

Uncontrolled Engineering: A Review of Nolfi and Floreano's *Evolutionary Robotics*

Hod Lipson
Computational Synthesis Lab

*Mechanical and Aerospace Engineering Department and
Faculty of Computing and Information
Cornell University, Ithaca, NY 14853-2801 USA*

YOU CAN TELL A FIELD IS REACHING MATURITY when its first textbook is written. One of the challenges in writing a textbook on such a leading edge and active topic as *Evolutionary Robotics*¹, is sorting through the wealth of new ideas and current research, and determining what material is sufficiently established to include, and what is still in the realm of speculation, yet to withstand proof or refutation. Invariably, some exciting new ideas are left out in the process, but Stefano Nolfi and Dario Floreano make a good job of this task. It is especially helpful that the authors are themselves pioneers in the field, and their own groundbreaking work helped shape its progress and achievements. The result is a well-written and well-balanced book, containing both theoretical principles and experimental support, as well as practical advice from over ten years of experience. This is one of these books from which you learn much more than you initially set out to – in many places it digresses to elaborate on by-topics and other fields, providing an overall well rounded account of *evolutionary design of physical robotic control systems*, with the two-wheeled *Khepera* robot as its main pedagogical platform. The entry cost for outsiders is low, and I would certainly use this book for teaching the topic in a graduate level course.

Here I will outline the main topics covered by the book, and add some high-level criticism where I believe certain important assumptions can be challenged. These notes are directed mainly at the field, rather than the book itself as it merely reflects the current state of the art, and as an Evolutionary Robotics practitioner myself I feel this is part an introspective confrontation. Specifically, I will challenge the overwhelming focus on evolution of brain rather than body *and* brain, or robot as a whole, as does evolution in nature. Second, I will challenge the presumption that evolution in physical *reality* is critically important over evolution in physical *simulation* (although we have done this in our lab at Brandeis as well, this is far from settled^{2,3}). Third, and most importantly, I question the ability of current evolutionary robotics techniques to scale to solve more complex tasks. Scaling is at the crux of success of evolutionary methods, and cannot be taken for granted, as the history of AI has taught us more than once before⁴.

The spirit conveyed by the subtitle of this book – *the biology, intelligence, and technology of self-organizing machines* – captures the essence of an exciting new paradigm of machine design by self-organization. But the authors quickly regress to a much narrower view of design of robot controllers:

“Evolutionary Robotics is a new technique for automatic creation of autonomous robots [...] the term has been introduced quite recently, but the idea of representing the control system of a robot as an artificial chromosome subject to the laws of genetics and of natural selection dates back to the 80’s...” (Preface)

In an excellent introduction, Nolfi and Floreano review the relationship between Evolutionary robotics and close fields in AI and engineering. Evolutionary Robotics shares many ideas with *Behavior-Based Robotics*, where global behavior that was traditionally pre-programmed into robot brains, is expected to emerge from local interaction between a set of basic behaviors of sensing and actuation⁵. Evolutionary robotics places an even more stringent constraint on the design process itself: The architecture or assembly plan of the basic building blocks should also emerge as a process of self-organization, typically involving variation and selection. Evolutionary robotics also shares the *machine learning* ideas that control systems can be trained using incomplete data and sparse reinforcement, to produce robust controllers that can generalize to unknown situations⁶. It is perhaps this focus on *behavior*, as well as practical constraints, that have led the field of Evolutionary robotics to focus primarily (wrongfully, in my opinion) on evolution of *control* systems alone, while keeping the morphology fixed, as if the morphology of the robot has nothing to do with its behavior or is ‘given’ as a constraint. Evolutionary robotics also shares a great deal with the field of Artificial Life of course, as a methodology for understanding life phenomena through their reproduction in artificial systems, with a focus on physical substrates. This biological perspective is especially interesting when addressing questions about the mechanisms that allow for the emergence of complexity. The book mentions various mechanisms, such as incremental evolution, lifetime learning as a guide to the evolutionary process, co-evolution, and developmental processes, including the evolution of evolvability and the representation problem⁷. The authors also point out an *engineering* perspective to this approach (more on this important relationship later).

Indeed, the focus of evolutionary robotics has been on ideas concerned primarily with evolution of robot controllers⁸, and even *Evolutionary Hardware* is primarily limited to *control* hardware. But in nature there is never a body without a brain, or a brain without a body. Like a chicken and egg neither came first, and neither can survive a variation that surpasses that adaptive capability of the other. So why neglect this important duality in evolution? Besides practical reasons (controllers are easier to change and test than bodies), and social reasons (computer scientists design brains, and mechanical engineers design bodies), the scope of problems that have been addressed by Evolutionary robotics to date are simple enough to solve only given a particular morphology (and this constraint enters the researcher’s mind when choosing a platform). But for more complex tasks in the future, any given morphology may be inadequate, or moreover, a fixed morphology may prove to be a constraint that is too difficult. I therefore believe that as we move to more complex tasks (and that is the ultimate challenge), this arbitrary body/brain boundary will necessarily be relaxed. This, in turn, brings additional aspects into play, for example now body *fabrication processes* need to be considered just as brain programming technologies have played a role in Evolutionary Robotics so far. Indeed, the control over both design *and* fabrication, so to speak, is an integral characteristic of all life forms⁹.

The book first covers evolutionary and neural techniques, overviewing basics in genetic algorithms, genetic programming, selection techniques and schema theory, and their relation to biological evolution. It then elaborates on neural network models, Hebbian learning, supervised, unsupervised and reinforcement learning. After this preliminary background, the authors dive into practical issues of “How to evolve robots”. Important aspects regarding experiment planning, mechanical robustness, energy supply and analysis of the results, are discussed in detail and exemplified using the Khepera platform¹⁰ (see picture). The authors offer invaluable insight and experience that, as any experimenter knows, can make or break a research project.



The Khepera Robot

Photo Alain Herzog

The next part of the book describes a set of experiments, ranging from “simple” reactive behaviors for navigation and light seeking, to problems that would traditionally require internal representation but can still be solved using reactive architectures, such as tasks that need to deal with ambiguity and aliasing, and ultimately problems that require internal states and learning. Anyone that has carried out physical robotic learning experimentation can appreciate the excruciating efforts involved in achieving each of these results.

The first set of experiments describe navigational tasks such as obstacle avoidance, from reactive approaches to more complex visually-guided setups, which require human-guided incremental evolution to break down the problem to co-development of sensory morphologies and control strategies, and then their binding. The authors point out that in some cases the evolved solution outperformed the hand-design solution by capitalizing on interactions between machine and environment that could not be captured by a model based approach (I assume because of their subtlety). These interactions also lead the authors to the view of the robots as *continuously adaptive* systems, rather than *optimal* solutions. One of the questions then is whether exploitation of these interactions leads just to refinement of strategies that could also be achieved grossly in simulation (and refined in reality in the last stages), or to new strategies that could not be found in simulation at all, and hence merit evolution in reality. The authors view the transition from simulation to reality as a form of incremental evolution, dependent primarily on the fidelity of noise in the simulation. Invariably, however, this dilemma will become more prominent as we will move to more complex tasks in the future, and is again at the crux of scaling evolutionary robotics techniques. As tasks become more complex, the discrepancies between simulation and reality will increase too, but so will the cost and duration of real-time physical experimentation. Evolution for variable or multiple environments will amplify this difficulty furthermore. Hence ultimately, there is no escape from scaling the simulation/experimentation process itself as well, by decomposing and abstracting it.

When moving to more complex tasks, limits of reactive architectures is exposed, but still the authors show how sensory-motor coordination is exploited by evolution to transform hard problems, including ones involving serial subtasks, so that they can be addressed

without internal state. Later on, the book discusses very complex tasks, namely garbage collection (locating barrels, picking them up and disposing of them over a wall), and recharging (locating and returning to recharging zones at certain charge battery levels). These tasks are used to study conditions permitting development of complex tasks. They show that emergent modular structures (where the existence modules, their architecture and arbitration mechanism was predetermined, but the combination weights evolved) allowed the decomposition of the global behavior into basic behaviors to emerge spontaneously. Moreover, the decomposition reached did not correspond to a distal decomposition an external designer would naturally expect, and outperformed other manually designed decompositions. The authors conclude that emergent modularity accelerates evolution and allows the development of more robust controllers, and that there is no correspondence between evolved modules and distal behaviors. The recharging experiments similarly show how the agents evolve an internal representation that has a structure different than a typical design based on distal modeling. These two experiments lead the authors through a discussion of the prospects of emergence of complex behaviors, and to address other mechanisms of adaptation such as learning plasticity, co-evolution and development, in the following chapters.

First, the authors investigate how learning can change the adaptive power of evolution by guiding it¹¹, even though characters acquired through lifetime learning are not inherited. They show how the combination of evolution and learning allows solving of problems that evolution alone cannot solve, even when the learning task is different from the task for which individuals are selected. Also, by learning, some aspects of the solution can be extracted from the environment rather than being encoded in the genotype. Next, the authors discuss competitive co-evolution, with the advantages and challenges it introduces. Specifically, co-evolution allows studying evolutionary process in a changing environment, allowing a much larger variety of behaviors to be examined. Furthermore, it provides a means for automatically producing incremental fitness criteria, with the hope of creating an “arms race” but equally often leading to cyclic dominance¹². Finally, they discuss developmental models such as growing and cellular encodings as a means to enhance the evolvability of architectures. Most of the coding schemes described use complex rules and building blocks supplied by the designer for a specific problem, and the experimental evidence is not yet strong enough to show the viability of this method compared to more directly specified architectures.

In the last two chapters of the book, the Khepera robot is relinquished in favor of more complex hardware. The first chapter discusses evolution of gait in various pedaling morphologies, like six-legged machines. Evolution of locomotion on these platforms is more difficult, and requires a staged (incremental) adaptation process where the designer decomposes the required functionality and each stage is learnt separately. Control is achieved using oscillating neurons as pattern generators. The book concludes with an overview of recent work in evolvable hardware and body-brain co-evolution, but argues that the technology for online implementation is still restricted.

The authors recognize the fact the ideas of evolutionary robotics can be used to assist in engineering design, but overlook the converse – what evolutionary robotics can *gain* from understanding what engineers have long considered as the *principles of design*^{13,14}. Evolutionary robotics is really a form of *design automation*, albeit arbitrarily limited to

robotics. But evolutionary engineering is divorced from the elegant parsimony of human engineering, praising itself that elegance merely reflects a human weakness for simplicity, as Kevin Kelly wrote¹⁵ after interviewing a dozen or so Alife researchers:

When Koza began to inspect the insides of his highly evolved prizes, he had the same shock that Sims and Ray did: The solutions were a mess! Evolution went the long way around. Or it burrowed through the problem by some circuitous loophole of logic. Evolution was chock-full of redundancy. It was inelegant. Rather than remove an erroneous section, evolution would just add a countercorrecting section, or reroute the main event around the bad sector. The final formula had the appearance of being some miraculous Rube Goldberg collection of items that by some happy accident worked. And that's exactly what it was, of course. But...parsimony is highly overrated...It is...a mere "human aesthetic". (Out Of Control, p. 337-8)

But Engineers have discovered long ago, that solutions that are a mess, *don't scale*. At first sight, engineering design appears a top-down process, as unrelated to evolution as can be. But the discovery of engineering knowledge – *the evolution of technology* – is, like the evolution of biology, a long self-organizing bottom-up process involving a historical co-evolution of design challenges and design solutions. However, unlike nature, the history of evolution of technology is much more accessible¹⁶, and its products and processes much easier to reverse-engineer¹⁷. Many of the principles we are now discovering in evolutionary computation as keys to scaling to higher complexity^{18,19,20}, resonate with accepted properties of good engineering design: *Functional modularity*, *structural regularity*, and *intermediate stability*. These principles ultimately coalesce into a single particularly useful functional and structural architecture: *Hierarchy*²¹. When working in simulation, then *functional abstraction* plays a critical role as well, because it allows the simulation to scale too by hiding lower-level details. While the authors mention modularity a little, and regularity enters indirectly through developmental processes, I think that this grand theme, which is at the crux of scaling, is overlooked.

Nolfi and Floreano make a superb job of presenting the state of the art of Evolutionary Robotics, weaving their own experimental results with current theoretical ideas. It is indeed a notable achievement, and will certainly be considered a hallmark in this leading edge field.

¹ Nolfi S., Floreano D. (2000), *Evolutionary Robotics – The Biology, Intelligence, and Technology of Self-Organizing Machines*, MIT Press, Cambridge

² Mataric, M. and Cliff, D. (1996). Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems*, Special Issue on Evolutional Robotics, 19(1):67–83

³ Ficici, S. G., Watson, R. A. and Pollack, J. B. (1999). Embodied Evolution: A Response to Challenges in Evolutionary Robotics. Eighth European Workshop on Learning Robots. Wyatt J.L., Demiris J., eds., 14-22

⁴ Moravec H., (1999) "Rise of the Robots", *Scientific American*, Vol. 281 No. 12

⁵ Brooks R.A. (1986) A robust layered control system for a mobile robot, *IEEE Journal of robotics and Automation* 2:14–23

-
- ⁶ Barto A.G., Bradtke S.J., Singh S.P. (1995) Learning to act using real-time dynamic programming, *Artificial Intelligence*, 6:105-122
- ⁷ Wagner G.P., Altenberg L. (1996), Complex adaptation and the evolution of evolvability, *Evolution* 50:967-976
- ⁸ Husbands P., Meyer J.A., Eds., (1998) *Evolutionary Robotics : First European Workshop, Evorobot'98*, Springer Verlag (LNCS Series)
- ⁹ H. Lipson and J. B. Pollack (2000), "Automatic design and Manufacture of Robotic Lifeforms", *Nature* 406, pp. 974-978
- ¹⁰ LAMI, Ecole Polytechnique Fédérale de Lausanne – EPFL, and K-Team, Inc.
- ¹¹ Hinton G.E., Nowlan S.. How learning can guide evolution. *Complex Systems*, 1:495-502
- ¹² Ficici, S. G. and Pollack, J. B. (1998). Challenges in Coevolutionary Learning: Arms-Race Dynamics, Open-Endedness, and Mediocre Stable States. Proceedings of the Sixth International Conference on Artificial Life. Adami, Belew, Kitano, Talor, eds. Cambridge: MIT Press, 1998
- ¹³ Ulrich K.T. and Eppinger S.D., (2000), *Product Design and Development*, 2nd Ed., McGraw-Hill, New York
- ¹⁴ Suh N.P. (1990) *The Principles of Design*, Oxford University Press, Oxford UK
- ¹⁵ Kelly K., (1994), *Out of Control: The New Biology of Machines, Social Systems and the Economic World*, Perseus Press.
- ¹⁶ Basalla G., (1989) *The Evolution of Technology*, Cambridge University Press
- ¹⁷ French M.J., (1995), *Invention and Evolution: Design in Nature and Engineering*, 2nd Ed., Cambridge University Press
- ¹⁸ Angeline, P. J. and Pollack, J. B (1992). Evolutionary Induction of subroutines, 14th Annual Cognitive Science Conference, 236-241
- ¹⁹ Hornby, G. S., Lipson, H. Pollack, J. B. (2001). Evolution of Generative Design Systems for Modular Physical Robots. *IEEE International Conference on Robotics and Automation*, In press
- ²⁰ Watson, R.A. and Pollack, J.B. (2000). Symbiotic Combination as an Alternative to Sexual Recombination in Genetic Algorithms. Proceedings of Parallel Problem Solving from Nature (PPSNVI), Schoenauer M., Deb K., Rudolph G., Yao X., Lutton E., Merelo J.J., Schwefel H.P. (Eds.), 2000. Springer Verlag, LNCS 1917
- ²¹ Alexander, C.W, (1970), *Notes on the Synthesis of Form*, Harvard University Press