, S., Resendiz, O.: Robustness and evolvability in genetic regulatory networks. Journal of Theoretical Biology **245**, 433–448 (2007)
2. Bertschinger, N., Natschläger, T.: Real-time computation at the edge of chaos in recurrent neural networks. Neural Computation **16**, 1413–1436 (2004)
3. Crutchfield, J.: The calculi of emergence: Computation, dynamics, and induction. Physica D **75**, 11–54 (1994)
4. Crutchfield, J., Young, K.: Computation at the onset of chaos. In: Complexity, Entropy, and Physics of Information. Addison Wesley (1990)
5. Edlund, J., Chaumont, N., Hintze, A., Koch, C., Tononi, G., Adami, C.: Integrated information increases with fitness in the evolution of animats. PLOS Computational Biology **7**(10), e1002236:1–e1002236:13 (2011)
6. Garattoni, L., Roli, A., Amaducci, M., Pinciroli, C., Birattari, M.: Boolean network robotics as an intermediate step in the synthesis of finite state machines for robot control. In: Liò, P., Miglino, O., Nicosia, G., Nolfi, S., Pavone, M. (eds.) Advances in Artificial Life, ECAL 2013, pp. 372–378. The MIT Press (2013)
7. Gell-Mann, M., Lloyd, S.: Information measures, effective complexity, and total information. Complexity **2**(1), 44–52 (1996)
8. Grassberger, P.: Randomness, information, and complexity, August 2012. arXiv:1208.3459
9. Hordijk, W.: The EvCA project: A brief history. Complexity **18**, 15–19 (2013)
10. Joshi, N., Tononi, G., Koch, C.: The minimal complexity of adapting agents increases with fitness. PLOS Computational Biology **9**(7), e1003111:1–e1003111:10 (2013)
11. Kauffman, S.: The Origins of Order: Self-Organization and Selection in Evolution. Oxford University Press, UK (1993)
12. Kinouchi, O., Copelli, M.: Optimal dynamical range of excitable networks at criticality. Nature Physics **2**, 348–351 (2006)
13. Langton, C.: Computation at the edge of chaos: Phase transitions and emergent computation. Physica D **42**, 12–37 (1990)

14. Legenstein, R., Maass, W.: Edge of chaos and prediction of computational performance for neural circuit models. Neural Networks **20**, 323–334 (2007)
15. Lopez-Ruiz, R., Mancini, H., Calbet, X.: A statistical measure of complexity. Physics Letters A **209**, 321–326 (1995)
16. Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptocz, A., Magnenat, S., Zufferey, J.C., Floreano, D., Martinoli, A.: The e-puck, a robot designed for education in engineering. In: Gonçalves, P., Torres, P., Alves, C. (eds.) Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions, vol. 1, pp. 59–65 (2009)
17. Nykter, M., Price, N., Aldana, M., Ramsey, S., Kauffman, S., Hood, L., Yli-Harja, O., Shmulevich, I.: Gene expression dynamics in the macrophage exhibit criticality. In: Proceedings of the National Academy of Sciences, USA, vol. 105, pp. 1897–1900 (2008)
18. Packard, N.: Adaptation toward the edge of chaos. In: Dynamic Patterns in Complex Systems, pp. 293–301 (1988)
19. Pinciroli, C., Trianni, V., O'Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Di Caro, G., Ducatelle, F., Birattari, M., Gambardella, L., Dorigo, M.: ARGoS: a modular, multi-engine simulator for heterogeneous swarm robotics. Swarm Intelligence **6**(4), 271–295 (2012)
20. Prokopenko, M., Boschetti, F., Ryan, A.: An information-theoretic primer on complexity, self-organization, and emergence. Complexity **15**(1), 11–28 (2008)
21. Ribeiro, A., Kauffman, S., Lloyd-Price, J., Samuelsson, B., Socolar, J.: Mutual information in random Boolean models of regulatory networks. Physical Review E **77**, 011901:1–011901:10 (2008)
22. Roli, A., Manfroni, M., Pinciroli, C., Birattari, M.: On the design of boolean network robots. In: Di Chio, C., Cagnoni, S., Cotta, C., Ebner, M., Ekárt, A., Esparcia-Alcázar, A.I., Merelo, J.J., Neri, F., Preuss, M., Richter, H., Togelius, J., Yannakakis, G.N. (eds.) EvoApplications 2011, Part I. LNCS, vol. 6624, pp. 43–52. Springer, Heidelberg (2011)
23. Roli, A., Villani, M., Serra, R., Garattoni, L., Pinciroli, C., Birattari, M.: Identification of dynamical structures in artificial brains: an analysis of boolean network controlled robots. In: Baldoni, M., Baroglio, C., Boella, G., Micalizio, R. (eds.) AI*IA 2013. LNCS, vol. 8249, pp. 324–335. Springer, Heidelberg (2013)
24. Serra, R., Villani, M.: Modelling bacterial degradation of organic compounds with genetic networks. Journal of Theoretical Biology **189**(1), 107–119 (1997)
25. Serra, R., Villani, M., Barbieri, A., Kauffman, S., Colacci, A.: On the dynamics of random Boolean networks subject to noise: Attractors, ergodic sets and cell types. Journal of Theoretical Biology **265**(2), 185–193 (2010)
26. Serra, R., Villani, M., Semeria, A.: Genetic network models and statistical properties of gene expression data in knock-out experiments. Journal of Theoretical Biology **227**, 149–157 (2004)
27. Shalizi, C.: Methods and techniques of complex systems science: An overview, March 2006. arXiv:nlin/0307015
28. Shmulevich, I., Dougherty, E.: Probabilistic Boolean Networks: The Modeling and Control of Gene Regulatory Networks. SIAM, Philadelphia (2009)
29. Shmulevich, I., Kauffman, S., Aldana, M.: Eukaryotic cells are dynamically ordered or critical but not chaotic. PNAS **102**, 13439–13444 (2005)
30. Strogatz, S.: Nonlinear dynamics and chaos. Perseus Books Publishing (1994)
31. Villani, M., Serra, R.: On the dynamical properties of a model of cell differentiation. EURASIP Journal on Bioinformatics and Systems Biology **4**, 1–8 (2013)