

Introduction to the Evolution of Physical Systems Special Issue

John Rieffel^{*,**}
Union College

Jean-Baptiste Mouret[†]
Inria

Nicolas Bredeche[‡]
Sorbonne Universités, UPMC

Evert Haasdijk[§]
Vrije Universiteit Amsterdam

I Introduction

We are delighted to introduce this Special Issue on the Evolution of Physical Systems, the culmination of a series of workshops organized around the topic that began with ALIFE XIII (2012) and ran through ECAL 2013, ALIFE XIV (2014), and ECAL 2015.

Inspired by our mutual interests, we coined the term *evolution of physical systems* (EPS) to describe evolutionary approaches that occur in real-world physical substrates rather than in simulation. We deliberately chose this term to be broad enough to encompass both parallel embodied evolution [19], in which evolution is distributed across a population of robots, and more classical evolutionary robotics work, where evaluation is serialized on a single robot, as by Floreano and Mondada [5].

2 A Brief Survey of the Evolution of Physical Systems

The evolution of physical systems has its roots in the embodiment philosophy of Rodney Brooks, who famously said “the world is its own best model” [3]. Brooks’ emphatic critique of the symbol system hypothesis argued that relatively simple systems, developed and grounded in suitably complex real-world environments, can lead to the emergence of complex behaviors. Contemporary approaches to the EPS aim at automating Brooks’ vision by exploiting evolutionary algorithms on systems that “live” and “evolve” in the real world, that is, whose behavior (and possibly their form) are as grounded as possible in the environment.

Embodied evolutionary algorithms first gained prominence with the “Sussex approach” of Harvey et al. [7] in the 1990s. The Sussex gantry-based robot [4, 6], driven by a neural network, was capable of robust behaviors in a noisy real-world environment. Their motivation for using the real-world rather than a simulated environment stemmed largely from the fact that in that era of limited computational power, it was faster and more efficient to test networks in situ than it was to develop and run a simulator that realistically modeled sensor noise. This notion, that it is sometimes more practical (and computationally efficient) to dispense with simulation entirely and evolve behaviors directly in the real world, persists at the heart of the EPS.

* Contact author.

** Union College, Schenectady, New York, U.S.A. E-mail: rieffel@union.edu

† Inria, CNRS, Université de Lorraine, UMR 7503, LORIA, Nancy, France. E-mail: jean-baptiste.mouret@inria.fr

‡ Sorbonne Universités, UPMC Univ Paris 06, CNRS, ISIR, Paris, France. E-mail: nicolas.bredeche@upmc.fr

§ Vrije Universiteit Amsterdam, Amsterdam, The Netherlands. E-mail: e.haasdijk@vu.nl

Two other results of the Sussex group are particularly noteworthy. First is Adrian Thompson's *silicon evolution* of a field-programmable gate array (FPGA) [17, 18]. The chosen task was to evolve a circuit capable of discriminating between 1- and 10-kHz signals without the use of a clock. By evolving in silicon, rather than in simulation, Thompson was able to generate nearly inscrutable solutions that exploited the analogue (and non-simulable) nature of FPGAs. This idea of using embodiment to arrive at “novel surprise” solutions is a further guiding principle of the EPS. The second is Jakobi's essential work on the *reality gap* [8], which highlighted the difficulty in transferring results evolved in simulation into the real world. Jakobi notes that the best results emerge when “the noise levels of the simulation have similar amplitudes to those observed in reality,” and points out that these approaches become less feasible as environments and sensors become more complex.

The practice of evolving outside of simulators grew from there. Floreano and Mondada [5] were able to evolve robust wall-following behavior in a tethered Khepera robot through embodiment. Watson et al. [19] took embodiment a step further by simultaneously embedding an entire population of evolving *TipperBots* into a shared arena—and demonstrated how the behaviors that emerged through this process were qualitatively different, and more effective, than those evolved in a simulated environment (their use of an electrified floor as an alternative to tethers was particularly innovative). Zykov et al. evolved dynamic open-loop gaits on a large hexapod robot, remarking upon the amount of labor required to reset the large robot between physical trials [21].

Several of our own research areas have been in this realm of problems that are easier to test physically than they are to simulate realistically. In their work on “evolutionary fabrication” [16, 12], Rieffel and colleagues ran an embodied evolutionary algorithm on a 3D printer, learning to produce a variety of shapes out of viscous extruded materials such as silicone. These efforts led to novel shapes (and novel means of producing the shapes) that no human would have designed. Similarly, and more recently, Rieffel's work on the evolution of vibration-based gaits for tensegrity robots [15] relied purely on physical trials, and led not only to fast and dynamically complex gaits, but also to a state-machine-based controller for a tensegrity robot able to chase targets [9].

While not using a simulator at all has many benefits, evolutionary robotics can take advantage of both simulators and real experiments to accelerate the evolutionary process while keeping a direct link with the real world. A first idea is to start evolution in simulation and then finish it on the real robot(s); this is, for instance, what Lipson and Pollack did to evolve the morphology and controllers of 3D-printed robots [13]. However, this approach means that designs for the real world are variants of those evolved in simulation, that is, robots are more grounded in the simulation than in the real world. A second, classic idea is to use data collected from tests on the real robot to improve the simulator, ideally to make the reality gap so small that it becomes unnoticeable [2, 20]. Nevertheless, to do so, simulators need to be generic enough to model almost anything that could happen, which, in turn, requires acquiring a considerable amount of data in the real world. More recently, a weaker version of this idea has emerged: Instead of trying to correct the simulator, one could use a few tests in the real world to learn to avoid the behaviors that exploit phenomena that are inaccurately simulated. Following this concept, Koos et al. managed to evolve controllers for six-legged robots and Khepera-like robots with only a dozen tests in the real world [10, 11, 14].

Despite numerous advances in both robotics and technology, Rodney Brooks' prescient words about the crucial role of embodiment in robotics remain as true as ever, and the EPS is as exciting as ever.

3 The Articles

We the editors are delighted to present the articles that were selected to be published in this issue. Collectively they provide an inspiring snapshot of the state of the evolution of physical systems.

3.1 Automating the Incremental Evolution of Controllers for Physical Robots

Jacobsen and colleagues describe an experimental setup that removes the burden of manually setting the environment (e.g., moving the robot(s) and object(s) back to their initial positions and orientations) from the human supervisor. This is particularly relevant in evolutionary robotics, where a large number of trials are generally required, all starting from similar initial conditions. The authors present a proof of concept, tested with a real robotics setup in an arena of 1 by 0.5 m, which makes use of an industrial robot arm, a camera, and fiducial markers that can be stuck on top of robots and objects alike. These fiducial markers are used both for monitoring the experiments, by tracking positions and orientations, and as visual anchors for the robot arm to grab so as to move back robots and objects to their initial positions. Beyond evolutionary robotics, the proposed setup may also be used for other robotics setups involving multiple trials (e.g., reinforcement learning and embodied evolution), whether for resetting the experimental setup, as demonstrated in this article, or for other, more dynamical interventions (e.g., by conditionally moving objects around or transporting robots during the course of an experiment). To ensure that their system can be adopted and easily deployed by others, the authors have included the source code and a detailed description of the technical aspects of their setup.

3.2 Abstraction, Sensory-Motor Coordination, and the Reality Gap in Evolutionary Robotics

In their contribution, Schepers and de Croon describe how choosing an appropriate level of abstraction can improve the transfer of evolved solutions across the reality gap. The task domain they have chosen is a decentralized formation flying task for a group of three small airborne vehicles controlled by neural networks. They compare the performance of a low-level controller that has precise control over vehicle dynamics (pitch rate, pitch angle, etc.) against a higher level controller that uses a more abstract velocity setpoint. After evolution (in simulation), the more abstract-level controller was better able to survive transfer into the real world. Schepers and de Croon suggest that the higher-level controller is better able to perform a kind of environmental exploration by inherently modeling the behavior of the other vehicles in the swarm. As they point out, an advantage of more abstracted controllers is that it reduces the necessity of high-fidelity simulation—an important advantage for evolutionary robotics tasks. We were impressed by the insights of this article in regard to the tradeoffs between simulator fidelity and real-world behavior. The flight videos on YouTube are also quite impressive.

3.3 Evolutionary Developmental Robotics: Improving Morphology and Control of Physical Robots

The work of Vujovic and colleagues is very much in the spirit of this special issue. They introduce a novel robot-building robot composed of a hot-melt extruder and an articulated arm, which acts as a sort of 3D printer, capable of fabricating the legs of simple modular robots. An evolutionary developmental approach is used to “grow” the designs of the legs, and once printed, the locomotive abilities of the printed robots are automatically evaluated in situ. This type of physically embodied integrated design-fabricate-test loop, in which one robot evolves and assembles other robots, is very much in the vein of Lipson and Pollack’s GOLEM robots [13], moving the field one step closer to 3D-printed robots that can walk themselves off the printer.

3.4 On Design Mining: Coevolution and Surrogate Models

Preen and Bull present an interesting theoretical analysis followed up with an actual physical case study. Evaluation of physical artefacts is commonly a time-consuming aspect of the EPS, so a thorough analysis of methods to minimize the number of required evaluations can greatly boost applicability. Preen and Bull employ the abstract NKCS model of coevolution to achieve this: It enables informed sampling of candidate solutions in a coevolutionary scenario, where various components of a physical system interact in combination with a neural-net-based surrogate model.

A successful case study to develop a heterogeneous set of vertical wind turbines for wind farms shows the applicability and relevance of their method to the evolution of designs for physical artefacts without relying on simulation.

3.5 Real-World Evolution of Robot Morphologies: A Proof of Concept

For obvious practical reasons, most of the work in the EPS has been focused on the evolution of controllers for existing robots, that is, for robots with a fixed morphology. Nevertheless, most of those who are interested in the field dream of a much more ambitious vision: robots whose morphology and brains can evolve on their own, much in the same ways as animals evolve.

In their contribution, Jelisavcic et al. investigate how such a vision could be achieved with a population of robots and 3D printing technology. Using the RoboGen system of 3D-printed modules [1], they first designed and 3D-printed two robots. In a second step, they put these two robots in an arena so that they could exchange their genome to create the genotype of a third robot, which they also 3D-printed. This first step of an evolutionary experiment is an interesting stepping stone to identify the challenges of setting up experiments in which a population of robots would evolve autonomously.

4 Looking Forward

Overall, these articles provide valuable contributions to the field of the evolution of physical systems, demonstrating that many of the current challenges of the field are more technological than scientific: To design machines that evolve in the physical world, we first need machines that can self-reproduce or, at least, autonomous machines that can in turn help other machines to evolve. And until high-fidelity computationally efficient simulators are developed, the role of physical embodiment in evolutionary robotics remains essential.

References

1. Auerbach, J., Aydin, D., Maesani, A., Kornatowski, P., Cieslewski, T., Heitz, G., Fernando, P., Loshchilov, I., Daler, L., & Floreano, D. (2014). Robogen: Robot generation through artificial evolution. In H. Sayama, J. Rieffel, S. Risi, R. Doursat, & H. Lipson (Eds.), *Artificial life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, EPFL-CONF-200995 (pp. 136–137). Cambridge, MA: MIT Press.
2. Bongard, J., Zykov, V., & Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science*, 314(5802), 1118–1121.
3. Brooks, R. A. (1990). Elephants don't play chess. *Robotics and Autonomous Systems*, 6(1–2), 3–15.
4. Cliff, D., Husbands, P., & Harvey, I. (1993). Evolving visually guided robots. In J.-A. Meyer, A. L. Roitblat, & S. W. Wilson (Eds.), *From animals to animats 2. Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 374–383). Cambridge, MA: MIT Press.
5. Floreano, D., & Mondada, F. (1994). Automatic creation of an autonomous agent: Genetic evolution of a neural network driven robot. In D. Cliff, P. Husbands, J.-A. Meyer, & S. W. Wilson (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, LIS-CONF-1994-003 (pp. 421–430). Cambridge, MA: MIT Press.
6. Harvey, I., Husbands, P., & Cliff, D. (1994). Seeing the light: Artificial evolution, real vision. In D. Cliff, P. Husbands, J.-A. Meyer, & S. W. Wilson (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior* (pp. 392–401). Cambridge, MA: MIT Press.
7. Harvey, I., Husbands, P., Cliff, D., Thompson, A., & Jakobi, N. (1997). Evolutionary robotics: The Sussex approach. *Robotics and Autonomous Systems*, 20(2–4), 205–224.
8. Jakobi, N., Husbands, P., & Harvey, I. (1995). Noise and the reality gap: The use of simulation in evolutionary robotics. In F. Morán, A. Moreno, J. J. Merelo, & P. Chacón (Eds.), *Advances in artificial life: Third European Conference on Artificial Life* (pp. 704–720). Berlin, Heidelberg: Springer.

9. Khazanov, M., Humphreys, B., Keat, W. D., & Rieffel, J. (2013). Exploiting dynamical complexity in a physical tensegrity robot to achieve locomotion. In P. Liò, O. Miglino, G. Nicosia, S. Nolfi, & M. Pavone (Eds.), *Advances in artificial life, ECAL 2013* (pp. 965–972). Cambridge, MA: MIT Press.
10. Koos, S., Cully, A., & Mouret, J.-B. (2013). Fast damage recovery in robotics with the T-resilience algorithm. *International Journal of Robotics Research (IJRR)*, 32(14), 1700–1723.
11. Koos, S., Mouret, J.-B., & Doncieux, S. (2013). The transferability approach: Crossing the reality gap in evolutionary robotics. *IEEE Transactions on Evolutionary Computation (TEC)*, 17(1), 122–145.
12. Kuehn, T., & Rieffel, J. (2012). Automatically designing and printing 3-D objects with Evofab 0.2. In C. Adami, D. M. Bryson, C. Ofria, & R. T. Pennock (Eds.), *Artificial life 13* (pp. 372–378). Cambridge, MA: MIT Press.
13. Lipson, H., & Pollack, J. B. (2000). Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799), 974–978.
14. Mouret, J.-B., Koos, S., & Doncieux, S. (2013). Crossing the reality gap: A short introduction to the transferability approach. arXiv preprint arXiv:1307.1870.
15. Rieffel, J., & Mouret, J.-B. (2017). Soft tensegrity robots. arXiv preprint arXiv:1702.03258.
16. Rieffel, J., & Sayles, D. (2010). Evofab: A fully embodied evolutionary fabricator. In G. Tempesti, A. M. Tyrrell, & J. F. Miller (Eds.), *Proceedings of the 9th International Conference on Evolvable Systems: From Biology to Hardware* (pp. 372–380). Berlin, Heidelberg: Springer.
17. Thompson, A. (1996). An evolved circuit, intrinsic in silicon, entwined with physics. In T. Higuchi, M. Iwata, & W. Liu (Eds.), *Proceedings of the First International Conference on Evolvable Systems: From Biology to Hardware* (pp. 390–405). London: Springer.
18. Thompson, A. (1996). Silicon evolution. In J. R. Koza, D. E. Goldberg, D. B. Fogel, & R. L. Riolo (Eds.), *Proceedings of the First Annual Conference on Genetic Programming* (pp. 444–452). Cambridge, MA: MIT Press.
19. Watson, R. A., Ficici, S. G., & Pollack, J. B. (2002). Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, 39(1), 1–18.
20. Zagal, J. C., & Ruiz-Del-Solar, J. (2007). Combining simulation and reality in evolutionary robotics. *Journal of Intelligent & Robotic Systems*, 50(1), 19–39.
21. Zykov, V., Bongard, J., & Lipson, H. (2004). Evolving dynamic gaits on a physical robot. In K. Deb, R. Poli, W. Banzhaf, H.-G. Beyer, E. Burke, P. Darwen, D. Dasgupta, D. Floreano, J. Foster, M. Harman, O. Holland, P. L. Lanzi, L. Spector, A. G. B. Tettamanzi, D. Thierens, & A. Tyrrell (Eds.), *Proceedings of Genetic and Evolutionary Computation Conference (GECCO)*, late-breaking paper, vol. 4. Berlin, Heidelberg: Springer.