An investigation of graceful degradation in Boolean network robots subject to online adaptation

Michele Braccini 0000−0002−2335−2514
Paolo Baldini 0000−0002−3625−5414
Andrea Roli 0000−0001−9891−5441

{m.braccini,p.baldini,andrea.roli}@unibo.it

Department of Computer Science and Engineering (DISI), Campus of Cesena,
Alma Mater Studiorum Università di Bologna, Italy
European Centre for Living Technology, Venice, Italy

Abstract

The ability to resist to faults is a desired property in robotic systems. However, it is also hard to obtain, having to modify the behavior to face breakdowns. In this work we investigate the robustness against sensor faults in robots controlled by Boolean networks. These robots are subject to online adaptation—i.e., they adapt some structural properties while they actually act—for improving their performance at a simple task, namely phototaxis. We study their performance variation according to the number of faulty light sensors. The outcome is that Boolean network robots exhibit graceful degradation, as the performance decreases gently with the number of faulty sensors. We also observed that a moderate number of faulty sensors—especially if located contiguously—not only produces a negligible performance decrease, but it can be sometimes even beneficial. We argue that online adaptation is a key concept to achieve fault tolerance, allowing the robot to exploit the redundancy of sensor signals and properly tune the dynamics of the same Boolean network depending on the specific working sensor configuration.

Keywords: Boolean networks, Robot, Online adaptation, Graceful degradation, Fault-tolerance.

Introduction

Research in robotics has often been influenced by biology. Sometimes the goal is to explore natural behaviors in a controlled setting, some others to bring the effectiveness of natural strategies to artificial applications [5, 9, 11, 17, 19].

Nevertheless, one often neglected property is the natural ability to survive to faults. The motivations are many, and rooted in the differences between artificial and biological bodies. For instance, organisms are able to recover from limited damages regenerating tissues, while machines are not. Even the perception capabilities differ, with biological bodies having many more redundant and degenerate sensors. This directly enhances the survival of an organism, since in case of faults essential information can be retrieved—and reconstructed—from the remaining signals. From this point of view, the gap between artificial and biological agents is still huge and far from being filled.

Although we cannot expect to achieve the same resilience of living organisms soon, we can still aim at improving the current state of robotics. Many robots work in a fragile equilibrium, where a single breakdown can lead to the catastrophic fault of the entire system. This problem is too often ignored, as it is frequently possible to reach the robot and recover the damage. Nevertheless, this binds the robot existence to us and prevents the creation of really autonomous robots, i.e., robots able to operate without the need of human intervention [18]. Autonomy is important especially in situations in which the robot have to act in an isolated environment, for example during rescue operations in hostile environments. We believe that the first step to obtain autonomous robots consists in making them less dependent on human support. This can be achieved by improving their resilience to faults, i.e., by making them able to continue their duties when subject to breakdowns. Obviously, the limited amount of sensors and actuators in a robot still makes it sensible to faults. Nevertheless, we aim at a graceful degradation of performance according to the amount of damage. This property alone is already an important improvement over the current state, where the lack of a single cue can lead to a critical decrease in performance.

Previous attempts to tackle the problem consisted in defining an internal representation of the robot structure, allowing to adapt the behavior as soon as some part of the body were detected to be damaged [4]. The robot synthesizes a new behavior simulating its self-model and finding the most suitable solution to the fault. This technique is very powerful, but also highly computationally demanding. This makes the approach unsuitable for computationally limited robots. In this work we build upon previous works of online adaptation in order to provide an alternative approach for minimally cognitive agents [1, 5, 6]. The idea is to continuously adapt the behavior with a simple strategy in order to maximize a performance metric. We expect this approach to lead to a graceful degradation of performance, as the reaction of the robot adapts to exploit the best information available. However, as the adaptation process is not instantaneous, we expect the performance to briefly drop before each recovery. The decrease directly correlates with the importance of the faulted signal in the reaction of the robot. Therefore, it is expected that different faults will recover in different amounts of time. In order to overcome this problem, we propose to analyze the performance of a robot that does not have any innate behavior or any sort of reactive mechanism. This allows us to assess the average time needed to recover, regardless of any previous known behavior.

Our investigation consists in evaluating the graceful degradation of performance in robots controlled by Boolean networks (BN). During the adaptation, only the couplings between robot sensors and BN nodes undergo changes, leaving the BN itself unchanged. The performance is evaluated in a phototaxis task, in which the robot has to reach a light source. The results aim at providing a first assessment as to whether online adaptation induces a graceful degradation of performance in BN-robots that perform phototaxis. Finally, we also try to understand to what extent the kinds of damage affect the performance.

This paper is structured as follows. In section 1 we give an introductory description of BNs and their working principle. In section 2 we describe the adaptive approach used in the experiments. Section 3 delves into experimental details. Following, section 4 presents and discusses the results of the study. Finally, in the conclusion section we summarize the work done and we illustrate future explorations on the topic.

1 Boolean networks

The core of the robot controller used in this work is a Boolean network (BN). BNs can be represented as a discrete set of Boolean variables and Boolean functions. The state of a variable depends on the state of the neighbors and on the transition function that controls it. Since their introduction as an abstract model of gene regulatory networks [14], BNs have been the subject of many works investigating their computational and dynamical properties. Notably, they effectively capture significant biological phenomena, such as cell differentiation [7, 8, 12, 13, 16, 23, 24]. One interesting application is in the context of robotics, where they have been used to develop adapting and evolving[1] robots [5, 22].

A specific class of BNs is Random Boolean Networks. Those are created randomly according to some rules, such as the (possibly average) number of neighbors K of each node. Also the transition functions are randomly generated, usually creating a mapping from the state of the neighbors to the resulting state of the variable. This can be done by creating a table enumerating the possible input states and by randomly setting the corresponding output. The setup can be influenced by a bias p representing the probability of an output to be set to 1. Different combinations of values for p and K change the dynamic of the network towards different dynamic regimes [15]. BNs possess an interesting dynamics that can be divided in two regimes: order and chaotic. The ordered region is characterized by a short propagation of the signals, and by the tendency of the system to return to a stable state. The chaotic regime enhances and favors the perpetuation of perturbations that may permanently change the state evolution of the system. The result is that the state of a chaotic network is never stable. A critical region lies at the boundary between these two regimes, separating them. In this condition the perturbations propagate more than in the ordered region, but fade away faster than in the chaotic one. Many works suggest critical BNs favors computation, and indeed have been successfully used to complete different

[1] This mostly thanks to their simple encoding and mutation.

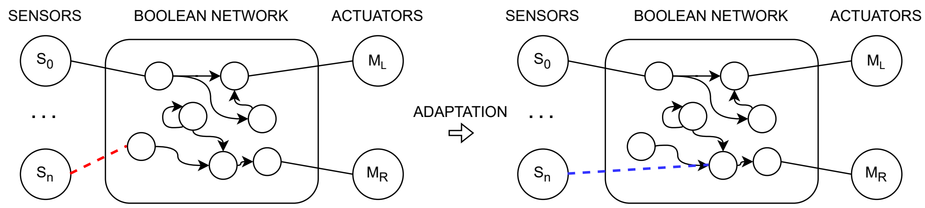

Figure 1: Example of an adaptation step. The subset of sensor-node couplings to be reconnected is indicated by a red dashed line on the left. The reconnected couplings are represented by a blue dashed line on the right. As visible, the adaptation affects only the node to which a sensor connects, and leaves unchanged the topology of the BN.

tasks, such as classification, filtering and control [21]. Even BNs evolved to solve specific tasks were often found to be critical [2].

2 Adaptation

We claim that the ability to adapt is a major facilitator to achieve a graceful degradation of performance. In this work, we use a previously proposed adaptive mechanism for the adaptation of the robot behavior [6]. This is designed to work on network based control systems, and has been already successfully used for the adaptation of robots controlled by random BNs.

In this schema, the BN mediates the signals from the sensors in order to produce an adequate output to control the actuators. The BN itself is random, and never changes during the experiment. Indeed, the adaptive process does not directly affect the BN. The only way to modify the behavior of the robot is thus by modifying the couplings between network nodes and sensors/actuators (see Figure 1). Specifically, in this work we only change the sensory couplings, i.e., those from the sensors to the BN. The idea is that it is possible to perturb the internal state of the BN in order to generate a desired response, i.e., a desired output which is used to properly control the actuators. Differently, we can state that the adaptation modifies the point of view of the robot, increasing or decreasing the focus on (i.e., effect of) some signals. When the robot behavior adapts, a sub-set of couplings between the robot light sensors and a BN controller is chosen and modified. The new couplings must connect to different nodes of the BN. Moreover, no more than one sensor can connect to a single node. Previous works demonstrated that this simple adaptive strategy is powerful enough to produce complex behaviors [5].

3 Experimental setting

The robot used in the experiment is a *foot-bot*, simulated in ARGoS [3, 20]. The foot-bots are equipped with 24 light sensors placed in a circular ring around the chassis of the robot, and can move by acting on the motors of two wheels. Before being passed to the BN, the values of the sensors are encoded in binary form according to a threshold $t = 0.2$. The outputs provided to the two wheels are either 0 or 1, converted either into a stopped wheel or a wheel that provides a fixed linear velocity ($2cm/s$ in our experiments). A modulation of the speed is possible by controlling the output at each step.

The task on which the performance degradation will be assessed is phototaxis. This consists in moving as fast as possible towards a light source, and staying close to it once it is reached. The experiment takes place in a squared arena with an edge equal to $100m$. The destination marked by the light is the center of the arena, while the starting point is in one of its corners. Robots undergo an adaptation process composed of 1200 epochs, each lasting one minute in simulation time. At the start of each epoch the robot adapts its controller, that is then evaluated at the end the epoch. At the end of each epoch the robot evaluates its performance as the difference between initial and final distance from the light. For ease of the experiments we provide this information directly to the robot. However, in principle robots are situated in their environment and should use only the information they can get from their sensors. For example, in this case we could use the variation of light intensity read by the light sensors. If the performance is not worse than the best one achieved so far, the incumbent mapping becomes the starting point for subsequent adaptations. Otherwise, the previous mapping is restored and adapted once more. In this experiment the adaptation can affect up to 6 connections from sensors to BN nodes.[2]

As we are in an online setting, besides the final distance to the light, we are also interested in the life performance of the robot, i.e., in the performance they achieve in the whole duration of a single experiment. For this reason, we assess the performance according to three different measures: *(a)* the distance from the light along the epochs; *(b)* the *run length distribution* (RLD), which represents the fraction of successful runs at each epoch;[3] *(c)* the final distance from the light.

Since a real robot may undergo many types of fault, we analyze how the performance changes in different situations. (i) The first case considers the event in which the sensors simply stop to produce any output. For instance, this may be due to the disconnection of a cable. Although the BN node to which the broken sensor connects can be changed, its signal will never affect the internal state of the BN controller. (ii) The second case consists in sensors producing random outputs. Differently from before, those still perturb the BN internal state. However, their signal can be considered noise, and is therefore expected to be detrimental. (iii) The third scenario considers the output of the

[2] The number of changes is randomly chosen in 1–6.

[3] We consider a run successful if the robot reaches a point in the arena at a distance less than or equal to a given threshold value d_θ.

faulted sensors to block in a fixed (random) state. This simulates the case of a short circuit in which the output connection touches the source voltage or the ground. All these type of damages can affect randomly picked sensors (case R) or a contiguous set of sensors (case A). The aim is to represent both random faults and breakdowns due to external events. For instance, a robot passing near a strong heat source or crashing against an obstacle may be subject to localized damages. In this case, the affected sensors would likely be physically contiguous.

We tested the working conditions with 0, 3, 6, 9, ..., 24 faulty sensors. The faulty sensors are chosen at random at the beginning of each replica of the experiment. We collected results on 300 replicas for each type of damage.

4 Results

We expect to identify a graceful degradation of performance, where it decreases according to the number of faulted sensors. Overall, the results support this hypothesis, as we can observe from the barplots depicting the average final distance from light reached at the end of the 1200 epochs (see Figure 2). This function estimates the capability of the robot to adapt to faults in the long run. Apart from some fluctuations due to variance in the experiments, the higher the number of faulty sensors, the higher the final distance from light. Notably, the performance decreases sensibly when the number of faulty sensors is fairly high.

The RLD is computed at a given target distance d_θ from the light and provides an estimation of the efficiency of the adaptive process to exploit the (limited) resources the robots can use: the steeper and the higher the curve, the better the performance at reaching for the first time a distance less than or equal to d_θ from the light. We computed the RLD with $d_\theta = 1\ m$ and $d_\theta = 5\ cm$. The plots are depicted in Figure 3 and Figure 4, respectively. Also in this case we observe a gradual decline in performance, as the number of defective sensors increases. Nevertheless, here a striking difference emerges between the cases R and A: when contiguous sensors are damaged, the RLD does not decrease until a large fraction of sensors is touched. In some cases (i and ii) it seems even better to have few sensors out of order. A possible explanation of this phenomenon is that when contiguous sensors are damaged the remaining sensors are sufficient to provide the robot the necessary information to navigate the arena correctly. Less inputs also imply a search among a lower number of working configurations, therefore the adaptation process can proceed faster.

Finally, the distance from light at each epoch provides a complementary view of robot performance, as it captures the performance in time, taking also into account the performance achieved in the adaptation attempts (i.e. the adaptation epochs). The distance from light—averaged across the replicas for each epoch—is plotted in Figure 5. We can observe a qualitative behavior analogous to the previous cases. It is remarkable that, also under this evaluation, the performance with few damaged sensors is still equivalent to—or even better than—the one in which all sensors are working. This observation reinforces our previous hypothesis, as it provides an online view of the impact of having less information. We conclude by observing that when all the 24 sensors do not work and just provide noise or fixed random values, the distance from light decreases anyway, as some lucky combinations can drive the robot towards the light even if it does not perceive it.

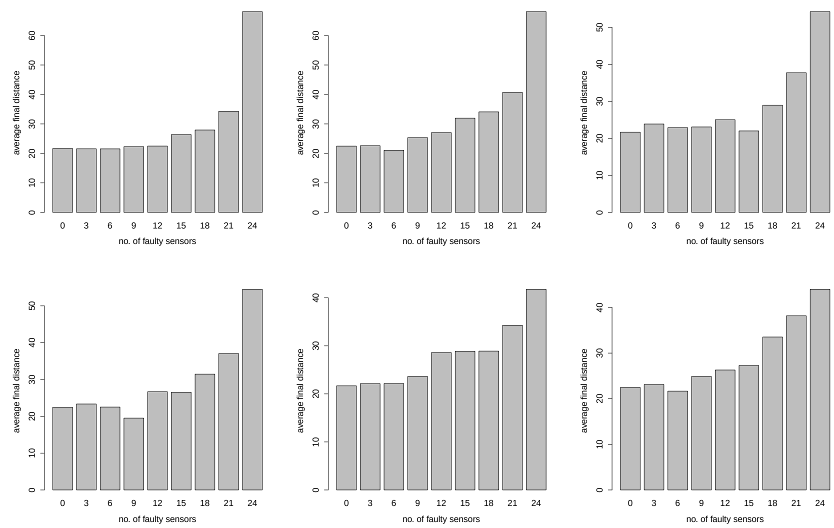

Figure 2: Barplots of final distance from the light, averaged across 300 replicas. Columns: left, case R – random picked sensors; right, case A – contiguous faulty sensors. Rows corresponds to kinds of damage (from the top, i: detached sensor, ii: noisy sensor, iii: random fixed value).

Conclusion

In this work we tested how a robot able to continuously modify its behavior adapts to faults in its sensors. The goal is to verify if an online adaptation allows the performance of the robot to recover to a similar level of the start. The expected result is that the performance will keep decreasing according to the amount of damage, possibly gracefully.

The results of our experiment suggest that indeed online adaptation induces a graceful degradation of performance. The distance of the robot from the target increases with the amount of damage, meaning that some lost information cannot be recovered to keep succeeding in the task as well as without being damaged. Nevertheless, this general situation has some notable exceptions. In fact, the results also show that in some conditions the lack of few sensors increases the performance of the robot (or the reduction is negligible). Our hypothesis is that removing some cues to the BN simplifies the correct combination of

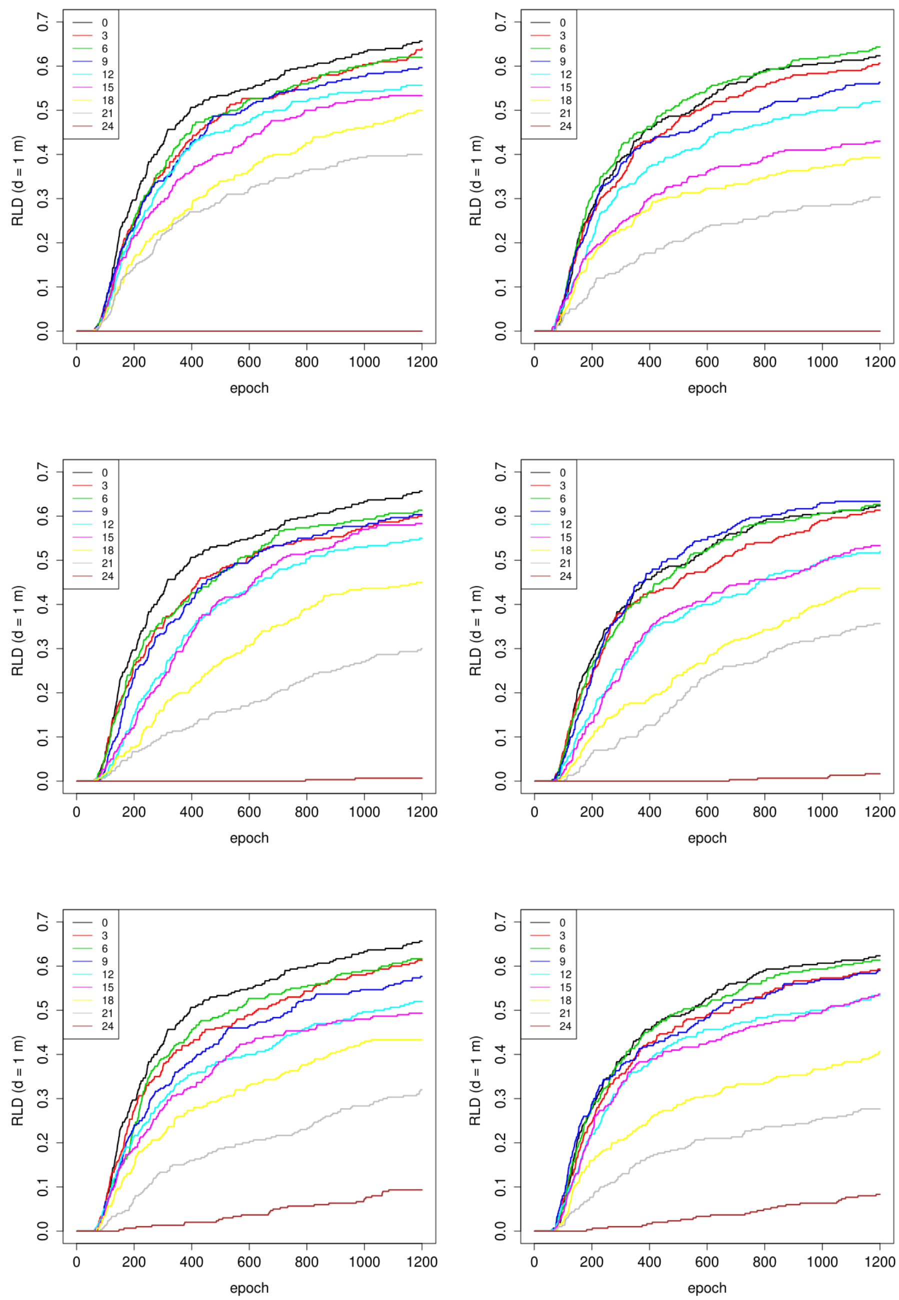

Figure 3: Run length distribution for target distance equal to **1 m**. A point (x, y) in the plot represents the fraction of replicas (y) that achieved a distance less than or equal to 5 cm at epoch x. Columns: left, case R – random picked sensors; right, case A – contiguous faulty sensors. Rows corresponds to kinds of damage (from the top, i: detached sensor, ii: noisy sensor, iii: random fixed value).

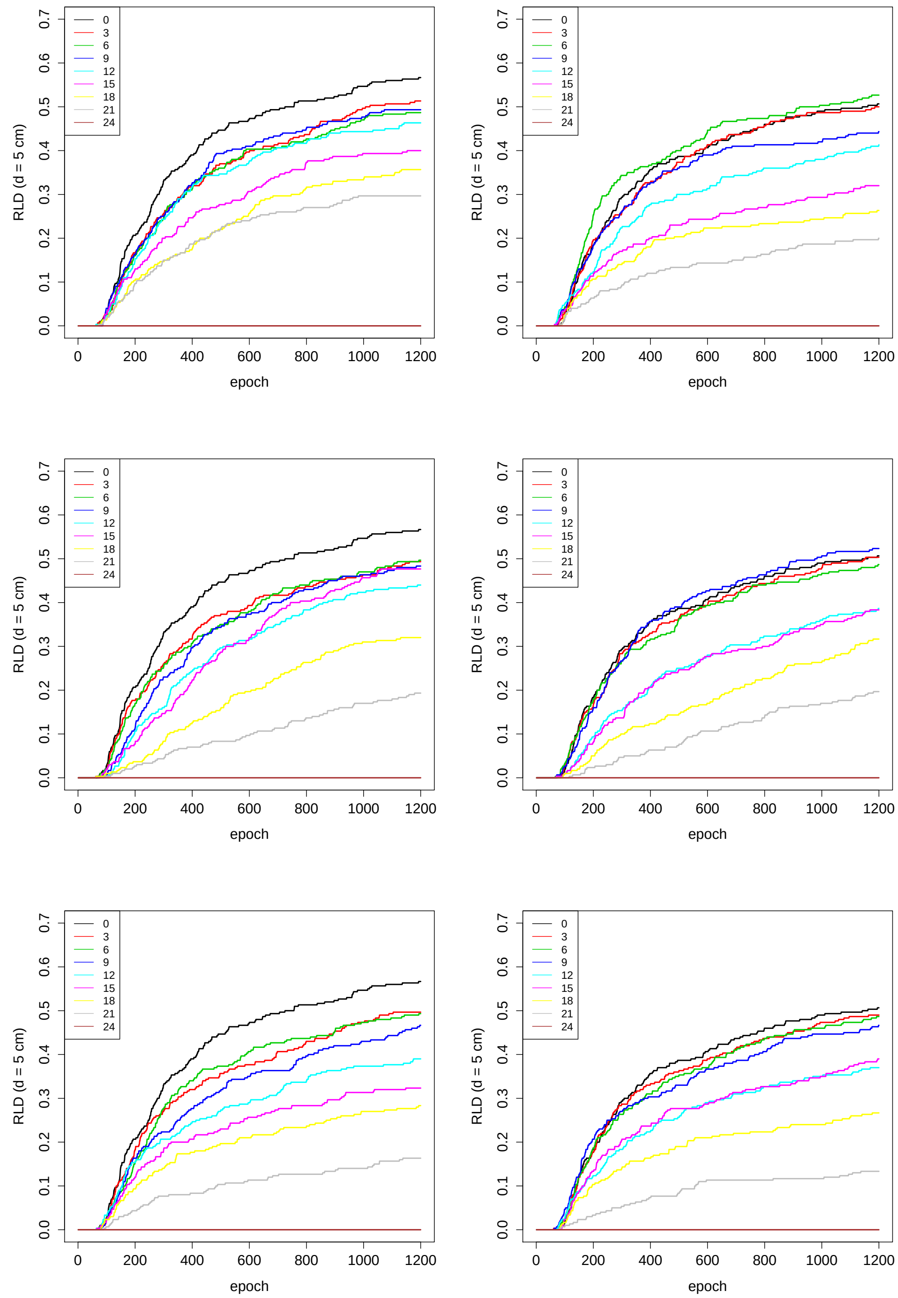

Figure 4: Run length distribution for target distance equal to **5 cm**. A point (x, y) in the plot represents the fraction of replicas (y) that achieved a distance less than or equal to 5 cm at epoch x. Columns: left, case R – random picked sensors; right, case A – contiguous faulty sensors. Rows corresponds to kinds of damage (from the top, i: detached sensor, ii: noisy sensor, iii: random fixed value).

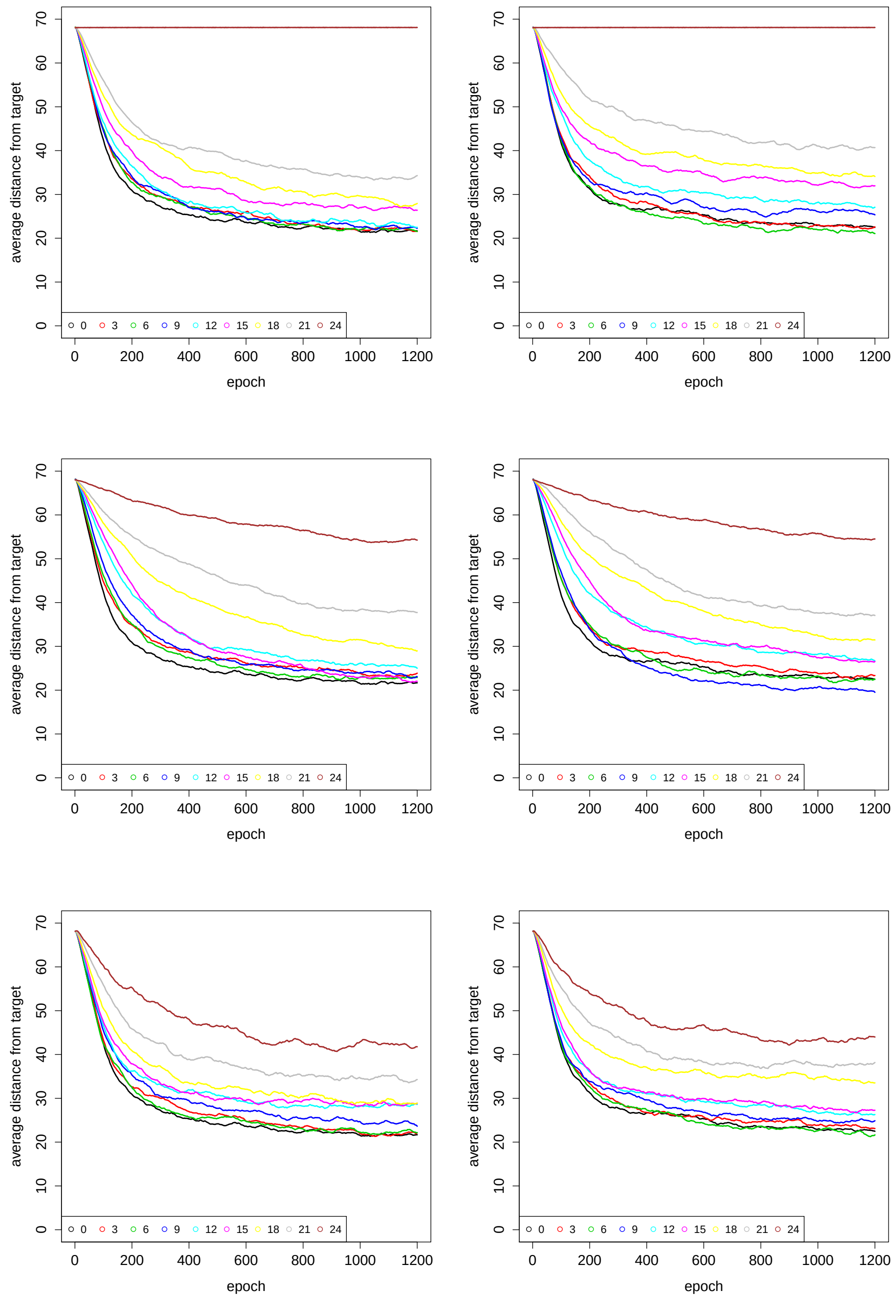

Figure 5: Distance from the light along adaptation epochs, averaged across 300 replicas. Columns: left, case R – random picked sensors; right, case A – contiguous faulty sensors. Rows corresponds to kinds of damage (from the top, i: detached sensor, ii: noisy sensor, iii: random fixed value).

the remaining. In other words, removing some irrelevant signals may help the controller to focus on more relevant ones. This last aspect addresses another crucial point in online adaptation: the contingent optimization of the resources available to the robot.

In a future work, we plan to consider the effect of faulty actuators on the performance. For instance, a wheel may simply start to turn slower due to consumption. In similar situations we expect the controller to adapt in order to modulate the output accordingly. Indeed, the output signal may change in intensity of frequency in order to mediate the fault. Finally, the results of this work push toward the analysis of possible *degeneracy* of sensors [10], i.e. the capability of compensate damages in some sensors with the integration of the information retrieved from other ones. The goal is to assess if different sensors and different *types* of sensors providing different information can be used to induce the same robot behavior.

Acknowledgements

AR acknowledges support from the PRIN 2022 research project of the Italian Ministry of University and Research titled *Org(SB-EAI) – An Organizational Approach to the Synthetic Modeling of Cognition based on Synthetic Biology and Embodied AI* (Grant Number: 20222HHXAX).

References

[1] Baldini, P., Braccini, M., Roli, A.: Online adaptation of robots controlled by nanowire networks: A preliminary study. In: Italian Workshop on Artificial Life and Evolutionary Computation. Springer (2022)

[2] Benedettini, S., Villani, M., Roli, A., Serra, R., Manfroni, M., Gagliardi, A., Pinciroli, C., Biratta ri, M.: Dynamical regimes and learning properties of evolved boolean networks. Neurocomputing **99**, 111–123 (2013)

[3] Bonani, M., Longchamp, V., Magnenat, S., Rétornaz, P., Burnier, D., Roulet, G., Vaussard, F., Bleuler, H., Mondada, F.: The marxbot, a miniature mobile robot opening new perspectives for the collective-robotic research. In: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, October 18-22, 2010, Taipei, Taiwan. IEEE (2010)

[4] Bongard, J., Zykov, V., Lipson, H.: Resilient machines through continuous self-modeling. Science (2006)

[5] Braccini, M., Roli, A., Barbieri, E., Kauffman, S.: On the criticality of adaptive boolean network robots. Entropy **24** (2022)

[6] Braccini, M., Roli, A., Kauffman, S.: A novel online adaptation mechanism in artificial systems provides phenotypic plasticity. In: Italian Workshop on Artificial Life and Evolutionary Computation. Springer (2021)

[7] Braccini, M., Roli, A., Villani, M., Montagna, S., Serra, R.: A simplified model of chromatin dynamics drives differentiation process in boolean models of grn. In: Artificial Life Conference Proceedings. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info ... (2019)

[8] Braccini, M., Roli, A., Villani, M., Serra, R.: Dynamical properties and path dependence in a gene-network model of cell differentiation. Soft Computing (2021)

[9] Dorigo, M., Trianni, V., Şahin, E., Groß, R., Labella, T.H., Baldassarre, G., Nolfi, S., Deneubourg, J.L., Mondada, F., Floreano, D., et al.: Evolving self-organizing behaviors for a swarm-bot. Autonomous Robots (2004)

[10] Edelman, G., Gally, J.: Degeneracy and complexity in biological systems. PNAS (2001)

[11] Floreano, D., Mattiussi, C.: Bio-inspired artificial intelligence: theories, methods, and technologies. MIT press (2008)

[12] Huang, S., Eichler, G., Bar-Yam, Y., Ingber, D.: Cell fates as high-dimensional attractor states of a complex gene regulatory network. Physical Review Letters (2005)

[13] Huang, S., Ernberg, I., Kauffman, S.: Cancer attractors: a systems view of tumors from a gene network dynamics and developmental perspective. In: Seminars in cell & developmental biology (2009)

[14] Kauffman, S.A.: Metabolic stability and epigenesis in randomly constructed genetic nets. Journal of theoretical biology (1969)

[15] Luque, B., Solé, R.: Phase transitions in random networks: Simple analytic determination of critical points. Physical Review E **55**(1) (1997)

[16] Montagna, S., Braccini, M., Roli, A.: The impact of self-loops on boolean networks attractor landscape and implications for cell differentiation modelling. IEEE/ACM transactions on computational biology and bioinformatics (2020)

[17] Nolfi, S., Floreano, D.: Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines. MIT press (2000)

[18] Pfeifer, R., Scheier, C.: Understanding Intelligence. The MIT Press (2001)

[19] Pfeifer, R., Lungarella, M., Iida, F.: Self-organization, embodiment, and biologically inspired robotics. science (2007)

[20] 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 (2012)

[21] Roli, A., Villani, M., Filisetti, A., Serra, R.: Dynamical criticality: overview and open questions. Journal of Systems Science and Complexity **31**(3), 647–663 (2018)

[22] Roli, A., Manfroni, M., Pinciroli, C., Birattari, M.: On the design of boolean network robots. In: Applications of Evolutionary Computation: EvoApplications 2011: EvoCOMPLEX, EvoGAMES, EvoIASP, EvoINTELLIGENCE, EvoNUM, and EvoSTOC, Torino, Italy, April 27-29, 2011, Proceedings, Part I. Springer (2011)

[23] Serra, R., Villani, M., Graudenzi, A., Kauffman, S.: Why a simple model of genetic regulatory networks describes the distribution of avalanches in gene expression data. Journal of Theoretical Biology (2007)

[24] Villani, M., Barbieri, A., Serra, R.: A dynamical model of genetic networks for cell differentiation. PloS one (2011)