Embodied evolution in a morphologically heterogeneous population of robots

Branen Salmon

May 19, 2003

Abstract

Holistic evolution (the concurrent evolution of control structure and morphology) and embodied evolution (the distribution of evolutionary function into asynchronous, autonomous robots operating in a task environment offer a great deal of possibility to the field of evolutionary robotics. Both methodologies have principlistic and practical ramifications, but the most radical—and important—may lie in their confluence.

This paper details the motivations behind an experimental work in progress.

1 Introduction

Over the past fifteen years, evolution has become a common part of the AI-roboticist's toolkit. Its searches across large-dimensional parameter spaces are automated and efficient, its variables (fitness functions) are powerful and elegant, and its domain generality is broad. These attributes have made it a highly popular and thoroughly researched tool in intelligent robotics. But despite this widespread adoption and refinement, evolutionary robotics remains underdeveloped in two critical areas: *embodied*[1] evolution and *holistic*[2] evolution. The importance of morphology in robotics has long been recognized [5, 8], but the development of holistic evolution has been significantly retarded by the difficulty of building morphologically dynamic hardware. Embodied evolution, on the other hand, is an almost completely unexplored territory, despite of—or perhaps because of—its potentially radical differences from the dominant paradigm of elitist[3] evolution.

---

[1] Note that I do not restrict "embodied" to mean "physically embodied." See [12] and [21].

[2] For speaking about the concurrent evolution of controller and morphology, I propose the term 'holistic.' This choice is a departure from the terms 'concurrent evolution' and 'co-evolution' used in previous literature, and is motivated by a desire to avoid confusion with the large body of work on *inter-special* "co-evolution."

[3] By "elitist evolution," I refer to the winner-take-all and winners-take-all Darwinian tournaments in which the "most fit"individuals of one discrete generation are replicated to form the test population of the next.

In this paper, I hope to argue the importance of these two frontiers. Furthermore, I will outline my belief in a possibility of their productive confluence in a holistic *and* embodied evolution.

2 Holistic evolution

Since the early 1990s, the literature has been filled with calls to the importance of holistic evolution. Since the topic was originally broached by Brooks [3,4], a number of roboticists have argued in principle that attention to morphological development is requisite to the continued vitality of intelligent robotics. More recently, researchers have demonstrated a number of practical benefits to these principles, both in simulated experiments and in actual hardware. The benefits offered are immense in both principle and practice.

2.1 A blow against dualism

The most important contribution that holistic evolution offers to the field may also be its most oblique: a fundamentally different way of thinking about intelligent systems. Brooks suggested in [3] that the path to intelligent robotics was marked by four points: *situatedness* in the dirty, continuous world (instead of in a symbol space); *embodiment* in a world-interactive form (such that robotic experience is the result of a world-robot dynamic); *intelligence* (or more precisely, behavior observed to be intelligent)[4]; and *emergence*—an understanding that the overall behavior of a robot is the distributed result of complex interactions between its control structures, its body, and its environment. This last point is key to the importance of holistic evolution. A robot's observed behavior is a result of the interaction between its actuators, its body, and the world. The activation of its actuators, in turn, is the result of its control structure. The state of its control structure depends on the activation of its sensors. Finally, a robot's sensory activation depends on the interaction between its body and the environment—the environment which is, itself, being modified by the robot's body. As Cliff, Husbands, and colleagues put it most eloquently:

> We view the networks we evolve as continuous dynamical systems, rather than as computational devices transforming between representations: inputs to the system might *perturb* the trajectory of the network in state space, so it enters a different state which might be interpreted by an external observer as a new behavior. [...] Separating morphology from control is a measure which is difficult to justify from an evolutionary perspective, and potentially misleading. [8]

furthermore,

---

[4] It is a fine point, but one worth making: the evaluation of intelligence can only be made through the evaluation of behavior, which in turn can only be made through observation.

> We regard the proper study of intelligence as an investigation into the interactions between autonomous agents and their environments. [...] We believe that ... the control systems needed [for adaptive, autonomous real-world robotics] will be of the complex dynamical systems variety, and these are inherently extremely difficult to design by traditional means. [...] Suitable sensor and actuator properties (including morphologies) are inextricably bound to the most appropriate "internal" dynamics of the control system and vice versa. [18]

It is this set of rich, chaotic relationships between controller, morphology, and environment that Brooks, Cliff, Husbands, and others have recognized as a watershed opportunity for intelligent robotics. By thinking of controller, morphology, and environment not as a discrete trinity, but rather as aspects of a larger dynamical system, we make it possible to exploit the structure of the *entire system* in our pursuit of intelligent behvaior.[5]

Yet, as Husbands et al. point out, even small autonomous systems (eg. insects) exhibit remarkably complex dynamics. They and others have concluded [15] that the design of robots that fully exploit the emergent interactions of the robot-environment system is simply out of the reach of manual engineering. The solution is holistic evolution.[6]

2.2 A provably productive tool

Unfortunately, the nagging difficultly of hardware implementation has thus far kept holistic evolution out of the robotics mainstream. Nevertheless, there is experimental evidence to support its theoretical promise. In simulation, Sims has evolved directed graphs that describe the neuromorphologies of creatures that compete for control of a block. [22] The strategies evolved include creatures that block opponents' movements, that fully enclose the block, and that topple over it—all of which suggest a close coupling of physical structure and control. Hornby and Pollack have evolved simulated creatures capable of 'naturalistic' locomotion [16]. And the DEMO Lab at Brandeis University has produced a system whereby locomoting robots are evolved in simulation and automatically fabricated in rapid-prototyping equipment. [20] (In fact, everything is automated except for the quick snapping-in-place of linear actuators into the fabricated robots.)

But perhaps most impressive of all have been the results from inquiries into the efficiencies—morphological and computational—of holistic evolution. In [6], Bugajska and Schultz found preliminary evidence that holistic evolution could optimize the quantity of sensors and the parameters for individual sensors on an autonomous micro air vehicle when fitness was partially determined by sensor count. More strikingly, Balakrishnn and Honavar found that in their

---

[5] Indeed, the potentials for robots that *by virtue of their design*, exploit structure latent in their environments to produce complex behavior is a central interest to the project of *innateness* [10].

[6] Descartes obviously did not build robots.

holistic evolution of a box-pushing robot, the sensor count was minimized *without* any bias in the fitness function towards individuals with fewer sensors, [1] which suggests, when considered alongside their robots' "counterintuitive" control structures, an intrinsic efficiency resulting from the synergetic coevolution of control and morphology. Additional sensors did not develop because the system was likely to find good solutions more quickly in simple spaces (few sensors) than in complex spaces (many sensors). This synergy is further echoed in work of Bongard and Paul, who found that when both systems were given favorable starting conditions, the holistic evolution of a bipedal robot could converge *more quickly* than the control-only evolution in a similar body—despite the exponentially larger search space of the former over the latter.[7] [2] So not only does holistic evolution offer the power of complex dynamical systems and the optimal coupling of brain, body, and environment; under some conditions, it offers super-efficient searches through its evolutional space.

3 Embodied evolution

Whereas holistic evolution got its start at the beginning of the 1990s, embodied evolution found its beginning at their end. The approach in which I take interest[8] was first articulated by Ficici, Watson, and Pollack of the DEMO Lab in 1999 [12, 23]. To date, there have been only a miniscule number of papers to even reference the approach from outside the DEMO Lab; nevertheless, it shows great promise.

3.1 Embodied evolution: what it is not

Before detailing just what embodied evolution *is*, it may be instructive to detail what it *is not*. Embodied evolution is not mainstream. Virtually all of evolutionary robotics employs one form or another of *elitist evolution*. Elitist evolution is characterized by a sequence of competitions between individuals in which the winners[9] of each step in the sequence—that is, of each generation—are replicated to form the next generation's population. (Alternatively, one could say that all but the 'best' are wiped out after each competition—thus, the term 'elitist.') There are several variations on the theme: most researchers employ a Darwinian selection, though a few have experimented with Larmarckian transmission of life learning. Some experiments have interbred pairs (or multiples) of the most fit individuals; others have reproduced them asexually, with only mutation thrown into the mix. A handful of researchers have even toyed with a 'kindler, gentler' evolution in which a less-fit individual will occasionally be spared the gauntlet. Elitist evolution comes in both single– and multiple-species varieties[10], and though the bulk of elitist competitions occur serially, techniques

---

[7] They attribute this to "extradimensional bypasses" linking suboptimal adaptive ridges through control– and morphology-space.

[8] This approach, in fact, happens to be the only approach of its kind in the literature.

[9] Generally, the winners are the individuals with the highest fitness evaluation.

[10] See [7, 9] for examples.

exist for their spatial distribution [17]. There even exist techniques for performing *embodied trials* in an elitist fashion, yet despite all of this variety, elitist evolution is still bound to some degree of centralized evaluation function; there is still some Nietzschean hand of god that reaches down and separates the wheat from the ostensible chaff.[11]

3.2 Embodied evolution: what it is

In the words of those who created it, embodied evolution is "evolution taking place within a population of real robots where evaluation, selection, and reproduction are carried out by and between the robots in a distributed, asynchronous, and autonomous manner [...] in the task environment." [12] In this sense, "embodied evolution" is not the evolution of embodied robots, but rather the *embodiment of evolution itself* into the robots themselves. The loss of a central evaluator means that evolution can no longer proceed in an elitist fashion; there is no more Nietzschean hand of god to select the cream of the crop. The development of the genome is up to the robots themselves. (I will broaden their definition slightly by claiming that 'real robots' need not be *physically* embodied, so long as they meet the requirements set forth in [21] as being an 'embodied system.'[12])

Given the current state of the art, it is very difficult to design robots capable of autonomously constructing their own offspring from the carcasses of robots that have ceased to function. Not surprisingly, then, the first implementation of embodied evolution [23] took a rather radical departure from the 'life and death' approach to cross-generational perpetuation; indeed, it jettisoned the notion of discrete generations from the outset. In its place, Watson et al. placed their *Probablistic Gene Transfer Algorithm*, which draws heavily upon the biological metaphor of microbial recombinance.[13] In the PGTA, robots repeatedly transmit their genes (with mutations) to other robots in the vicinity at a rate determined by the sending robot's fitness. Robots that receive gene transmissions may ignore the transmission or incorporate it into their genome; the likelyhood that the latter will happen increases as fitness goes down.[14] By this mechanism, genes from fit individuals will tend to overwrite genes from unfit individuals... without the need for any inter-robot state to be maintained.

---

[11] Note that this language is not meant to imply that elitist evolution is somehow 'bad methodology,' but merely to make the text a little more flavorful.

[12] In this sense, robots may be embodied in software simulation or even in software environments completely incommensurate with our own physical world [11], so long as said embodiment is of a sufficiently complex dynamic to be interesting.

[13] Recombinance is the primary mechanism by which asexual microbes exchange genes amongst individuals. Parts of a species's genome exist as genes arranged in relatively atomic packets called *plasmids*. These plasmids can be swapped amongst microbes in proximity of one another without disrupting the rest of the microbes' genes. Plasmids are considered to be a major factor in the speed with which asexual microbes are able to adapt to new environments—particularly to the presence of antibiotics.

[14] Interestingly, this ability to reject foreign genes appears to be a requirement for the eventual stabilization of a population's genome. [23]

3.3 Embodied evolution: plays well with others

Embodied evolution renders moot a number of issues that are insurmountable obstacles in elitist evolution. Most notable is that of speed. In elitist approaches, individuals are must be compared at every *synchronous* generation. Whether this comparator is global or distributed (as in [17]), it still must involve necessarily the serial evaluation of multiple individuals. Though for small populations, this is not likely to be a problem, the computational complexity of these comparisons grows exponentially as the size of the population increases. In embodied evolution, on the other hand, there are no synchronous generations to require central evaluation. Furthermore, there is no need for any central functionality, as all reproductive functions are distributed, autonomous, and asynchronous. In adding an individual, one also adds everything necessary to execute that individual's evolutionary function. Thus, embodied evolution no incurs no intrinsic performance penalty as the population increases. (For this reason, and since embodied evolution takes place in a population's task environment [12], embodied evolution seems like an excellent tool for the investigation of collective, cooperative, and swarm robotics, as well.)

3.4 Embodied evolution: speed demon

Furthermore, I suspect that the performance of a population may have the potential to increase super-linearly with population size. As the number of gene transactions will increase exponentially with population size, so too will the amount of computational power being spent on this distributed, autonomously parallelized search through the evolutionary parameter space. Additionally, as genes are perpetuated on a sub-individual *'plasmid'* level, I suspect that the intermixing of genes between individuals will create an opportunity for inter-dimensional synergy, a phenomenon in which the optimization of individual gene regions will tend to retard the absolute rate of search in their regions, increasing the effective rate of search in others, and leading to a more expeditious convergence of the entire genome on an optimal configuration.[15]

3.5 Embodied evolution: pliable, yet robust

Perhaps the most exciting aspect of embodied evolution is that it is explicitly designed for adaptation in the task environment; evolution is recast from being a preparatory measure to being an integral, ongoing part of a population's existence. By virtue of this fact, a population need not ever cease to be pliable; indeed, the stabilization of such a population's genome will come not from the cessation of the evolutionary process, but from a stability which emerges from the optimal interaction of population and environment. Regardless of the health of the population, the evolutionary facilities of a population's individuals are always active. When the bulk of the population is fit, this means that genes will be

---

[15] Indeed, this is what would be suggested by the microbial analogy discussed in footnote 13.

transmitted frequently, but rarely used. However, should the environment suddenly change and the health of the population start to drop, said transmissions will begin to be accepted, and the entropy of their mutations will restart the evolutionary search for a new optimal genome. This is, of course, an incredible benefit: once a population has adapted to an environment, it will stabilize, but as soon as that adaptation is rendered suboptimal by environmental alteration, the adaptive process will resume. Stability comes as a completely emergent result of embodied evolution; nowhere need there be an explicit 'adaptation switch.'

4 Holistic, embodied evolution

On their own, holistic and embodied approaches to evolution each have a great deal to offer. However, I believe that even greater benefits could be derived from their application in tandem than could be derived from either one alone. Embodied evolution's distributed, asynchronous search and inter-spacial genetic synergy could couple with the extradimensional bypasses[16] of holistic evolution to result in extremely fast adjustments to environmental changes. Add embodied evolution's emergent stability and holistic evolution's coupling of mind, body, and environment to the mix, and the result could be a population capable of both extremely agile adaptation and robust stability in a wide variety of tasks and task environments.

Additionally, holistic, embodied evolution seems a natural platform by which to incrementally develop increasing levels of neuromorphological complexity. In [5], Brooks suggests that in the beginning stages of evolution, *"the robot should initially be operated with only some of its sensors and perhaps only some of its actuators."* Once the robot has evolved basic behaviors, he says, *"additional sensors and actuators can be made available so thta higher level behaviors can be evolved."* In [13], Funes and Pollack present a multi-stage fitness function designed to lead the evolutionary process through increasing degrees of complexity. Despite successes with each of these two approaches, they both suffer from their distal[17] notions of 'simplicity' and 'complexity.' Indeed, a fitness function which steps through what we perceive (distally) to be logically contiguous levels of complexity need not be taking the most efficient [proximal] path from initial to desired condition, and the distally-selected disabling of a robot's sensors could well wind up complicating its initial evolutionary search.

A population of robots employing holistic, embodied evolution would render the second concern irrelevant, for the selective disabling of a robot's sensors would be controlled by the same evolutionary process that drove the development of the control structures Brooks was originally hoping to optimize. The

---

[16] See footnote 7.

[17] In this paper, I use "distal" and "proximal" to refer to the human-intelligible interpretation of a system (eg. going towards the light) and the system's internal dynamics, which need not make sense to us (eg. responding to arbitrary input with arbitrary output). This usage is borrowed from [19].

distal nature of the fitness function, on the other hand, is impossible to dismiss. Nevertheless, the compounded problems of a multi-stage fitness function might at least be eased, if not entirely alleviated, by the confluence of emergent stability and punctuated equilibrium. In their paleontological theory of punctuated equilibrium, Gould and Eldgridge claimed that inter-speciation has tended to happen rapidly, and is generally bordered by extremely long periods of phenotypic stability. [14] I suspect that a similar pattern could emerge in populations of robots employing embodied evolution. The population could indeed find its own stable 'steps' on the way to a more optimal genome, periodically 'regrouping' on minima of increasingly greater scope. These periods of stabilization could serve as population-wide 'self-synchronizations,' in which the least effective genes would be flushed from the population, setting the stage for further genetic adaptations.

5 Conclusion

Both holistic evolution and embodied evoltion offer a plethora of benefits to the field of evolutionary robotics. These benefits are both principlistic and practical, conservative and radical. The most radical possibilities—but perhaps the most important—arise at the intersection of the two methodologies, and it is precisely this intersection which I seek to explore.

References

[1] K. Balakrishnan and V. Honavar. On sensor evolution in robotics. In J. R. Koza, D. E. Goldberg, D. B. Fogel, and R. L. Riolo, editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 455–460, Stanford University, CA, USA, 28–31 1996. MIT Press.

[2] J. C. Bongard and C. Paul. Making evolution an offer it can't refuse: Morphology and the extradimensional bypass. *Lecture Notes in Computer Science*, 2159:401–??, 2001.

[3] R. A. Brooks. Intelligence without reason. In J. Myopoulos and R. Reiter, editors, *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)*, pages 569–595, Sydney, Australia, 1991. Morgan Kaufmann publishers Inc.: San Mateo, CA, USA.

[4] R. A. Brooks. Intelligence without represention. Number 47 in Artificial Intelligence, pages 139–159. 1991.

[5] R. A. Brooks. Artificial life and real robots. In *European Conference on Artificial Life*, pages 3–10, 1992.

[6] M. D. Bugajska and A. C. Schultz. Co-evolution of form and function in the design of autonomous agents: Micro air vehicle project. In *Proc. Workshop*

on Evolution of Sensors GECCO 2000, pages 240–244, Las Vegas, Nevada, USA, 8 2000.

[7] S. Bullock. Co-evolutionary design: Implications for evolutionary robotics.

[8] D. Cliff, P. Husbands, and I. Harvey. Analysis of evolved sensory-motor controllers. Technical Report Cognitive Science Research Paper CSRP264, Brighton BN1 9QH, England, UK, 1992.

[9] D. Cliff and G. F. Miller. Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In *European Conference on Artificial Life*, pages 200–218, 1995.

[10] J. Elman, L. Bates, M. Johnson, A. Karmiloff-Smith, D. Parisi, and K. Plunkett. *Rethinking innateness: A connectionist perspective on development*. MIT Press, 1996.

[11] O. Etzioni. Intelligence without robots: A reply to brooks. *AI Magazine*, 14(4):7–13, 1993.

[12] S. Ficici, R. Watson, and J. Pollack. Embodied evolution: A response to challenges in evolutionary robotics, 1999.

[13] P. Funes and J. B. Pollack. Evolutionary body building: Adaptive physical designs for robots. *Artificial Life*, 4(4):337–357, 1998.

[14] S. J. Gould and N. Eldridge. Punctuated equilibria: the tempo and mode of evolution reconsidered. *Paleobiology*, (3):115–151, 1977.

[15] I. Harvey, P. Husbands, and D. Cliff. Issues in evolutionary robotics. Technical Report Cognitive Science Research Paper CSRP219, Brighton BN1 9QH, England, UK, 1992.

[16] G. S. Hornby and J. B. Pollack. Evolving L-systems to generate virtual creatures.

[17] P. Husbands. Distributed coevolutionary genetic algorithms for multi-criteria and multi-constraint optimisation. In *Evolutionary Computing, AISB Workshop*, pages 150–165, 1994.

[18] P. Husbands, I. Harvey, D. Cliff, and G. F. Miller. Artificial evolution: A new path for artificial intelligence? *Brain and Cognition*, 34:130–159, 1997.

[19] S. Nolfi. Using emergent modularity to develop control system for mobile robots, 1997.

[20] J. Pollack, H. Lipson, G. Hornby, and P. Funes. Three generations of automatically designed robots, 2002.

[21] T. Quick, K. Dautenhahn, C. L. Nehaniv, and G. Roberts. On bots and bacteria: Ontology independent embodiment. In *European Conference on Artificial Life*, pages 339–343, 1999.

[22] K. Sims. Evolving 3D morphology and behavior by competition. In *Artificial Life IV Proceedings*, pages 28–39. MIT Press, 1999.

[23] R. A. Watson, S. G. Ficici, and J. B. Pollack. Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In P. J. Angeline, Z. Michalewicz, M. Schoenauer, X. Yao, and A. Zalzala, editors, *Proceedings of the Congress on Evolutionary Computation*, volume 1, pages 335–342, Mayflower Hotel, Washington D.C., USA, 6-9 1999. IEEE Press.