

The Role of Embodiment in Open-Ended Evolution

Josh Bongard¹, Nick Cheney², Zahra Mahoor³, Joshua Powers⁴

¹ The University of Vermont, josh.bongard@uvm.edu

² The University of Wyoming, ncheney@uwyo.edu

³ The University of Vermont, zahra.mahoor@uvm.edu

⁴ The University of Vermont, joshua.powers@uvm.edu

Abstract

Open-ended evolution (OEE) has long served as a goal and guiding principle for Artificial Life researchers. However, despite the common use of embodied agents in OEE simulations, the role of embodiment in facilitating OEE has received relatively little attention. Here we introduce three recent research projects that demonstrate how the body of a learning and/or evolving agent can provide more opportunity for approaching OEE than an equivalent yet non-embodied system. In the first example, evolution of morphology, if done correctly, along with the neural control of embodied machines is shown to sustain OEE for over 6000 generations. In the second example, morphology is shown to be a way to resist catastrophic interference, a phenomenon that often frustrates OEE. The final example deals with the crowdsourcing of Artificial Life studies, one way to scale up such studies to realize OEE (assuming sustained crowd involvement), and how morphology can help or hinder this form of crowdsourcing. We hope that these examples will inspire other researchers to discover yet other ways in which embodiment may increase the probability of realizing OEE in future simulations or real world settings.

1. Introduction

Open-ended evolution (OEE) is the long-term goal for many Artificial Life researchers. If the agents being considered are embodied and are thus capable of performing multiple tasks in multiple environments, one can consider algorithms that realize less or more OEE as points embedded in a three-dimensional space (Fig. 1). The position of a given method in this space indicates how well it is able to improve the behavioral performance of agents (dimension 1), how many environments those agents can perform in (dimension 2), and how many tasks they can perform well in those environments (dimension 3).

In this paper we demonstrate how carefully considering not just the neural controller but also the morphology of the evolving agents can improve agents along all three of these dimensions. To do so, we describe three recent experiments that improve agents for one task in one environment (Sect. 2), in multiple environments (Sect. 3), and for multiple tasks (Sect. 4).

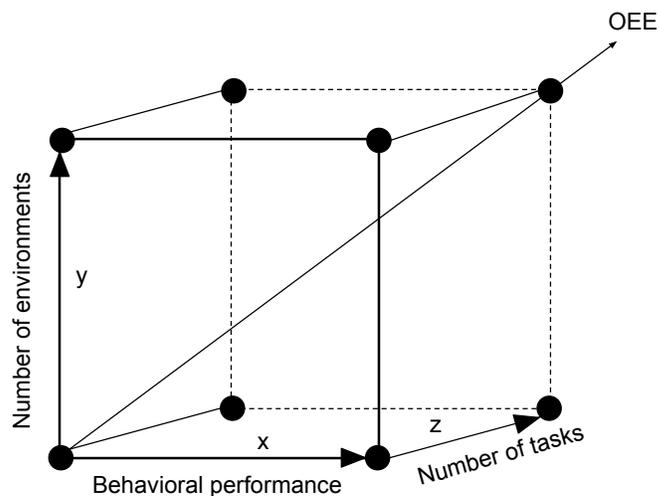


Figure 1: Visualizing OEE systems that evolve embodied agents as embedded within a three-dimensional space. A given system may be worse or better at improving an agent’s performance (x) in multiple environments (y) of different tasks (z).

2. Morphology

Arguably the first demonstration of a new path toward open-ended evolution with embodied agents was the pioneering work of Karl Sims, who showed a number of populations successfully evolving body plans and control policies against several objective functions (Sims (1994)). Since then, a number of studies have followed in which morphology and control are simultaneously evolved in an attempt to realize ever increasing complexity in the face of ever increasing task environments and objective functions (e.g. Komosinski and Rotaru-Varga (2000); Hornby and Pollack (2002); Hotz (2004); Lipson and Pollack (2000); Bongard and Pfeifer (2003); Brodbeck et al. (2015)).

However, it was not clear from those studies how much useful morphological complexity could indeed be generated, or whether such complexity would continue to increase in-

definitely, given the appropriate evolutionary selection pressures and sufficient computational resources. This is due to the fact that such evolutionary systems tend to have low evolvability: any morphological mutation tends to have a much larger behavioral impact than control mutations do, and thus the former are less likely to yield beneficial mutations. This results in a system where evolution becomes trapped in local optima that consist of a fixed morphology; evolution can only realize behavioral improvement through changes in control. In effect then, evolution is forced to perform random search along the morphological dimensions of the fitness landscape (Cheney et al. (2016)).

Yet morphology holds particular promise for realizing OEE beyond that afforded by optimizing the topologies of neural networks in non-embodied agents or embodied agents with fixed morphologies. This is because changes to an agent’s body can expose it to novel sensorimotor experiences. For example, the control policy of a morphologically-fixed bipedal robot can be tuned to realize a large set of gaits, but object manipulation will forever be beyond the agent’s reach unless evolution can exapt legs and feet into arms and hands. A non-embodied deep learning system can continuously improve its ability to categorize objects in images, yet never learn about their affordances (Xia et al. (2018)) unless it is able to evolve sensors and manipulators that allow it to interact with the objects in those images in diverse ways (Cariani (1993); Broersma et al. (2017)).

However, it is likely that potential morphological innovations are often lost during evolutionary simulations because selection removes them from the population before control policy changes are found that realize that potential: hands and feet are useless unless appropriate controllers can enable them to perform object manipulation. To this end, we recently introduced an evolutionary algorithm that lessens selection pressure on embodied agents (and their descendants) that have recently experienced morphological mutations (Cheney et al. (2017)). This method builds on a multi-objective optimization method that protect members of new lineages against members of older lineages within a multi-lineage population Schmidt and Lipson (2011). In (Cheney et al. (2017)), the age of the lineage objective is replaced with the number of generations since the ancestor in a lineage experienced a morphological mutation.

It was shown that if this method is applied to evolving locomotion strategies for soft robots, not only is evolvability increased, but effective open-ended evolution (EOEE) was sustained for over 6000 generations (Fig. 2). Interestingly, EOEE is not achieved if the method is inverted such that potential control innovations are protected for a time to determine whether morphologies can be found by evolution to realize those potentials (Fig. 3). This suggests that morphological change may provide unique opportunities for subsequent evolutionary exploitation that are not provided by control changes.

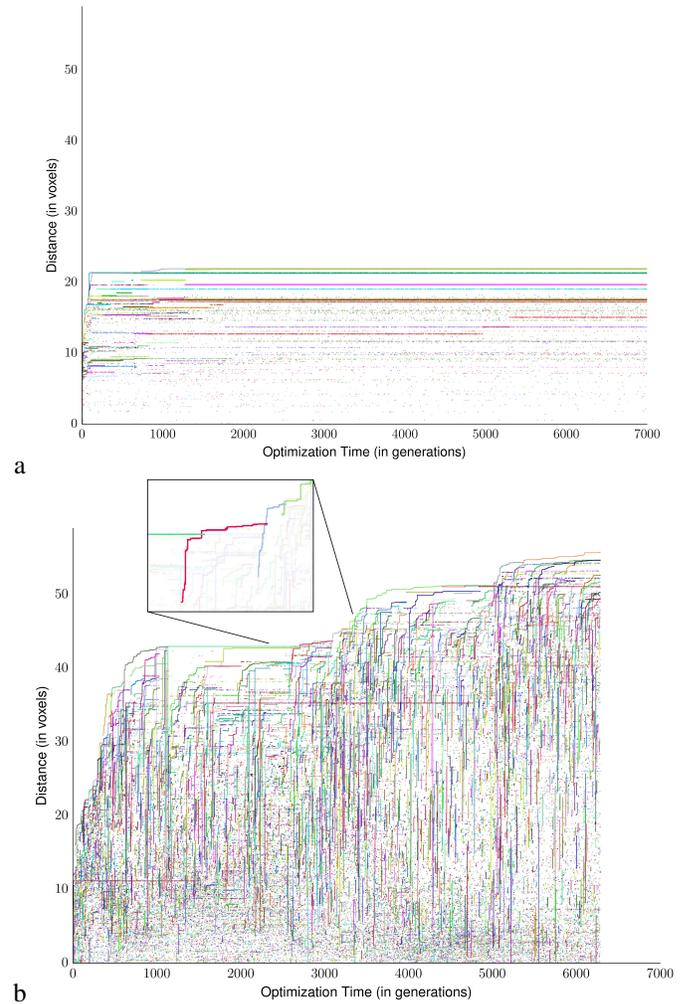
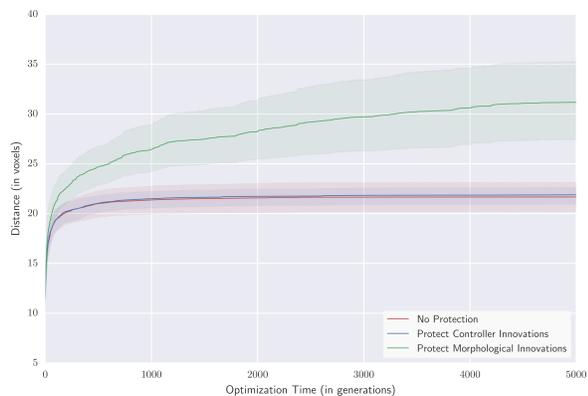


Figure 2: Effective open-ended evolution achieved via morphological protection. In (a), a population of soft robots was evolved to locomote rapidly. In (b), selection pressure was decreased on lineages that recently suffered a morphological mutation. Each colored line corresponds to a collection of robots with identical morphologies but differing control policies.

This work is but one exemplar of a larger body of work (Bernatskiy and Bongard (2018)) that suggests that simulating adaptive processes at different time scales may be one approach to eventually achieve open-ended evolution. More specifically, morphological adaptations occurring at slower time scales may create opportunities to adapt control on faster time scales. In Cheney et al. (2017), both intra-lineage and inter-lineage competition occurred at evolutionary time scales, but multiple timescale simulations that include evolutionary, developmental and learning time scales may yield systems even more capable of exhibiting open-ended evolution (Kriegman et al. (2017)).



a

Figure 3: Protecting potential control innovations to enable evolutionary time for compensating morphological adaptations to be found (blue line) is equivalent to protecting neither morphological nor control innovations (red line), and does not realize effective open-ended evolution in the same way that protecting morphological innovations does (green line).

3. Catastrophic Interference

One roadblock on the path to open-ended evolution is the concept of catastrophic forgetting. This phenomenon is often observed in sequential learning tasks: an agent learns (or is evolved to perform) one task and then, when trained to perform a new task, ‘forgets’ the previous task. This phenomenon is also known as catastrophic interference when an agent is trained to perform multiple tasks simultaneously. A number of methods to resist catastrophic interference have been proposed in the literature, including modularity (Lipson et al. (2002); Kashtan and Alon (2005); Espinosa-Soto and Wagner (2010); Clune et al. (2013); Bongard et al. (2015)) and selective synaptic plasticity (Velez and Clune (2017); Kirkpatrick et al. (2017)).

When methods are reported that successfully resist catastrophic interference, it is usually because the agents have been given some ability to capture features common across their training environments in their neural controllers. In modularity approaches, this facilitates the rewiring of neural feature detectors when their environments change. In synaptic plasticity approaches, plasticity is suppressed for synapses that handle features common across multiple environments. However, both approaches focus on neural rather than morphological adaptation to resist catastrophic interference.

In recent work we have shown that morphology can also help to resist this phenomenon (Powers et al. (2018)). We investigated this in an evolutionary setting in which embodied agents with three different body plans (Fig. 4) were evolved to perform phototaxis in a number of different environments. Catastrophic interference was measured as the

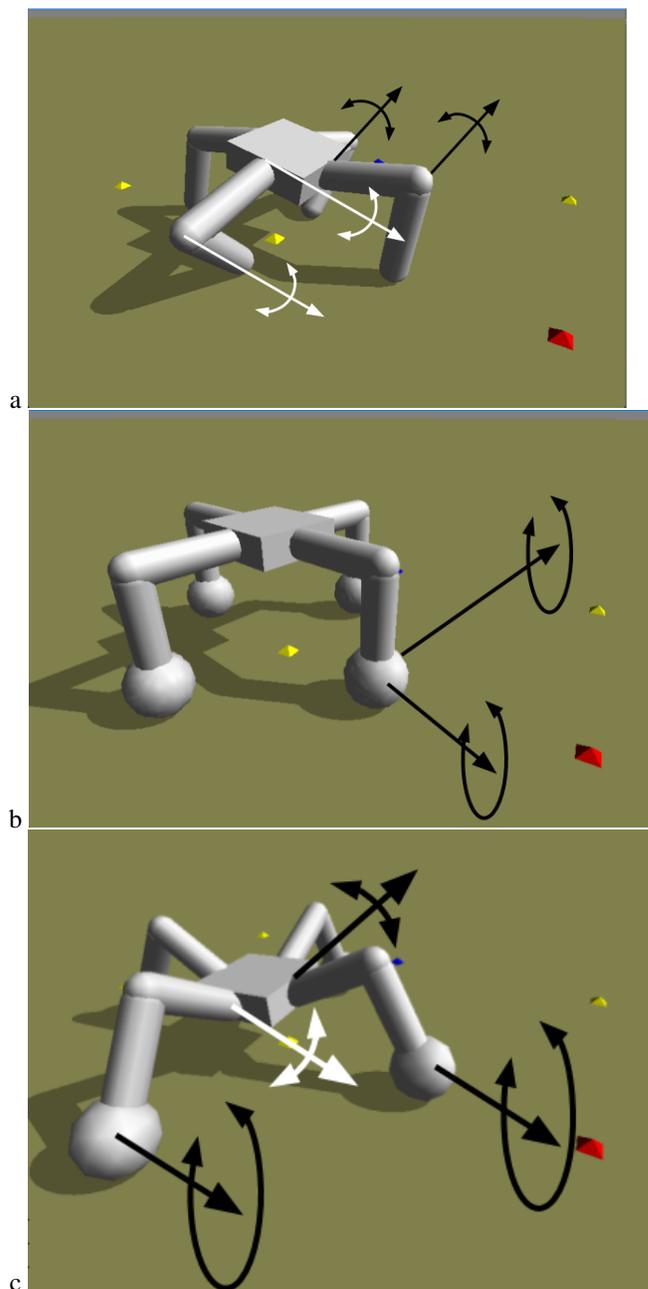


Figure 4: Three classes of robots (legged (a), wheeled (b), and whegged (c)) were evolved using Pyrosim (ccapelle.github.io/pyrosim) to perform phototaxis in multiple environments. Each robot has eight degrees of freedom, as depicted by the black and white arrows which indicate the axis (straight) and direction (curved) of rotation for a particular hinge joint (a, c) or wheel (b, c). Video of all three robot types can be seen at youtu.be/yY7Vi7fw7Ik.

amount of improvement that could be made in an offspring agent across all environments without adversely impacting its parent’s original competency in any one of the environ-

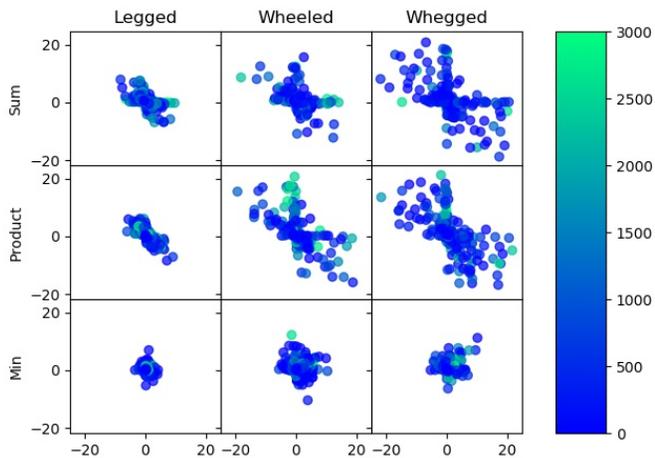


Figure 5: The nine panels correspond to the best individuals produced from nine evolutionary simulations. In the left, center, and right columns, the legged, wheeled, or whegged robot was employed respectively (see Fig. 4). In the top, center, and bottom rows, performance in both environments was summed, averaged, or performance in the worst environment was used to calculate fitness respectively. Each point corresponds to a mutation: a child controller’s fitness relative to its parent controller in the first and second environments is plotted against the horizontal and vertical axes respectively. Points that fall in the upper right of each panel thus indicate child controllers that resisted catastrophic interference: they achieved an increase in fitness in both environments, relative to their parent controllers. Color indicates the generation during which that mutation occurred. As can be seen, the wheeled and whegged robots were much more resistant to catastrophic interference than the legged robot when the best fitness function (*min*) was employed. Indeed these two experimental variants (bottom-middle and bottom-right panels) were significantly more evolvable than the other seven variants. More details about this work are provided in Powers et al. (2018).

ments (Fig. 5). We found that some body plans enabled evolution to find neural changes that achieved this much more often than in other body plans.

The reasons why one morphology resists catastrophic interference better than another is not yet clear. One possibility suggests that, because differing morphologies provide embodied agents with the ability to generate different kinds of sensorimotor feedback loops, some morphologies may allow an agent to align itself with its environment in such a way that two seemingly different environments are perceived, from the point of view of the agent, as similar. Agents that lack this particular morphology may not be able to generate this perception and thus have to evolve or learn

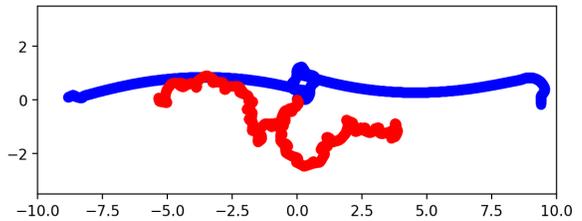


Figure 6: A tracing of a typical whegged robot (blue) and legged robot (red) drawn from the experimental variants reported in the lower left and lower right panels of Fig. 5, respectively (see video youtu.be/uWy33A5HZGM). In environment 1 the light source was placed at (9, 0). In environment 2 the light source was placed at (−9, 0).

different strategies for these ‘different’ environments.

Anecdotal evidence to support this argument can be seen in Fig. 6: the whegged robot produces rotationally symmetric trajectories in two environments, suggesting that its perception over time in these two environments is very similar. A legged robot equipped with an evolved controller, on the other hand, performs two very different trajectories when approaching the light source in these environments, suggesting evolution has had to discover neural networks capable to generating two different sensorimotor feedback loops to accomplish phototaxis in these environments.

Future work will involve investigating how to measure this ability of an embodied agent to induce sensory convergence across seemingly different environments, thus reducing catastrophic interference and facilitating open-ended evolution against a continuously expanding set of task environments. Another future strand of inquiry will involve evolving morphology and control to maximize two objectives: fitness and sensory convergence.

4. Crowdsourcing

Section 2 demonstrated how morphology could may facilitate the eventual realization of OEE for embodied agents operating in one environment; the previous section demonstrated how morphology could facilitate OEE for embodied agents in a growing number of environments. Another dimension along which to pursue OEE is to evolve robots capable of performing an increasing number of tasks.

In evolutionary robotics, this requires construction of multiple fitness terms that select for the various desired behaviors. However, designing a fitness function that does not succumb to perverse instantiation (Bostrom (2014)) is an open challenge in the field: highly-fit embodied agents often evolve that exhibit surprising behavior that does not fulfill the investigator’s original intent (Lehman et al. (2018)). This difficulty is exacerbated when a complex fitness function must simultaneously select for multiple desired behaviors, yet guard against perverse instantiation in all of them.

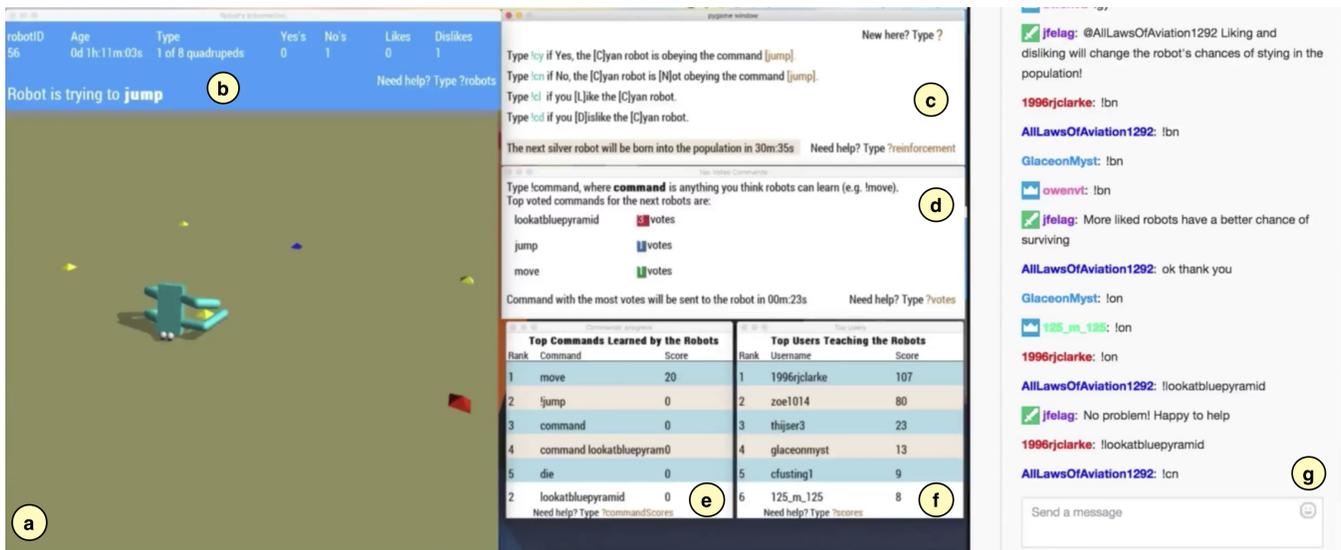


Figure 7: The Twitch Plays Robotics front end, as seen by a participant. **a)** A robot is simulated for thirty seconds under the current command, ‘*jump*’. **b)** A panel listing information about the current robot including its ID, age, and type, the number of yes’s and no’s provided during the current evaluation, and the number of likes and dislikes the current robot has received overall. **c)** A panel explaining how users can reinforce the current robot. **d)** A panel prompting users to propose or vote for the next command for the next three minute window. **e)** A top five commands by score list plus the current proposed command and score. **f)** A top five users by score list plus the last active user in the chat and their score. **g)** The live chat session where users enter reinforcement, commands, help inquiries, and other messages.

In response to this, some have argued for objective-free evolutionary algorithms: only behavioral novelty is selected for (Lehman and Stanley (2008)). However, this approach has proved inadequate for many problems as it weights exploration much too heavily compared to exploitation. In response to this, behavioral novelty is often embedded in a multiobjective optimization framework that includes ‘traditional’ fitness functions that select for the desired behavior as the other objectives (Mouret (2011)).

An alternative path toward realizing agents that can succeed in a growing number of environments while avoiding perverse instantiation is to enable a crowd of non-experts to collectively formulate fitness functions and evaluate embodied agents against them. The field of human-robot interaction (HRI) has long studied ways of enabling non-experts to train robots. However, with the recent advent of free, large-scale, easy-to-use, and abuse-hardened web infrastructure, it has become relatively easy for a small research group to build and deploy systems that enable large numbers of non-experts to interact with simulated or physical robots via the web.

Besides combating perverse instantiation, the crowdsourcing of robotics also offers the promise of scaling up robotics in a similar way that large stores of easy-to-access labeled data sets enabled the scaling up of machine learning methods. Although large amounts of high-dimensional, labeled input data such as video and LIDAR is available for

training robots with few mechanical degrees of freedom (e.g. quadcopters and autonomous cars), high degree-of-freedom (DOF) robots have not benefited from this data revolution because large amounts of labeled *motor* data does not exist: sensor signals generated by a robot’s movements to which humans have attached sensorimotor accounts such as ‘walking’.

In pursuit of the dual goals of combating perverse instantiation and scaling up evolutionary robotics to facilitate OEE, we have been developing an HRI paradigm called ‘Twitch Plays Robotics’¹ (TPR). Twitch.tv is particularly popular among people live streaming themselves playing video games as well as people who enjoy watching others play those games. Each Twitch channel contains a live video stream from a content provider and a real-time chat interface that enables viewers to interact with the content provider and/or the stream itself. Twitch is an attractive option for crowdsourcing studies because it has a very large user base (15 million daily active users); it is hardened against bot and human abuse; it is simple to stream content to, and collect text responses from it; and, given its size, significant numbers of unpaid participants can be recruited, assuming the task set for them is relatively engaging.

The method is summarized visually in Fig. 7 (the system’s front end) and Fig. 8 (the system’s back end) and is

¹twitch.tv/twitchplaysrobotics

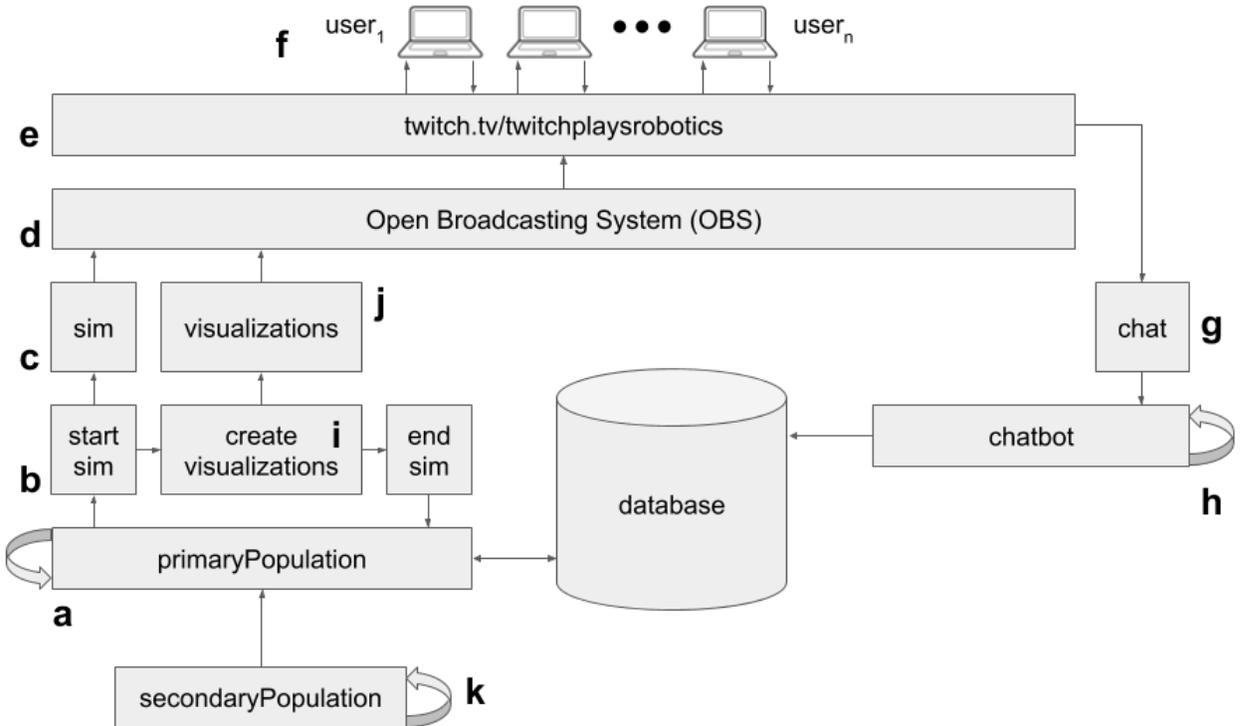


Figure 8: Overview of the Twitch Plays Robotics back end. Three processes (a, h, and k) run continuously and in parallel on a local server. They serve content to, and receive data back from, a third-party web service (twitch.tv). (a) primaryPopulation maintains a population of simulated robots and optimizes them to maximize obedience to a human crowd. Every thirty seconds, a random robot is selected (b), rendered inside of a physics engine (c), and streamed via a local broadcasting application (d) to the web service (e) where the simulation is seen by zero or more participants simultaneously (f). In response to what they see, the participants may provide feedback by typing in text to a chat window (g) which is captured by a chatbot (h) and stored in a database. Additional visualizations are created (i) and overlaid (j) on the streamed video: this includes instructions to the users about the system, as well as a high score board designed to motivate participation. A secondary population (k) continuously evolves robots that exhibit different behaviors when issued with user commands to ensure that they respond when shown to the crowd issued differing commands. Periodically, a robot from the secondary population is injected into the primary population. This also ensures the crowd is shown a continual flow of novel robots and behaviors to retain their interest.

compromised of five parts. First, robots with differing morphologies are shown to human participants, who may interact with them in various ways. Second, the crowd may issue commands to the robots using natural language, which is converted into numbers or vectors that the robots ‘hear’: the encoded language is fed into the input layer of the robot’s neural controller. Third, users may provide positive or negative reinforcement to a given robot who is performing some action in response to a crowd-issued command. Fourth, The robots evolve to collect as much positive reinforcement as possible. Finally, a critic model is trained to take as input a crowd-issued command and sensor data generated by a robot responding to that command, and output a prediction about how much positive reinforcement that robot is likely to collect, if it was exposed to the crowd.

Two experiments have been conducted using the TPR platform to date. In the first (Anetsberger and Bongard

(2015)), we demonstrated that, even though unpaid, participants tended to provide, on average, honest reinforcement. This was proven by showing that the critic could successfully predict crowd reinforcement from an unseen test set when trained on crowd reinforcement from a non-overlapping training set. (If the crowd tended to type in random positive or negative reinforcement, critic predictions would be no better than chance.)

In more recent work (Mahoor et al. (2017)) we exposed the crowd to a set of 10 morphologically-distinct robot ‘species’ (Fig. 9). It was found that, for 10 critics trained against each of these species’ data sets, some of the critics were more predictive than others (Fig. 10). This indicates that morphology is important when attempting to scale up robotics using human trainers: depending on the morphology, a participant may have an easier or harder time determining how to reinforce a particular controller for that mor-

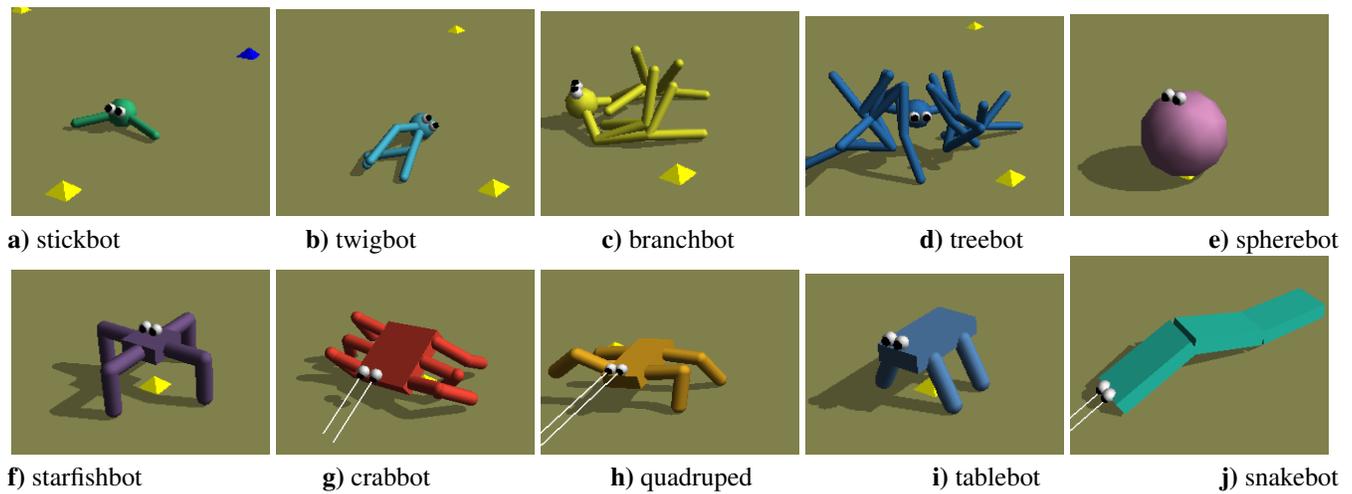


Figure 9: The ten robot species employed in the crowdsourcing experiment reported in (Mahoor et al. (2017)).

phology.

Two challenges facing any crowdsourcing of robotics project is how to elicit commands from the crowd that robots can be evolved to obey, while simultaneously holding the interest of these unpaid participants. In future work we plan to address this challenge by not exposing the crowd to 10 robots with fixed morphologies, but populations of robots with evolving body plans and control policies. By starting with morphologically simple and thus behaviorally constrained agents, we predict the crowd will only issue simple, motoric commands (such as ‘move’ and ‘jump’). As the agents evolve to obey these commands, we will allow gradual morphological complexification. This may suggest new behavioral opportunities to the observing crowd, who may begin to issue more complex commands (such as ‘move toward’ or ‘grasp’). This may in turn lead to a virtual cycle of increasingly complex, capable, and interesting robots, thus better engaging the crowd and growing our number of participants. If this cycle can be sustained indefinitely, crowdsourcing of robotics may prove to be a road toward open-ended evolution.

5. Conclusions

This paper has summarized three experiments which demonstrate how morphology can facilitate the future realization of open-ended evolutionary experiments along three fronts: evolving agents that are more competent, in more environments, at successfully performing more tasks.

Acknowledgements

This work was supported by the NASA Space Technology Research Fellowship #NNX13AL37H, the National Science Foundation award EAGER-1649175, the Army Research Office contract W911NF-16-1-0304, and the Defense Ad-

vanced Research Projects Agency contract HR0011-18-2-0022.

References

- Anetsberger, J. and Bongard, J. (2015). Robots can ground crowd-proposed symbols by forming theories of group mind. In *ALIFE 15: The Fifteen Conference on the Synthesis and Simulation of Living Systems*.
- Bernatskiy, A. and Bongard, J. (2018). Evolving morphology automatically reformulates the problem of designing modular control. *Adaptive Behavior*, 26(2):47–64.
- Bongard, J. C., Bernatskiy, A., Livingston, K., Livingston, N., Long, J., and Smith, M. (2015). Evolving robot morphology facilitates the evolution of neural modularity and evolvability. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, pages 129–136. ACM.
- Bongard, J. C. and Pfeifer, R. (2003). Evolving complete agents using artificial ontogeny. In *Morpho-functional Machines: The new species*, pages 237–258. Springer.
- Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. OUP Oxford.
- Brodbeck, L., Hauser, S., and Iida, F. (2015). Morphological evolution of physical robots through model-free phenotype development. *PLoS one*, 10(6):e0128444.
- Broersma, H., Miller, J. F., and Nichele, S. (2017). Computational matter: Evolving computational functions in nanoscale materials. In *Advances in Unconventional Computing*, pages 397–428. Springer.
- Cariani, P. (1993). To evolve an ear. epistemological implications of gordon pask’s electrochemical devices. *Systems Research and Behavioral Science*, 10(3):19–33.
- Cheney, N., Bongard, J., SunSpiral, V., and Lipson, H. (2016). On the difficulty of co-optimizing morphology and control in evolved virtual creatures. In *Proceedings of the Artificial Life Conference*, volume 2016, pages 226–234.

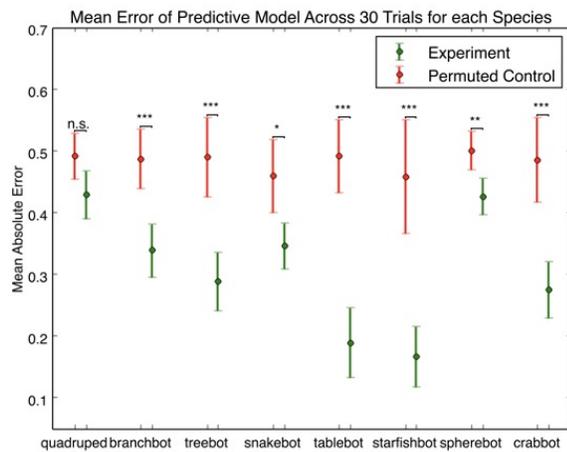


Figure 10: Prediction results for eight species of the 10 species trained by the crowd under the commands ‘move’ and ‘stop’. (Two of the species not receive sufficient reinforcement under this command.) The ‘experiment’ results report the critic’s ability to predict unseen reinforcement from the crowd. The ‘permuted control’ results indicate the predictive ability of critics trained on the same data set but in which the reinforcement signals supplied by participants were permuted among all of the controllers for that species. The p -value between the experimental treatment of the starfishbot and others are: quadruped=***, branchbot=***, treebot=**, snakebot=***, tablebot=n.s., spherebot=***, crabbot=*. The p -value between the experimental treatment of the tablebot and others are: quadruped=***, branchbot=***, treebot=n.s., snakebot=***, starfishbot=n.s., spherebot=***, crabbot=n.s. (***= $p < 0.001$, **= $p < 0.01$, *= $p < 0.05$, and n.s.=otherwise). Error bars report a 99% confidence interval.

Cheney, N., Bongard, J., SunSpiral, V., and Lipson, H. (2017). Scalable co-optimization of morphology and control in embodied machines. *arXiv preprint arXiv:1706.06133*.

Clune, J., Mouret, J.-B., and Lipson, H. (2013). The evolutionary origins of modularity. *Proc. R. Soc. B*, 280(1755):20122863.

Espinosa-Soto, C. and Wagner, A. (2010). Specialization can drive the evolution of modularity. *PLoS computational biology*, 6(3):e1000719.

Hornby, G. S. and Pollack, J. B. (2002). Creating high-level components with a generative representation for body-brain evolution. *Artificial life*, 8(3):223–246.

Hotz, P. E. (2004). Asymmetric cell division and its integration with other developmental processes for artificial evolutionary systems. In *Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Artificial Life*, volume 9, page 387. MIT Press.

Kashtan, N. and Alon, U. (2005). Spontaneous evolution of modularity and network motifs. *Proceedings of the National Academy of Sciences of the United States of America*, 102(39):13773–13778.

Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526.

Komosinski, M. and Rotaru-Varga, A. (2000). From directed to open-ended evolution in a complex simulation model. *Artificial Life*, 7:293–299.

Kriegman, S., Cheney, N., and Bongard, J. (2017). How morphological development can guide evolution. *arXiv preprint arXiv:1711.07387*.

Lehman, J., Clune, J., Misevic, D., Adami, C., Beaulieu, J., Bentley, P. J., Bernard, S., Belson, G., Bryson, D. M., Cheney, N., et al. (2018). The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *arXiv preprint arXiv:1803.03453*.

Lehman, J. and Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*, pages 329–336.

Lipson, H. and Pollack, J. B. (2000). Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799):974.

Lipson, H., Pollack, J. B., and Suh, N. P. (2002). On the origin of modular variation. *Evolution*, 56(8):1549–1556.

Mahoor, Z., Felag, J., and Bongard, J. (2017). Morphology dictates a robot’s ability to ground crowd-proposed language. *arXiv preprint arXiv:1712.05881*.

Mouret, J.-B. (2011). Novelty-based multiobjectivization. In *New horizons in evolutionary robotics*, pages 139–154. Springer.

Powers, J., Kriegman, S., and Bongard, J. (2018). Form, fitness and forgetting: The effects of morphology and fitness on catastrophic interference. In *Proceedings of the 2018 Conference on Artificial Life*. MIT Press. to appear.

Schmidt, M. and Lipson, H. (2011). Age-fitness pareto optimization. In *Genetic Programming Theory and Practice VIII*, pages 129–146. Springer.

Sims, K. (1994). Evolving 3d morphology and behavior by competition. *Artificial life*, 1(4):353–372.

Velez, R. and Clune, J. (2017). Diffusion-based neuromodulation can eliminate catastrophic forgetting in simple neural networks. *PLoS one*, 12(11):e0187736.

Xia, F., R. Zamir, A., He, Z., Sax, A., Malik, J., and Savarese, S. (2018). Gibson env: real-world perception for embodied agents. In *Computer Vision and Pattern Recognition (CVPR), 2018 IEEE Conference on*. IEEE.