10 Morphological Pretraining: Adaptation from the Inside Out

Josh Bongard

Pretraining in Nature

If every component at every level in a biological system is directly exposed to evolutionary selection pressure at all times, the species to which that organism belongs must navigate a rugged fitness landscape with many local optima. That is, most mutations are likely to cause phenotype change at some point in the organism's lifetime that endangers its survival or chance of producing offspring. For this reason, life abounds with examples in which adaptations have evolved that temporarily reduce or completely damp selection pressure.

One notable example of reduced selection pressure in nature is the massive amounts of neutrality that exist in genotype networks: networks in which genotypes of actual organisms are represented as nodes and edges represent single-point mutations that are neutral or nonneutral. Wagner (2014) has demonstrated how such massive neutrality can allow evolution to walk along neutral paths in such networks. The larger the neutral subnetwork within the overall genotype network, the more genotypes can be reached by single-point, nonneutral mutations at the periphery of the neutral subnetwork. This increase in nonneutral neighbors increases the probability that at least one of them will embody a nonneutral beneficial mutation, leading the species to which it belongs to a better fit within its ecological niche.

Other examples of genetic mechanisms that adaptively reduce selection pressure are heat shock proteins (HSPs), which are expressed when the organism experiences stressful conditions. In some cases, HSPS interact with the genome in ways that reduce the phenotypic expression of genetic mutations (Chen et al., 2018). This can result in an adaptive response in which reduced mutational impact occurs when the organism and its progeny are in a beneficial environment, and increased mutational impact occurs when they are in stressful environments. The net result is that the organism and its kin explore alternative phenotypes more when stressed and stay near ancestral phenotypes when not.

Cognition also provides opportunities for reducing selection pressure by, as stated by Popper (1978), simulating the repercussions of contemplated actions: hypotheses can die instead of the organism. Importantly, such mental rehearsal of potentially dangerous behaviors occurs within the organism. It has been assumed that most or all of such mental simulation occurs within neural tissue and that it is particularly concentrated in the central nervous system of higher animals (Tanaka et al., 2020; Wolpert et al., 1998). But work in the area of aneural intelligence (Blackiston et al., 2024; Levin, 2019; Shreesha & Levin, 2024) suggests the possibility that a wider range of organisms, including aneural and plant species, may be capable of internally simulating potentially dangerous actions. Some likely internal arenas for practicing risky or uncertain responses to novel threats are the gut microbiome and the immune system.

All of the examples above, however (with the possible exception of mental rehearsal), do not reduce selection pressure by shielding risky actions (or contemplated actions) within the organism. Embodiment provides additional opportunities for reduction of selection pressure by mechanically shielding internal processes. At large size scales, the within-egg and mammalian uterine environments provide many damping processes: completely fluid-filled environments dampen high acceleration movements experienced by the outer agent (e.g., the mother hitting or being hit by external objects) as well as dampening accelerations caused by the inner agent itself (e.g., the developing organism's own movements). Similarly, the tight packing of the organism as it nears term within the egg or uterus does not provide sufficient runway for high acceleration movements to reach dangerous velocities. At smaller size scales, the vast majority of cells within high cell count multicellular organisms are not in contact with the outside world: Their environments are other cells. Presumably, cells have evolved adaptive mechanisms in which they can shield or dampen the potentially dangerous actions of their neighbors, and/or provide learning gradients for those neighbors to learn from increasingly high risk but high reward actions.

During these periods of reduced selection pressure, evolution, and/or development, the organism itself is free to practice risky moves in genetic space, metabolic space, morphospace, or the space of possible behaviors. These can all be considered forms of pretraining: biologicals can explore where various behaviors fall within a risk/opportunity spectrum. This can help them learn to avoid high-risk/low-opportunity behaviors and how to carefully enter into high-risk/high-opportunity ones. Some of these concepts have been directly or indirectly incorporated into machine learning (ML) methods. These are reviewed in the next section. However, the additional paths via which embodied systems can pretrain themselves to grapple with the danger and uncertainty of their external environments by simulating parts of that world within themselves, has yet to be explored in robotics. This new set of opportunities for robotics is explored in the "Pretraining in Robotics" section. By building robots out of biological rather than technological components, the resulting "biobots" may be able to inherit morphological pretraining mechanisms that have already been installed into those biological components by natural selection. This opportunity is explored in the "Pretraining in Biobotic" section. Concluding remarks about the future of morphological pretraining are provided in the final section.

Pretraining in ML

ML has, since its inception, acknowledged that a careful trade-off must be struck between risky behavior and behavior likely to yield useful behavior. This is commonly referred to as the exploration/exploitation trade-off. Intuitively, and in practice, machine learning methods are designed to favor exploration early and exploitation later.

Pretraining in Premodern Machine Learning

This can be seen in early ML methods such as simulated annealing (Bertsimas & Tsitsiklis, 1993), in which a "temperature" hyperparameter is gradually reduced during training: high temperatures allow candidate solutions to a given problem to make large jumps in solution space; as the temperature falls, the magnitude of solution changes is gradually reduced, relying on the assumption that solutions are increasingly likely to have found a slope leading up to a global optimum as training continues. Other norms within the ML community that point to its awareness that practicing risky play within a protected space is important is the separation of training and testing regimes. Freezing artificial intelligence (AI) model parameters during testing limits the model to exploiting what it has learned during training.

Evolutionary computation is a branch of ML, and evolution strategies are themselves a particular form of evolutionary computation (Hansen et al., 2015). In ES, each genetic parameter that encodes some part of a candidate solution to a given problem has a step size associated with it: this denotes the amount that mutation may vary that parameter when it produces offspring solutions. This early form of automatically tuning the amount of adaptation that occurs during optimization inspired adaptive learning rates (Liu et al., 2019), a subject of study that continues to be of interest within the ML community.

Pretraining in Modern ML

Finally, pretraining of foundation models has risen to prominence in the deep learning community. Once again, these views indirectly acknowledge the value of enabling risky play before exposing a learner to a real-world challenge. Pretraining methods have followed the observation that it is often easier to provide a relatively small amount of domain-specific training data to a foundation model previously trained on a large amount of domain-agnostic training data, compared to training a foundation model for a given domain ab initio.

However, all of the above examples operate on passive, nonembodied learners, like bit strings in evolutionary algorithms or neural networks in ML. Embodied machines offer literal internal spaces within which risk play can occur to prepare it for risky encounters with the outside world. The beginnings of such exploitation in embodied machines like robots is explored in the next section.

Pretraining in Robotics

All current robots, including autonomous vehicles, drones, and semiautonomous machines, like energy harvesters, have limited internal mechanical complexity in the sense of a large number of semi-independent components. As a simple example, if the wheel of an autonomous vehicle blows out, the wheel itself does not contain semi-independent components that can migrate to the site of damage and heal the wound, thus saving the vehicle as a whole from having to adapt to the injury. However, much progress is being made on self-healing materials (e.g., Tan et al., 2021; Wen et al., 2021). This example is but one of many that suggests developing machines within machines could yield robots that are more useful and safer than current machines: if the machine as a whole is overwhelmed by a surprise from within, it would have no recourse but to fail unexpectedly because nothing within can provide a partial or complete compensation.

Nonembodied Rehearsal: Rigid Robots

Grappling with surprise is a primary concern for roboticists, as real-world deployment of physical machines is likely to bring that machine into contact with environments very different from those in which it was trained (see figure 10.1a). Roboticists have therefore considered incorporating internal rehearsal into autonomous machines. The author participated in one such study in which internal rehearsal was cognitive rather than physical: a legged robot constantly trained physical simulations to increasingly capture its own body plan (Bongard et al., 2006). When the robot suffered an injury such as the loss of a leg, the robot updated its "self-models" to reflect this fact. The robot could then use these self-models to rehearse potentially risky actions before attempting them in reality. Subsequent work with a different legged robot (Cully et al., 2015) demonstrated that sufficient rehearsal with a self-model before deployment into real worlds, in which damage did occur, could confer adaptation onto the machine with it having to adapt self-models during deployment. This latter approach thus, in retrospect, can be viewed as an important step toward morphological pretraining, even if no physical components within the body of the robot participated in rehearsals.

Embodied Rehearsal: Soft Robots

Rigid robots, however, even if filled with semi-independent components, would by definition resist internal dynamics ever influencing how the robot exerts force on its external environment due to its rigid shell. Soft robots, on the other hand, by their very nature propagate internal forces to their surfaces. If designed properly, they could modulate which internal forces propagate to their surface, as illustrated in figure 10.1b. Finally, they may be designed or learn to suppress forces arising from risky internal actions and allow vetted internal actions to escape and guide whole-body actions. One initial step in this direction is particle robots (Savoie et al., 2019; S. Li et al., 2019; Ma et al., 2024). This class of autonomous machine is composed of a number of semi-independent robots surrounded by a deformable container. The actions of the internal robots are indirectly aggregated and directed by their interactions with each other and their container. The most well-known exemplar from this class (S. Li et al., 2019) was composed of carefully designed cylindrical components that could dynamically attach and detach from one another. This enabled the overall robot that contained them to achieve locomotion and other simple sensorimotor behaviors, but it also provides the opportunity for subsets of mostly or completely independent subassemblies within the body of the robot to practice risky behaviors. If successful behaviors are found internally, such subassemblies could recruit more or all of the remaining components within the particle robot's body to magnify this successful behavior to the robot as a whole.

A Richer Inner Life: Fluid Robots

If the size of independent components within a particle robot can be sufficiently reduced and the number of them sufficiently increased, not only do more opportunities for internal rehearsal become possible but the robot as a whole can begin to act like a fluid, as shown in figure 10.1c. S. Li et al. (2019) explored scaling of internal particles, and Bielawski et al. (2024) explored optimizing the behavior of "fluid robots" in a differentiable simulator using gradient descent.

Considering a robot as a fluid rather than a collection of a few rigid or soft components may offer even more opportunities for morphological pretraining than nonfluid soft robots do. The actions of one or a few particles within a fluid or granular material can trigger waves that propagate through the material. Careful orchestration of where and when waves are triggered can cause them to reinforce or negate each other via positive or negative interference. One can then envisage prospective actions being encoded as vibrations within a fluid robot. These vibratory prospective actions could be interrogated by other waves that encode environmental features: prospective actions that increasingly survive interrogatory waves would propagate through more of the robot's body. Prospective actions overwhelmed by interrogatory vibrations are failing internal tests and do not propagate to the robot's surface. This hypothetical process would recapitulate motor imagery, in which candidate actions arising in the motor cortex descend to muscle groups, where they become actualized. It would also embody the concept of neural Darwinism (Edelman, 1993), in which ideas putatively compete in the brain for control of the motor system.

Exploiting vibrational dynamics within a granular or fluid material to simulate actions and test them using self-models or world models that are also encoded as vibration may seem counterintuitive. However, a growing body of work shows that vibration may be an attractive modality with which to compute because multiple computations can be performed in the same place at the same time. F. Li et al. (2014) showed that results of Boolean operations can be observed arriving at the same particle but at different points on the frequency spectrum. Parsa et al. (2022a, 2022b, 2024) showed that evolutionary algorithms can be employed to tune the features of grains within a granular material such that two specific Boolean operations arrive at the same place at the same time but at different points on the frequency spectrum. Beaulieu et al. (2024) took this concept of "polycomputing" one step further by showing how AI can further exploit vibrations within a granular material to further increase the material's computational density.

Levin and Bongard (2023) have speculated about how and why polycomputing may exist in biological materials. Finally, there is evidence that neurons may employ soliton waves in addition to electricity for transporting information within neural tissue (Heimburg & Jackson, 2005), and new evidence shows that neurons mix frequencies (Luff et al., 2024). Although this biological evidence does not prove that organisms use vibration to internally rehearse risky actions, it does suggest that there may be no clean separation between electrical and mechanical internal rehearsal in particular, and thought and action in general in biology, and that there may be a good adaptive reason for this: electrically simulating prospective action within the body, no matter how accurate, is no substitute for a dress rehearsal in which physical actions are played out within the body before being allowed to propagate outward to repeat themselves using the body as a whole.

From the Inside Out: Everting and Autotomic Robots

Morphological pretraining could take many forms in biological systems of which we are aware. Similarly, there are many ways morphological pretraining could be incorporated into technological artifacts, some of which would be of adaptive benefit to the machines and others which would not. One way of categorizing different forms of morphological pretraining is to consider the way in which internal ruminations become external reality (see figure 10.1d). An internally considered action, one that is either simulated cognitively or acted out mechanically, may be directly or indirectly communicated to the robot as a whole. Or, internal-to-external transmission may be more direct: internal components that participate in a rehearsal become exposed to the outside world via unintentional or intentional morphological change. Several robots already exist that are capable of this inside-becomes-outside dynamic. So-called vine robots that evert the internal surfaces of their cylindrical bodies are capable of this (Blumenschein et al., 2020), as are autotomic robots (Davis et al., 2023). Autonomic robots can intentionally jettison parts of their body for adaptive purposes. Several species are capable of autotomy, usually to sacrifice a limb rather than their life to a predator (Emberts et al., 2019).

In everting or autonomic robots, parts of their internal structure can become exposed to the outside environment. At present, this property is binary: either a part of a robot is internal or it is directly exposed to the outside world. However, in nature, many body plans, or parts of body plans, of plants or animals are fractal in nature. In such fractal morphologies, there can be a gradation between inside and outside: oxygenated blood, alveoli, alveolar sacs, bronchioles, bronchi, the trachea, and the mouth and nose describe a gradual progression from inside to outside. In plants, flower buds may be shielded from sun, wind, and pests by surrounding leaves on nearby twigs. ML has taught us that gradients of all kinds can be helpful in enabling the learner, be it a nonembodied machine, embodied machine, or organism, to confront potentially risky behavior with a series of environments that increasingly resemble the real environment for which they are being trained. One such gradient is passing a behavior along a series of body parts that are increasingly in contact with the outside world.

Robots capable of having plans buried in internal material offer new vistas for future robotics. One can imagine the two halves of a robot that has been cut down the middle forming two smaller, geometrically and behaviorally similar versions of the original robot. One can imagine a robot everting usually shielded internal sensors to the outside environment and then retracting them to present that stimuli to internal rehearsal processes. Or it may evert internal manipulators that grasp bits of an unknown surface and draw them inward for internal prospective legs to try walking across. One can imagine a robot that grows to increase its fractal dimension, thereby smoothing its internal-to-external gradient, providing more intermediate testing grounds for internally germinated prospective actions.

From the Outside In: Thermodynamically Open Robots

The converse of autotomic robots is thermodynamically open robots. This class of machine is capable of incorporating new energy and matter into itself. This process has been approximated for a long time in robotics, in the form of modular robots: robots that have the ability to temporarily or permanently attach to one another, as illustrated in figure 10.1e. During attachment, what were previously externally facing surfaces become internal components of a composite robot. During detachment, like autonomic robots, internal surfaces or parts can become directly exposed to external environments. Modular robots usually transition rapidly between fully attached and detached states. However, more complex morphologies in the future may enable modular robots to modulate how much or how little components are connected, thereby creating another form of inside-to-outside gradients along which internally rehearsed actions can migrate into external actions that directly affect the robot's external environment.

Most modular robotics platforms adhere to the same topological change: components lock together two- or three-dimensionally like Lego bricks. However, the rise of robots with radically different body plans, like robotic fabrics (Buckner & Kramer-Bottiglio, 2018), suggests other ways that robots may combine: a robot fabric could envelop another robot, thereby modulating the behaviors arising from the engulfed machine. Again, this would shield the actions of the engulfed robot from the outside world, perhaps allowing it to practice risky behaviors. If it finds one, it could instruct the engulfing robot to withdraw. Interestingly, one robot engulfing another would serve as a morphological form of fine-tuning described in the section "From the Inside Out: Everting and Autotomic Robots": it may be easier to realize a robot with a desired behavior by engulfing a robot that already approximates that behavior with a robot fabric that "fine-tunes" that behavior into the desired one.

It's Robots All the Way Down: Fractal Robotics

Nature abounds with organisms that exhibit fractal structure, from corals to most plants to the vascular, skeletal, muscular, pulmonary, and nervous systems of higher animals. Many of these biological fractals are not just structurally self-similar but behaviorally self-similar as well: alveoli, alveolar sacs, bronchioles, bronchi, and the trachea each individually, and collectively, facilitate gas exchange. Intuitively, an organism or machine could increase its utility, or continue operation in a wider range of environments, if it learns to compose itself together with others of its kind in a way that yields a large version of themselves, capable of exhibiting the same behavior but at larger size scales, as shown in figure 10.1f. To achieve this, a set of robots or organisms could practice aggregation and disaggregation in different ways to improve the fidelity of structural and behavioral recapitulation at larger size scales: they could learn to improve at achieving structural and behavioral self-similarity. Fractal robots were imagined (Moravec et al., 1996) and have since been built (Kriegman et al., 2021), but the latter are evaluated in either a disaggregated or fractal state: the ways in which they assemble and disassemble into fractal forms and functions has yet to be studied. Certain fractals, like branching structures, do not have a clear inside and outside. Other fractals, such as the Menger sponge, do possess such a distinction. One could imagine a nonfractal robot practicing to create smaller versions of its own form and function internally. Once that is learned, by definition, the smaller copies would also be capable of internal assembly of even smaller copies.

The adaptive advantage of fractal robots over nonfractal robots is not yet clear, but its ubiquity in nature suggests it does have such an advantage. One potential usefulness is the compactness of encoding the design, and thus indirectly the function, of fractal robots. Another potential use might be to achieve a good fit between the fractal robot and fractal environments: a fractal robot may be best able to efficiently navigate a cave system with fractally arranged tunnels or inspect a coral reef. A third potential use might be the ability to rapidly exhibit some useful behavior at very different size scales by aggregation, de-aggregation, or eversion (exposing internally assembled smaller versions to the external environment).

Pretraining in Biobotics

If robots could benefit from a complex internal environment within which to practice risky behavior, one approach is to create bioinspired robots: robots composed of technological components simulate the complex internal processes that occur in organisms. Another, arguably more direct

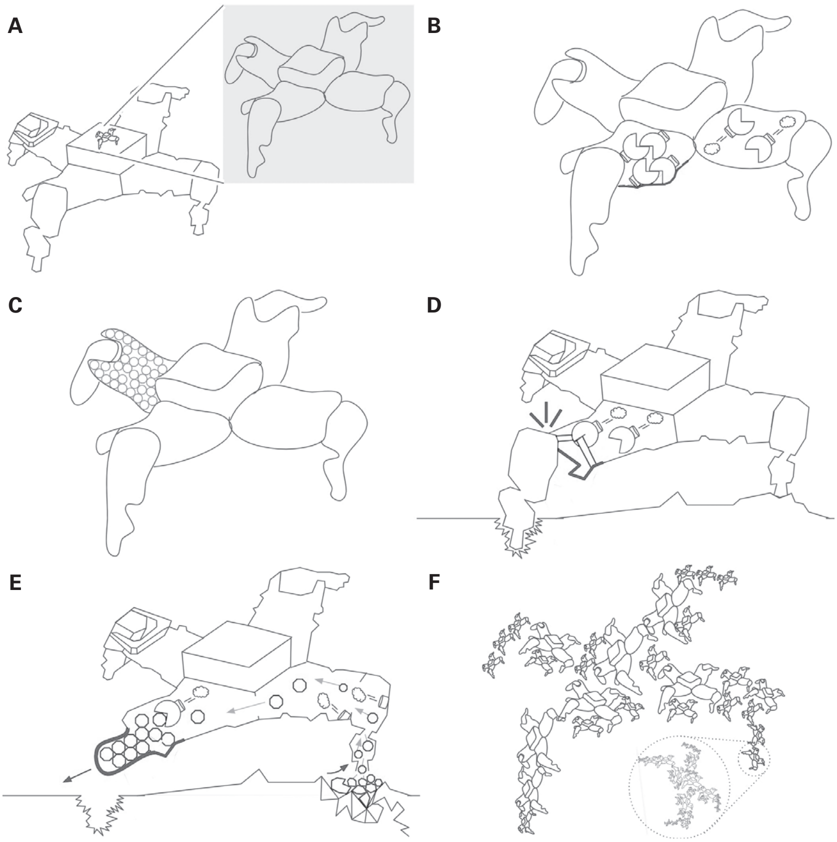

Figure 10.1
Morphological pretraining in present and future robots. (a) Rigid robots are capable of cognitively rehearsing risky actions using self-models. (b) Future soft robots could contain robots that physically rehearse actions internally. (c) Fluid robots could contain enough small robots to act as a fluid. (d) Autonomic robots could lose body parts, enabling well-rehearsed internal robots to become exposed to the outside world. (e) Thermodynamically open robots could absorb external material, involve it in internal processes, and direct it to (re)grow lost/new body parts. (f) Robots could learn to assemble into fractal structures such that the aggregate looks and acts like its components. *Source*: the author.

approach, is to create biobots: robots composed of biological rather than technological components. The latter approach has the benefit that, even though biological parts may have to be harvested from living hosts and reconfigured in new ways, this transformation may not disrupt the resulting biobot's ability to recover and maintain a complex inner environment.

If this is true, biobots can "inherit" part or all of the morphological pretraining capabilities of the biological donor. There is some evidence to support this latter speculation, as existing biobots do show the ability to survive radical reconfigurations that differ greatly from their host's phenotype. Whether or how much morphological pretraining facilitates survival across the radical transformation remains to be seen.

Pretraining in Xenobots

Two classes of biobots have been introduced to the literature to date: Xenobots, assembled from frog cells (Blackiston et al., 2021; Kriegman et al., 2020, 2021), and anthrobots, grown from human cells (Gumuskaya et al., 2023). Xenobots were the first class of biobot to be built. They were constructed from several thousand frog epithelial cells. If constructed appropriately, cells on the surface of the multicellular assembly grow cilia, which beat against the surrounding water and often result in motility. These cilia grow on the surface of wild-type frogs, facilitating the sloughing of skin-related pathogens. The fact that Xenobots still grow cilia despite the fact that their internal structure differs greatly from the tadpole or adult frog—the interior of a Xenobot is only more epithelial cells—suggests that morphological pretraining may have helped here, as these two frog phenotypes share a similar exterior but survive with two very different internal environments.

Much additional work is required to determine whether the short period of development that the embryonic frog skin cells enjoyed before being harvested prepared them to survive when reconfigured into Xenobots. Similarly, it would be of interest to discover whether these cells participate in internal rehearsal of risky actions preharvest and whether they continue to do so in Xenobots. If hallmarks of morphological pretraining can be found in embryonic frog tissues, this generates a testable hypotheses: embryos from species that similarly engage in morphological pretraining are more amenable to AI-dictated reconfigurations than embryos from species that do not.

Pretraining in Xenomics

The above-described research questions depend on a crisper definition of morphological pretraining than simply "internal rehearsal of risky action" and the ability to observe such dynamics at work in living tissues. Recent work has begun interrogation along these lines by looking for adaptive internal migration of cells within Xenobots: internal mechanical rearrangement, rather than cognition within statically fixed internal tissue. Some evidence has pointed to internal cellular migration and adaptation (Blackiston et al., 2024), but a link between these two phenomena has yet to be found. Local injuries were visited on Xenobots by puncturing them, and imaging was conducted pre- and postinjury. It was observed that cells near the point of damage migrate inward to fill the wound. It was also found that for most of the injured Xenobots, there was a common change in the way information flows within the Xenobot pre- and postdamage, where information flow was inferred from calcium activity within and between the cells.

Next steps will involve seeking cellular migration that resembles the migratory patterns that were observed postdamage: do Xenobots in particular, or frog or other embryos in general, "rehearse" injuries and responses to them, before they ever occur? A positive answer to this question could facilitate a broader research program in which organisms are observed to recover from dangerous situations with the help of internal changes and those internal changes are observed within the organism long before that dangerous situation arises.

If a growing number of biological examples of morphological pretraining can be discovered, they would collectively form a training set that AI could employ to learn the common patterns underlying them. The AI could then learn how to embody those dynamics in an ever-widening array of robots that are composed of technological components, biological components, or increasingly exotic admixtures of both.

Conclusion

Much ink has been spilled pointing to differences between machines and organisms. This chapter has explored one overlooked distinction: the potential ability of organisms to internally rehearse risky actions using dynamic components, not just cognitively using static internal components. Whether or how extensively this phenomenon is at work in nature has yet to be determined, but it is not difficult to formulate ways in which such an ability could confer adaptive advantages on the organism. Almost all current technologies, including state-of-the-art autonomous robots, do not have internal components capable of significant rearrangement. Thus, their ability to perform morphological pretraining is greatly limited or impossible. For this reason, effort in robotics should be expended to create machines capable of complex internal change, and efforts in embodied intelligence (EI) should be dedicated to determining how such machines can exploit this internal change to rehearse responses to the myriad challenges the external world will throw at them, long before those challenges ever arrive.

References

Beaulieu, S., Welch, P., Parsa, A., O'Hern, C., Kramer-Bottligio, R., & Bongard, J. (2024). Refractive Computation: Parallelizing logic gates across driving frequencies in a mechanical polycomputer. *Artificial Life Conference Proceedings*, *36*(1), p. 72.

Bertsimas, D., & Tsitsiklis, J. (1993). Simulated annealing. *Statistical Science: A Review Journal of the Institute of Mathematical Statistics*, *8*(1), 10–15. https://doi.org/10.1214/ss/1177011077

Bielawski, K., Rozlivek, J., Hoffmann, M., & Bongard, J. (2024). Best practices for differentiable soft robot modeling and optimization with the material point method. In *Artificial Life Conference Proceedings*, *36*(1), p. 11. https://doi.org/10.1162/isal_a_00724

Blackiston, D., Dromiack, H., Grasso, C., Varley, T. F., Moore, D. G., Srinivasan, K., Sporns, O., Bongard, J., Levin, M., & Walker, S. I. (2024). Revealing non-trivial information structures in aneural biological tissues via functional connectivity. *bioRxiv*. https://doi.org/10.1101/2024.05.09.593467

Blackiston, D., Lederer, E., Kriegman, S., Garnier, S., Bongard, J., & Levin, M. (2021). A cellular platform for the development of synthetic living machines. *Science Robotics*, *6*(52), eabf1571. https://doi.org/10.1126/scirobotics.abf1571

Blumenschein, L. H., Coad, M. M., Haggerty, D. A., Okamura, A. M., & Hawkes, E. W. (2020). Design, modeling, control, and application of everting vine robots. *Frontiers in Robotics and AI*, *7*, 548266. https://doi.org/10.3389/frobt.2020.548266

Bongard, J., & Levin, M. (2023). There's plenty of room right here: Biological systems as evolved, overloaded, multi-scale machines. *Biomimetics*, *8*(1), 110. https://doi.org/10.3390/biomimetics8010110

Bongard, J., Zykov, V., & Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science*, *314*(5802), 1118–1121. https://doi.org/10.1126/science.1133687

Buckner, T. L., & Kramer-Bottiglio, R. (2018). Functional fibers for robotic fabrics. *Multifunctional Materials*, *1*(1), 012001. https://doi:10.1088/2399-7532/aad378

Chen, B., Feder, M. E., & Kang, L. (2018). Evolution of heat-shock protein expression underlying adaptive responses to environmental stress. *Molecular Ecology*, *27*(15), 3040–3054. https://doi.org/10.1111/mec.14769

Cully, A., Clune, J., Tarapore, D., & Mouret, J.-B. (2015). Robots that can adapt like animals. *Nature*, *521*(7553), 503–507. https://doi.org/10.1038/nature14422

Davis, Q. T., Woodman, S., Landesberg, M., Kramer-Bottiglio, R., & Bongard, J. (2023). Subtract to adapt: Autotomic robots. *2023 IEEE International Conference on Soft Robotics (RoboSoft)*, *100* (pp. 1–6). IEEE. https://doi.org/10.1109/ROBOSOFT55895.2023.10122102

Edelman, G. M. (1993). Neural Darwinism: Selection and reentrant signaling in higher brain function. *Neuron*, *10*(2), 115–125.

Emberts, Z., Escalante, I., & Bateman, P. W. (2019). The ecology and evolution of autotomy: Ecology and evolution of autotomy. *Biological Reviews of the Cambridge Philosophical Society*, *94*(6), 1881–1896. https://doi.org/10.1111/brv.12539

Gumuskaya, G., Srivastava, P., Cooper, B. G., Lesser, H., Semegran, B., Garnier, S., & Levin, M. (2023). Motile living biobots self-construct from adult human somatic progenitor seed cells. *bioRxiv*. https://doi.org/10.1101/2022.08.04.502707

Hansen, N., Arnold, D. V., & Auger, A. (2015). Evolution strategies. In J. Kacprzyk & W. Pedrycz (Eds.), *Springer handbook of computational intelligence* (pp. 871–898). Springer.

Heimburg, T., & Jackson, A. D. (2005). On soliton propagation in biomembranes and nerves. *Proceedings of the National Academy of Sciences*, *102*(28), 9790–9795. https://doi.org/10.1073/pnas.0503823102

Kriegman, S., Blackiston, D., Levin, M., & Bongard, J. (2020). A scalable pipeline for designing reconfigurable organisms. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(4), 1853–1859. https://doi.org/10.1073/pnas.1910837117

Kriegman, S., Blackiston, D., Levin, M., & Bongard, J. (2021). Kinematic self-replication in reconfigurable organisms. *Proceedings of the National Academy of Sciences of the United States of America*, *118*(49), e2112672118. https://doi.org/10.1073/pnas.2112672118

Kriegman, S., Nasab, A. M., Blackiston, D., Steele, H., Levin, M., Kramer-Bottiglio, R., & Bongard, J. (2021). Scale invariant robot behavior with fractals. *arXiv*. http://arxiv.org/abs/2103.04876

Levin, M. (2019). The computational boundary of a "self": Developmental bioelectricity drives multicellularity and scale-free cognition. *Frontiers in Psychology*, *10*, 2688. https://doi.org/10.3389/fpsyg.2019.02688

Li, F., Anzel, P., Yang, J., Kevrekidis, P. G., & Daraio, C. (2014). Granular acoustic switches and logic elements. *Nature Communications*, *5*(1), 5311. https://doi.org/10.1038/ncomms6311

Li, S., Batra, R., Brown, D., Chang, H.-D., Ranganathan, N., Hoberman, C., Rus, D., & Lipson, H. (2019). Particle robotics based on statistical mechanics of loosely coupled components. *Nature*, *567*(7748), 361–365. https://doi.org/10.1038/s41586-019-1022-9

Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., & Han, J. (2019). On the variance of the adaptive learning rate and beyond. *arXiv*. http://arxiv.org/abs/1908.03265

Luff, C. E., Peach, R., Mallas, E.-J., Rhodes, E., Laumann, F., Boyden, E. S., Sharp, D. J., Barahona, M., & Grossman, N. (2024). The neuron mixer and its impact on human brain dynamics. *Cell Reports*, *43*, 114274. https://doi.org/10.1016/j.celrep.2024.114274

Ma, D., Chen, J., Cutler, S., & Petersen, K. (2024). Smarticle 2.0: Design of scalable, entangled smart matter: 16th international symposium. In J. Bourgeois, J. Paik, B. Piranda, J. Werfel, S. Hauert, A. Pierson, H. Hamann, T. L. Lam, F. Matsuno, N. Mehr, & A. Makhoul (Eds.), *Distributed autonomous robotic systems*, *28* (pp. 509–522). Springer Nature.

Moravec, H., Easudes, J., & Dellaert, F. (1996, December). *Fractal branching ultra-dexterous robots (Bush robots)* (PR No. 10-86888). NASA Advanced Concepts Research Project. https://www.ri.cmu.edu/publications/fractal-branching-ultra-dexterous-robots-bush-robots-/

Parsa, A., O'Hern, C. S., Kramer-Bottiglio, R., & Bongard, J. (2024). Gradient-based design of computational granular crystals. *arXiv*. http://arxiv.org/abs/2404.04825

Parsa, A., Wang, D., O'Hern, C. S., Shattuck, M. D., Kramer-Bottiglio, R., & Bongard, J. (2022). Evolution of acoustic logic gates in granular metamaterials. In *International Conference on the Applications of Evolutionary Computation (Part of EvoStar)* (pp. 93-109). Springer International Publishing. https://doi.org/10.1007/978-3-031-02462-7_7

Parsa, A., Wang, D., O'Hern, C. S., Shattuck, M. D., Kramer-Bottiglio, R., & Bongard, J. (2022). Evolving programmable computational metamaterials. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '22)* (pp. 1–8). ACM. https://doi.org/10.1145/3512290.3528861

Popper, K. (1978). Natural selection and the emergence of mind. *Dialectica*, *32*(3–4), 339–355.

Savoie, W., Berrueta, T. A., Jackson, Z., Pervan, A., Warkentin, R., Li, S., Murphey, T. D., Wiesenfeld, K., & Goldman, D. I. (2019). A robot made of robots: Emergent transport and control of a smarticle ensemble. *Science Robotics*, *4*(34), eaax4316. https://doi.org/10.1126/scirobotics.aax4316

Shreesha, L., & Levin, M. (2024). Stress sharing as cognitive glue for collective intelligences: A computational model of stress as a coordinator for morphogenesis. *Biochemical and Biophysical Research Communications*, *731*, 150396. https://doi.org/10.1016/j.bbrc.2024.150396

Tan, Y. J., Susanto, G. J., Anwar Ali, H. P., & Tee, B. C. K. (2021). Progress and roadmap for intelligent self-healing materials in autonomous robotics. *Advanced Materials*, *33*(19), 2002800. https://doi.org/10.1002/adma.202002800

Tanaka, H., Ishikawa, T., Lee, J., & Kakei, S. (2020). The cerebro-cerebellum as a locus of forward model: A review. *Frontiers in Systems Neuroscience*, *14*, 19. https://doi.org/10.3389/fnsys.2020.00019

Wagner, A. (2014). *Arrival of the fittest: Solving evolution's greatest puzzle*. Penguin.

Wen, N., Song, T., Ji, Z., Jiang, D., Wu, Z., Wang, Y., & Guo, Z. (2021). Recent advancements in self-healing materials: Mechanicals, performances and features. *Reactive & Functional Polymers*, *168*, 105041. https://doi.org/10.1016/j.reactfunctpolym.2021.105041

Wolpert, D. M., Miall, R. C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences*, *2*(9), 338–347. https://doi.org/10.1016/S1364–6613(98)01221–2

This is a section of doi:10.7551/mitpress/15999.001.0001

Embodied Intelligence

Multidisciplinary Perspectives on Natural, Artificial, and Hybrid Systems

Edited by: Sheila L. Macrine, Jennifer M. B. Fugate, Arsen Abdulali, Josie Hughes

Citation:

Embodied Intelligence: Multidisciplinary Perspectives on Natural, Artificial, and Hybrid Systems
Edited by: Sheila L. Macrine, Jennifer M. B. Fugate, Arsen Abdulali, Josie Hughes
DOI: 10.7551/mitpress/15999.001.0001
ISBN (electronic): 9780262053501
Publisher: The MIT Press
Published: 2026

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open

The MIT Press

The MIT Press
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
mitpress.mit.edu

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data is available.

ISBN: 978-0-262-05349-5

EU Authorised Representative: Easy Access System Europe, Mustamäe tee 50, 10621 Tallinn, Estonia | Email: gpsr.requests@easproject.com