

# How morphological development can guide evolution

Sam Kriegman<sup>1,\*</sup>, Nick Cheney<sup>2</sup>, and Josh Bongard<sup>1</sup>

<sup>1</sup>University of Vermont, Department of Computer Science, Burlington, VT, USA

<sup>2</sup>University of Wyoming, Department of Computer Science, Laramie, WY, USA

\*sam.kriegman@uvm.edu

## ABSTRACT

Organisms result from adaptive processes interacting across different time scales. One such interaction is that between development and evolution. Models have shown that development sweeps over several traits in a single agent, sometimes exposing promising static traits. Subsequent evolution can then canalize these rare traits. Thus, development can, under the right conditions, increase evolvability. Here, we report on a previously unknown phenomenon when embodied agents are allowed to develop and evolve: Evolution discovers body plans robust to control changes, these body plans become genetically assimilated, yet controllers for these agents are not assimilated. This allows evolution to continue climbing fitness gradients by tinkering with the developmental programs for controllers within these permissive body plans. This exposes a previously unknown detail about the Baldwin effect: instead of all useful traits becoming genetically assimilated, only traits that render the agent robust to changes in other traits become assimilated. We refer to this as *differential canalization*. This finding also has implications for the evolutionary design of artificial and embodied agents such as robots: robots robust to internal changes in their controllers may also be robust to external changes in their environment, such as transferal from simulation to reality or deployment in novel environments.

## Introduction

The shape of life changes on many different time scales. From generation to generation, populations gradually increase in complexity and relative competency. At the individual level, organisms grow from a single-celled egg and exhibit extreme postnatal change as they interact with the outside world during their lifetimes. At a faster time scale still, organisms behave such as to survive and reproduce.

Many organisms manifest different traits as they interact with their environment. It seems wasteful not to utilize this extra exploration to speed the evolutionary search for good genotypes. However, to communicate information from these useful but temporary traits to the genotype requires inverting the generally very complex, nonlinear and stochastic mapping from DNA to phenotype. Inverting such a function would be exceedingly difficult to compute. However, organisms can pass on their particular capacity to acquire certain characteristics. Thus phenotypic plasticity can affect the direction and rate of evolutionary change by influencing selection pressures. Although this phenomenon was originally described by Baldwin<sup>1</sup>, Morgan<sup>2</sup> and Waddington<sup>3</sup>, among others, it has become known as ‘the Baldwin effect’. In Baldwin’s words: ‘the most plastic individuals will be preserved to do the advantageous things for which their variations show them to be the most fit, and the next generation will show an emphasis of just this direction in its variations’<sup>1</sup>. In a fixed environment, when the ‘advantageous thing’ to do is to stay the same, selection can favor genetic variations which more easily, reliably, or quickly produce these traits. This can lead to the genetic determination of a character which in previous generations needed to be developed or learned.

Thirty years ago, Hinton and Nowlan<sup>4</sup> provided a simple computational model of the Baldwin effect that clearly demonstrated how phenotypic plasticity could, under certain conditions, speed evolutionary search without communication to the genotype. They considered the evolution of a bitstring that is only of value when perfectly matching a predefined target string. The search space therefore has a single spike of high fitness with no slope leading to the summit. In such a space, evolution is no better than random search.

Hinton and Nowlan then allowed part of the string to randomly change at an additional (and faster) developmental time scale. When the genetically specified portion of the string is correct, there is a chance of discovering the remaining portion in development. The speed at which such individuals tend to find the good string will be proportional to the number of genetically determined bits. When the target string is found, development stops and the individual is rewarded for the amount of remaining developmental time. This has the effect of creating a gradient of increasing fitness surrounding the correct specification that natural selection can easily climb by incrementally assimilating more correct bits to the genotype.

Hinton and Nowlan imagined the bitstring as specifying the connections of a neural network in a very harsh environment. We are also interested in this interaction of subsystems unfolding at different time scales, but consider an embodied agent situated in a physically-realistic environment rather than an abstract control system. This distinction is important as it grounds our hypotheses in the constraints and opportunities afforded by the physical world. It also allows us to investigate how changes in morphology and control might differentially affect the direction or rate of evolutionary search. More specifically, it exposes the previously unknown phenomenon of differential canalization reported here.

Inspired by Hinton and Nowlan, Floreano and Mondada<sup>5</sup> explored the interaction between learning and evolution in mobile robots with a fixed body plan but plastic neural control structure. They noted that the acquisition of stable behavior in ontogeny did not correspond to stability (no further change) of individual synapses, but rather was regulated by continuously changing synapses which were dynamically stable. In other words, agents exploited this ontogenetic change for behavior, and this prevented its canalization. In this paper, we structure development in a way that restricts its exploitation for behavior and thus promotes the canalization of high performing static phenotypes. Also, the robot's body plan was fixed in Floreano and Mondada's experiments<sup>5</sup>, whereas in the work reported here, evolution and development may modify body plans.

Several models that specifically address morphological development of embodied agents have been reported in the literature<sup>6-9</sup>. However, the relationship between morphological development and evolvability is seldom investigated in such models. Moreover, there are exceedingly few cases that considered postnatal change to the body plan (its resting structural form) as the agent behaves and interacts with the environment through physiological functioning (at a faster time scale).

We are only aware of four cases reported in the literature in which a simulated robot's body was allowed to change while it was behaving. In the first two cases<sup>10,11</sup>, it was not clear whether this ontogenetic morphological change facilitated the evolution of behavior. Later, Bongard<sup>12</sup> demonstrated how such change could lead to a form of self-scaffolding that smoothed the fitness landscape and thus increased evolvability. This ontogenetic change also exposed evolution to much more variation in sensor-motor interactions which increased robustness to unseen environments. Recently, Kriegman *et al.*<sup>13</sup> showed how development can sweep over a series of body plans in a single agent, and subsequent heterochronic mutations canalize the most promising body plan in more morphologically-static descendants.

We are not aware of any cases reported in the literature to date in which a simulated robot's body and control are simultaneously allowed to change while it is behaving. In this paper, we investigate such change in the morphologies and controllers of soft robots as they are evolved for coordinated action in a simulated 3D environment. By morphology we mean the current state of a robot's shape, which is slowly changed over the course of its lifetime by a developmental process. We distinguish this from the controller, which sends propagating waves of actuation throughout the individual, which also affects the instantaneous shape of the robot but to a much smaller degree. We here refer to these two processes as 'morphology' and 'control'. As both processes change the shape, and thus behavior, of the robot, this distinction is somewhat arbitrary. However, the central claim of this paper, which is that some traits become canalized while others do not, is not reliant on this distinction.

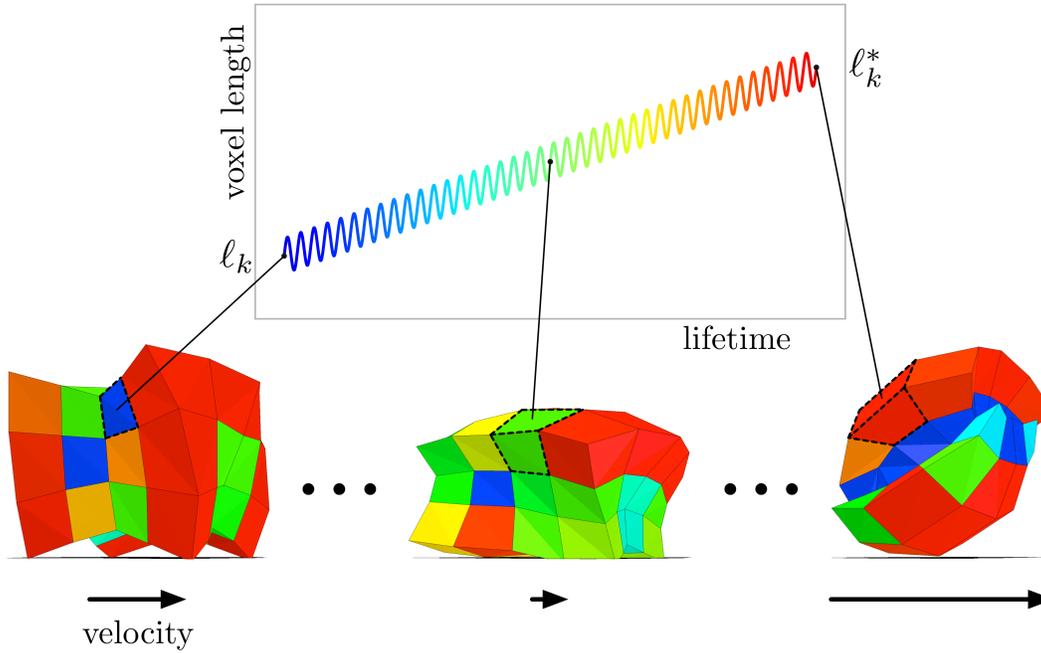
We use soft robots because they provide many more degrees of morphological freedom compared to traditional robots composed of rigid links connected by rotary or linear actuators. This flexibility allows soft robots to accomplish tasks that would be otherwise impossible for their rigid-bodied counterparts, such as squeezing through small apertures<sup>14</sup> or continuously morphing to meet different tasks. Recent advancements in materials science are enabling the fabrication of 3D-printed muscles<sup>15</sup> and nervous systems<sup>16</sup>. However, there are several challenges to the field of soft robotics, including an overall lack of design intuition: What should a robot with nearly unbounded morphological possibility look like, and how can it be controlled? Controllability often depends on precision actuation and feedback control, but these properties are difficult to maintain in soft materials in which motion in one part of the robot can propagate in unanticipated ways throughout its body<sup>17</sup>.

We present here a minimally complex but embodied model of morphological and neurological development. This new model represents an alternative approach to the challenging problem of soft robot design and presents an *in silico* testbed for hypotheses about evolving and developing embodied systems. This model led to the discovery of differential canalization and how it can increase evolvability.

## Results

All experiments were performed in the open-source soft-body physics simulator *Voxelyze*, which is described in detail by Hiller and Lipson<sup>18</sup>. We consider a locomotion task, over flat terrain, for soft robots composed of a  $4 \times 4 \times 3$  grid of voxels (Fig. 1). Robots are evaluated for 40 actuation cycles at 4 Hz, yielding a lifetime of ten seconds. Fitness is taken to be distance traveled measured in undeformed body lengths (four unit voxels, i.e. 4 cm). Example robots are shown in Figures 1 and 5, and Supplementary Video S1.

The morphology of a robot is given by the resting length of each voxel. However the shape and volume of each voxel is changed by external forces from the environment and internal forces via behavior. The morphology of a robot is denoted by the  $4 \times 4 \times 3 = 48$ -element vector  $\ell$ , where each element is the resting length of that voxel (with possible values within  $1.0 \pm 0.75$  cm). Like most animals, our robots are bilaterally symmetrical. The lefthand  $2 \times 4 \times 3 = 12$  resting voxel lengths are



**Figure 1. Modeling development.** An evolved soft robot changes its shape during its lifetime (postnatal development), from a walking quadruped into a rolling form. Evolution dictates how a robot’s morphology develops by setting each voxel’s initial ( $\ell_k$ ) and final ( $\ell_k^*$ ) resting length. The length of a single voxel  $k$  is plotted to illustrate its (slower) growth and (faster) actuation processes. Voxel color indicates the current length of that cell: the smallest voxels are blue, medium sized voxels are green, and the largest voxels are red. As robots develop and interact with a physically realistic environment, they generate heterogeneous behavior in terms of instantaneous velocity (bottom arrows). Soft robot evolution, development and physiological functioning can be seen in Supplementary Video S1.

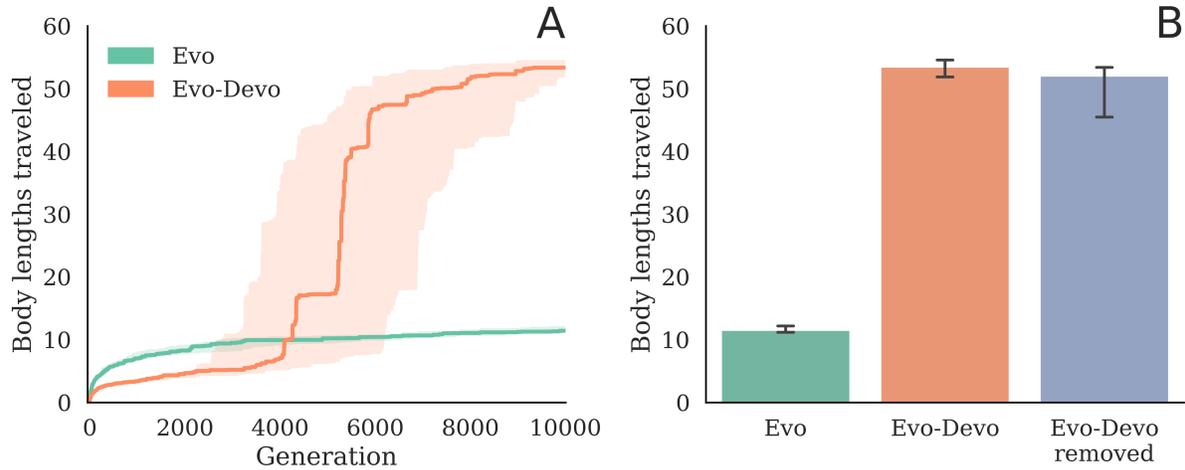
reflected on the other, righthand side of the midsagittal line, yielding 24 independent resting lengths. The controller, however, is not constrained to be symmetrical since many behaviors, even for symmetric morphologies, consist of asymmetric gaits, and is given by the phase offset of each voxel from a global oscillating signal with an amplitude of 0.14 cm. The controller is denoted by the 48-element vector  $\phi$ , where each element is the phase offset of that voxel (with possible values within  $0 \pm \pi/2$ ).

We investigated the impact of development in this model by comparing two experimental variants: Evo and Evo-Devo (schematized in Supplementary Fig. S2). The control treatment, **Evo**, lacks development and therefore maintains a fixed morphology and control policy in a robot as it behaves over its lifetime. Two parameters per voxel are sufficient to specify an evolved robot at any time  $t$  in its lifetime: its morphology  $\ell_k$ , and controller  $\phi_k$ . An evolutionary algorithm optimizes 24 morphological and 48 control parameters.

The experimental treatment **Evo-Devo** evolves a developmental program rather than a static phenotype (Fig. 1). For each parameter in an Evo robot, an Evo-Devo robot has two: its starting and final value. The evolutionary algorithm associated with the Evo-Devo treatment thus optimizes 48 morphological and 96 control parameters. The morphology and controller of the  $k^{\text{th}}$  voxel change linearly from starting to final values, throughout the lifetime of a developing robot. The starting and final points of development are predetermined by a genome which in turn fixes the direction (compression or expansion) and rate of change for each voxel. Development is thus *ballistic* in nature rather than adaptive, as it cannot be influenced by the environment.

For both treatments we conducted 30 independent evolutionary trials. At the end of evolutionary optimization, the non-developmental robots (Evo) tend to move on average with a speed of 10 body lengths in 10 seconds, or 1 length/sec. The evolved and developing robots (Evo-Devo) tend to move at over 5 lengths/sec (Fig. 2A). To ensure evolved and developing robots are not exploiting some unfair advantage conferred by changing body plans and control policies unavailable to non-developmental robots, we manually remove their development by setting  $\ell^* = \ell$  and  $\phi^* = \phi$ , which fixes the structure of their morphologies and controllers at birth ( $t = 0$ ). The resulting reduced robots suffer only a slight (and statistically non-significant) decrease in median speed and still tend to be almost five times faster than the systems evolved without development (Fig. 2B, treatment ‘Evo-Devo removed’). Ballistic development is therefore beneficial for search but does not provide a behavioral advantage in this task environment.

To investigate this apparent search advantage, we trace development and fitness across the 30 lineages which produced a



**Figure 2. Evolvability and development.** Morphological development drastically increases evolvability (A), even when development is manually removed from the optimized systems by setting the final parameter values equal to their starting values ( $\ell_k^* = \ell_k$  and  $\phi_k^* = \phi_k$ ), in each voxel (B). Median fitness is plotted with 95% bootstrapped confidence intervals for three treatments: evolving but non-developmental robots (Evo), evolving and developing robots (Evo-Devo), and evolving and developing robots evaluated at the end of evolution with their development removed (Evo-Devo removed). Fitness of just the final, evolved populations are plotted in B.

‘run champion’: the robot with highest fitness at the termination of a given evolutionary trial (Supplementary Fig. S1). The *developmental window* is defined separately for morphology (Equation 3) and control (Equation 4) as the absolute difference in starting and final values summed across the robot and divided by the total amount of possible development, such that 0 and 1 indicate no and maximal developmental change, respectively. Evo robots by definition have development windows of zero, as do Evo-Devo robots that have had development manually removed. An Evo-Devo robot with a small developmental window has thus become *canalized*<sup>3</sup>. In terms of fitness, there were two observed basins of attraction in average velocity: a slower design type which either trots or gallops at a speed of less than 1 length/sec (Fig. 5A,B and Fig. 3A\*), and a faster design type that rolls at 5-6 lengths/sec (Fig. 5C). After ten thousand generations, 25 out of a total of 30 Evo-Devo trials (83.3%) find the faster design, compared to just 6 out of 30 Evo trials (20%).

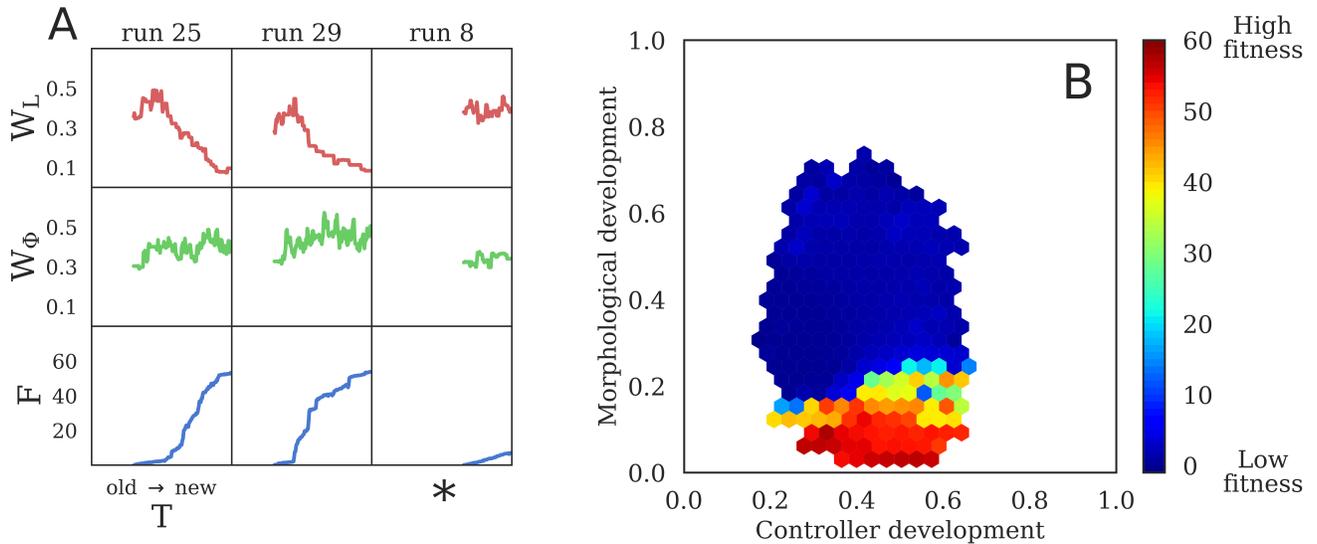
### Differential canalization.

Modular systems are more evolvable than non-modular systems because they allow evolution to improve one subsystem without disrupting others<sup>19,20</sup>. Modularity may be a property of the way a system is built, or it may be an evolved property. The robots evolved here are by definition modular because the genes which affect morphology are independent of those which affect its control. However the more successful Evo-Devo lineages evolved an additional form of modularity, which we term differential canalization: Some initially developmentally plastic traits become integrated and canalized, while other traits remain plastic.

In the successful Evo-Devo trials, morphological traits were canalized while control traits were not. Evidence for this is provided in Supplementary Fig. S1 (which is summarized in Fig. 3A). Trajectories of controller development (green curves) do not follow any discernible pattern and appear upon visual inspection to be consistent with a random walk or genetic drift. The trajectories of morphological development (red curves), however, follow a consistent pattern. The magnitude of morphological development increases slightly, but significantly ( $p < 0.001$ ), before decreasing all the way to the most recent descendant, which is the most fit robot from that trial (the run champion). The morphological development window of the most fit robot is significantly less than the starting morphological development window ( $p < 0.001$ ), but there is no significant difference between starting and final control development windows. Furthermore, this pattern tends to correlate with high fitness: in trials in which this pattern did not appear (runs 6, 8, 16-18), fitness did not increase appreciably over evolutionary time.

This process within the lineages of the run champions is consistent with a more general correlation found in all designs explored during optimization across all runs: Individuals with the highest fitness values tend to have very small amounts of morphological development, while their control policies are free to develop (Fig. 3B). However, despite the fact that morphological development tends to be canalized in the most fit individuals, it cannot simply be discarded as the non-developmental systems have by definition small morphological windows, and small controller windows, but also low fitness.

To test the sensitivity of the canalized morphologies to changes in their control policies, we applied a random series of control mutations to the Evo and Evo-Devo run champions for each of the 30 evolutionary trials. For each run champion, we



**Figure 3. Differential canalization.** Developmental windows (i.e. the total lifetime developmental change) for morphology,  $W_L$  (see Equation 3), and controller,  $W_\phi$  (see Equation 4), alongside fitness  $F$ . (A) Three representative lineages taken from Supplementary Fig. S1, which displays the lineages of all 30 Evo-Devo run champions. Evolutionary time  $T$  moves from the oldest ancestor (left) to the run champion (right). A general trend emerges wherein lineages initially increase their morphological development in  $T$  (rising red curves) and subsequently decrease morphological development to almost zero (falling red curves). Five of the 30 evolutionary trials, annotated by \*, fell into a local optima. (B) Median fitness as a function of morphology and controller development windows ( $W_L$ ,  $W_\phi$ ), for all Evo-Devo designs evaluated. Overall, the fastest designs tend to have small amounts of morphological development, but are free to explore alternative control policies.

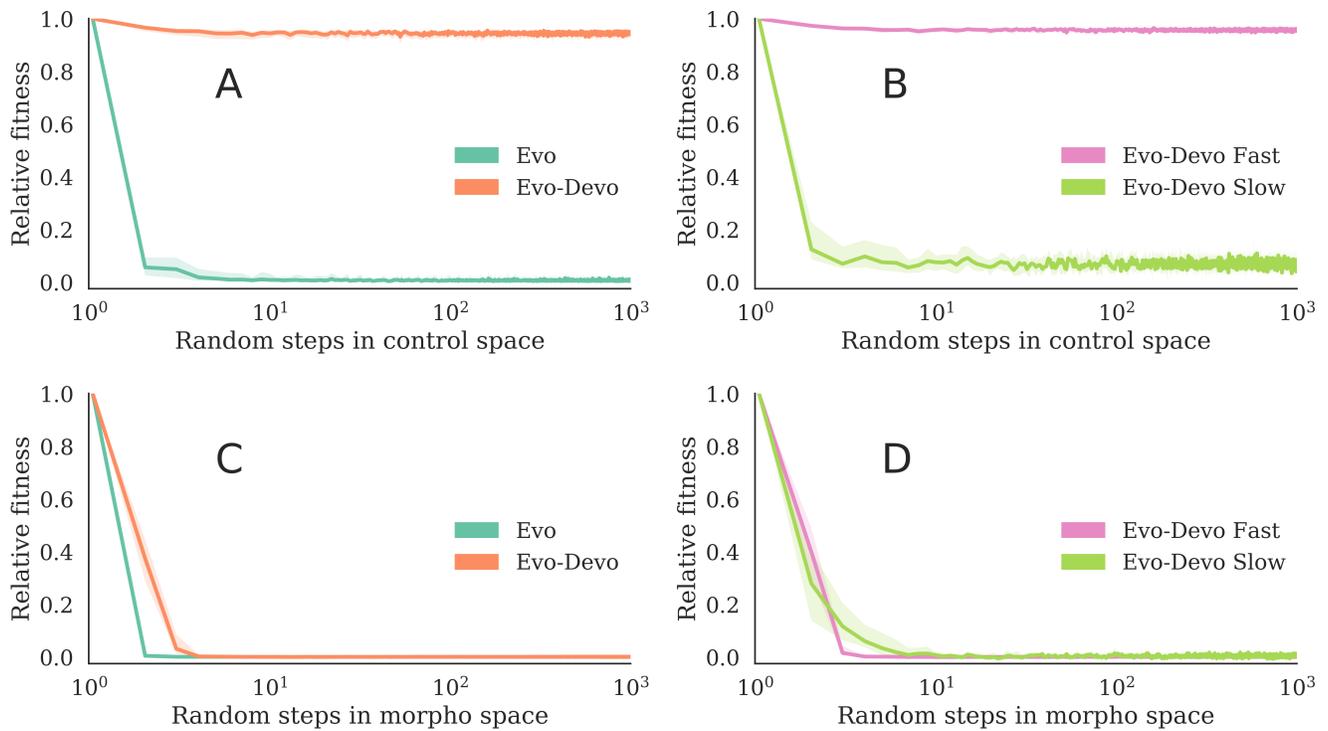
perform 1000 subsequent random controller mutations that build upon each other in series (a Brownian trajectory in the space of controllers)—and repeat this process ten times for each run champion, each with a unique random seed. It was found that optimized Evo-Devo robots tend to possess body plans that are much more robust to control mutations than those of Evo robots (Fig. 4A). The first control mutation to optimized Evo robots tends to immediately render them immobile, whereas optimized Evo-Devo robots tend to retain most of their functionality even after 1000 successive random changes to their controllers. Within Evo-Devo designs, the functionality of the 25 fast designs are minimally affected by changes to their control, whereas the five slow designs also tend to break after the first control mutation (Fig. 4B). Thus it can be concluded that these five robots are non-modular: their non-canalized morphologies evolved a strong dependency on their controllers. The Evo robots are similarly non-modular: they are brittle to control mutations.

To test the sensitivity of the evolved controllers to changes in their morphologies, we applied the same procedure described in the previous paragraph but with random morphological mutations rather than control mutations. It was found that both developmental and non-developmental systems tend to evolve controllers that are very sensitive to morphological mutations (Fig. 4C). The first few morphological mutations to optimized robots, in both treatments, tend to immediately render them immobile. Within Evo-Devo design types, neither of which canalized development in their controllers (Supplementary Fig. S1), both the fast and slow designs possessed controllers sensitive to changes in their morphologies (Fig. 4D). Thus it can be concluded that the non-canalized controllers evolved a strong dependency on their morphologies. The same is true of the non-developmental systems (and is consistent with the findings of Cheney *et al.*<sup>21</sup>). Therefore, the only trait which was successfully canalized was also the only trait which was robust to changes in other traits.

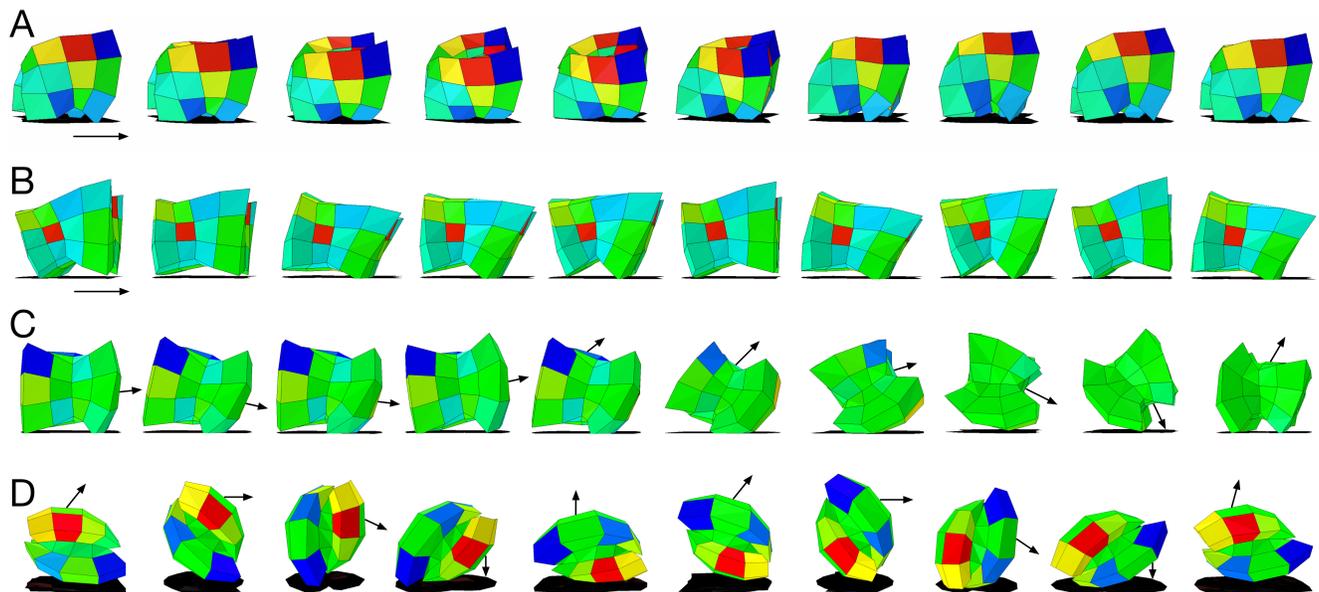
### Heterochrony in morphological development.

The evolutionary algorithm can rapidly discover an actuation pattern that elicits a very small amount of forward movement in these soft robots regardless of the morphology. There is then an incremental path of increasing locomotion speed that natural selection can climb by gradually growing legs to reduce the surface area touching the floor and thus friction, and simultaneously refining controller actuation patterns to better match and exploit the morphology (Fig. 5A,B).

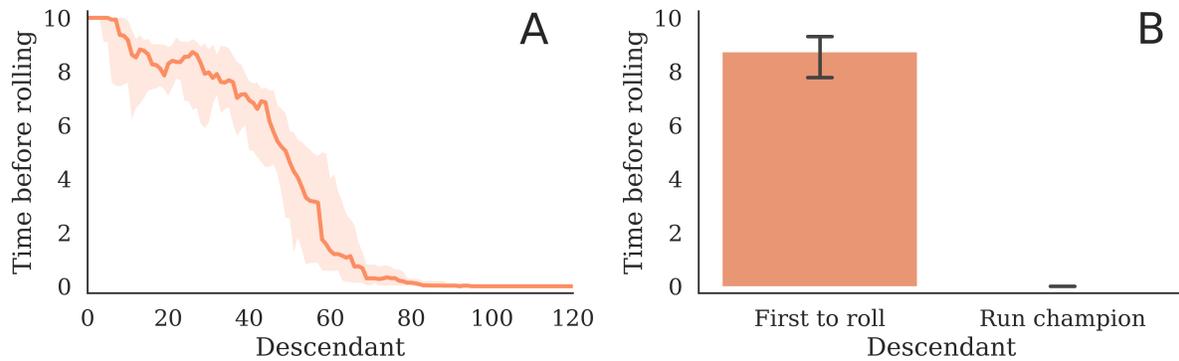
There is, however, a vastly superior design partially hidden from natural selection—a ‘needle in the haystack’, to use Hinton and Nowlan’s metaphor<sup>4</sup>. On flat terrain, rolling can be much faster and more efficient than walking, but finding such a design is difficult because the fitness landscape is deceptive. Rolling over once is much less likely to occur in a random individual



**Figure 4. Sensitivity to morphological and control mutations.** Ten random walks were taken from each run champion. (A) Successive *control* mutations to the Evo and Evo-Devo run champions. (B) The previous Evo-Devo results separately for fast and slow design types. (C) Successive *morphological* mutations to the Evo and Evo-Devo run champions. (D) The previous Evo-Devo results separately for fast and slow design types. Medians plotted with 99% confidence intervals. The faster Evo-Devo robots tend to possess body plans that are robust to control mutations.



**Figure 5. Evolved behavior.** (A) An evolved trotting soft quadruped with a two-beat gait synchronizing diagonal pairs of legs (moving from left to right). (B) A galloping adult robot which goes fully airborne mid-gait. (C) A galloping juvenile robot which develops into a rolling adult form. (D) A rolling juvenile robot at 10 points in ontogeny immediately after birth. Arrows indicate the general directionality of movement, although this is much easier to observe in Supplementary Video S1. Voxel color here indicates the amount of subsequent morphological development remaining at that cell: blue indicates shrinking voxels ( $l_k > l_k^*$ ), red indicates growing voxels ( $l_k < l_k^*$ ), green indicates little to no change either way ( $l_k \cong l_k^*$ ).



**Figure 6. Late onset discoveries.** Ontogenetic time before the discovery of rolling over, taken from the lineages of the best robot from the each of the 25 Evo-Devo trials that produced a rolling design. Median time to discovery, with 95% C.I.s, for (A) the lineage from the most distant ancestor ( $T = 0$ ) to more recent descendants, and (B) the first ancestor to roll over compared to the final run champion. Rolling over is measured from the first time step the top of the robot touches the ground, rather than after completely rolling over. The first ancestors to roll over tend to do so at the end of their lives, their descendants tend to roll sooner in life, and the final run champions all begin rolling immediately at birth. These results are a consequence of dependent time steps: because mutational changes affect all downstream steps, their phenotypic impact is amplified in all but the terminal stages of development. Thus, late onset changes can provide exploration in the search space without breaking rest-of-life functionality, and subsequent evolution can gradually assimilate this trait to the start of development.

than shuffling forwards slightly. And as a population continues to refine walking morphologies and gaits, lineages containing rocking individuals which are close to rolling over, or roll over just once, do not survive long enough to eventually produce a true rolling descendant.

Development can alter the search space evolution operates in because individuals sweep over a continuum of phenotypes, with different velocities, rather than single static phenotype that travels at a constant speed (Supplementary Fig. S2E,J). The lineages which ultimately evolved the faster rolling design initially increased their morphological plasticity in phylogenetic time as evidenced by the initial upward trends in the red curves in Supplementary Fig. S1 (exemplified in Fig. 3A) which contain a statistically significant difference between their starting and maximum developmental window sizes ( $p < 0.001$ ). This exposes evolution to a wider range of body plans and thus increases the chance of randomly rolling at least once at some point during the evaluation period.

The peak of morphological plasticity in Supplementary Fig. S1 (and Fig. 3A) generally lines up with the start of an increasing trend in fitness (blue curves) and marks the onset of differential canalization. Rolling just once allows an individual to move further (1 body length) than some early walking behaviors but they incur the fitness penalty of having fallen over and thus not being able to subsequently walk for the rest of the trial. Therefore this tends to happen at the very end of ontogeny (Fig. 6), as individuals evolve to ‘dive’ in the last few time steps of the simulation of their behavior, thus incurring an additional increase of fitness over their parent, which does not exhibit this behavior. Since more rolling incurs more fitness than less rolling, a form of progenesis occurs as heterochronic mutations move  $\ell_k$  closer to  $\ell_k^*$  for each voxel. This gradually earlifies rolling from a late onset behavior to one that arises increasing earlier in ontogeny (Supplementary Video S1). As more individuals in the population discover and earlify this rolling behavior, the competition stiffens until eventually individuals which are not born rolling from the start are not fast enough to compete (Fig. 5C,D).

### Generality of results.

For the results above, the mutation rate of each voxel was under evolutionary control (self-adaptation). In an effort to assess the generality of our results, we replicated the experiment described above for various fixed mutation rates (Supplementary Fig. S3). The fastest walking behaviors are produced with the lowest mutation rate we tested. The fastest rolling behaviors are also produced by the lowest mutation rates, but higher mutation rates facilitate the discovery of this superior phenotype. Without development, as in Hinton and Nowlan’s case<sup>4</sup>, the search space has a single spike of high fitness. One can not do better than random search in such a space. At the highest mutation rate, optimizing Evo morphologies is random search, and this is the only mutation rate where Evo does not require significantly more generations than Evo-Devo to find the faster design type. This can be observed in Supplementary Fig. S3, for a given mutation rate, at the generations where the slopes of the fitness curves tend to increase dramatically. However, the best two treatments, as measured by the highest median speed at the end of optimization, have development, and they are significantly faster than random search (with and without development) ( $p < 0.01$ ).

## Discussion

In these experiments, the intersection of two time scales—slow linear development and rapid oscillatory actuation, as from a central pattern generator—generates positive and negative feedback in terms of instantaneous velocity: the robot speeds up and slows down during various points in its lifetime (Supplementary Fig. S2J). Prior to canalization, unless all of the phenotypes swept over by an individual in development keep the robot motionless, there will be intervals of relatively superior and inferior performance. Evolution can thus improve overall fitness in a descendant by lengthening the time intervals containing superior phenotypes and reducing the intervals of inferior phenotypes. However, this is only possible if such mutations exist.

We have found here that such mutations do exist in cases where evolutionary changes to one trait do not disrupt the successful behavior contributed by other traits. For example, robots that exhibited the locally optimal trotting behavior (Fig. 5A) exhibited a tight coupling between morphology and control, and thus evolution was unable to canalize development in either one, since mutations to one subsystem tended to disrupt the other. Brief ontogenetic periods of rolling behavior (Fig. 5C), on the other hand, could be temporally extended by evolution through canalization of the morphology alone (Fig. 5D), since these morphologies are generally robust to the pattern of actuation. The key observation here is that only phenotypic traits that render the agent robust to changes in other traits become assimilated, a phenomenon we term differential canalization.

This insight was exposed by modeling the development of simulated robots as they interacted with a physically realistic environment. Differential canalization may be possible in disembodied agents as well, if they conform to appropriate conditions described in Supplementary Discussion.

This finding of differential canalization has important implications for the evolutionary design of artificial and embodied agents such as robots. Computational and engineered systems generally maintain a fixed form as they behave and are evaluated. However, these systems are also extremely brittle when confronted with slight changes in their internal structure, such as damage, or in their external environment such as moving onto a new terrain<sup>22–24</sup>. Indeed, a perennial problem in robotics and AI is finding general solutions which perform well in unseen environments<sup>24,25</sup>. Our results demonstrate how incorporating morphological development in the optimization of robots can reveal, through differential canalization, characters which are robust to internal changes. Robots that are robust to internal changes in their controllers may also be robust to external changes in their environment<sup>12</sup>. Thus, allowing robots to change their structure as they behave might facilitate evolutionary improvement of their descendants, even if these robots will be deployed with static phenotypes or in relatively unchanging environments.

These results are particularly important for the nascent field of soft robotics in which engineers cannot as easily presuppose a robot's body plan and optimize controllers for it because designing such machines manually is unintuitive<sup>17,26</sup>. Our approach addresses this challenge, because differential canalization provides a mechanism whereby static yet robust soft robot morphologies may be automatically discovered using evolutionary algorithms for a given task environment. Furthermore, future soft robots could potentially alter their shape to best match the current task by selecting from previously trained and canalized forms. This change might occur pneumatically, as in Shepherd *et al.*<sup>27</sup>, or it could modulate other material properties such as stiffness (e.g. using a muscular hydrostat).

We have shown that for canalization to occur in our developmental model, some form of paedomorphosis must also occur. However, there are at least two distinct methods in which such heterochrony can proceed: progenesis and neoteny. Progenesis could occur through mutations which move initial parameter values ( $\ell$ ,  $\phi$ ) toward their final values ( $\ell^*$ ,  $\phi^*$ ). Neoteny could instead occur through mutations which move final values ( $\ell^*$ ,  $\phi^*$ ) toward their initial values ( $\ell$ ,  $\phi$ ). Although a superior phenotype can materialize anywhere along the ontogenetic timeline, late onset mutations are less likely to be deleterious than early onset mutations. This is because our developmental model is linear in terms of process, and interfering with any step affects all temporally-downstream steps. Since the probability of a mutation being beneficial is inversely proportional to its phenotypic magnitude<sup>28</sup>, mutational changes in the terminal stages of development require the smallest change to the developmental program. Hence, late-onset discoveries of superior traits are more likely to occur without breaking functionality at other points in ontogeny, and these traits can become canalized by evolution through progenesis: mutations which reduce the amount of ontogenetic time prior to realizing the superior trait (by moving  $\ell \rightarrow \ell^*$  and/or  $\phi \rightarrow \phi^*$ ). Indeed progenesis was observed most often in our trials (Fig. 6): late onset mutations which transform a walking robot into a rolling one are discovered by the evolutionary process, and are then moved back toward the birth of the robots' descendants through subsequent mutations.

Finally, we would like to note the observed phenomenon of *increased* phenotypic plasticity prior to genetic assimilation. Models of the Baldwin effect usually assume that phenotypic plasticity itself does not evolve, although it has been shown how major changes in the environment can select for increased plasticity in a character that is initially canalized<sup>29</sup>. In our experiments however, there is no environmental change. There is also a related concept known as 'sensitive periods' of development in which an organism's phenotype is more responsive to experience<sup>30</sup>. Despite great interest in sensitive periods, the adaptive reasons why they have evolved are unclear<sup>31</sup>. In our model, increasing the amount of morphological development increases the chance of capturing an advantageous static phenotype, which can then be canalized, once found. There is a balance, however, as a phenotype will not realize the globally optimal solution by simply maximizing development. This would merely produce a cube of maximum volume shrinking to a cube of minimum volume, or vice versa.

The developmental model described herein is intentionally minimalistic in order to isolate the effect of morphological and neurological change in the evolutionary search for embodied agents. The simplifying assumptions necessary to do so make it difficult to assess the biological implications. For example, we model development as an open loop process and thus ignore environmental queues and sensory feedback<sup>32,33</sup>. We also disregard the costs and constraints of phenotypic plasticity<sup>34,35</sup>. However, by removing these confounding factors, these results may help generate novel hypotheses about morphological development, heterochrony, modularity and evolvability in biological systems.

## Methods

### Resting length.

Voxelyze<sup>18</sup> simulates soft materials as a lattice of spring-like beam elements and point masses. Voxels are centered on these point masses and their individual properties depend on attributes of the beams connecting them to neighboring voxels. By default, prior to environmental or behavioral forces, all (resting) beam lengths are equal, and thus voxels are rendered as cubes of equal size (see for example Cheney *et al.*<sup>36</sup>). In these experiments, however, we alter the resting shape of the voxels by changing the underlying beam lengths. For convenience we refer to voxels by their resting ‘length’, but the actual resting shape of a single voxel depends on its neighbors.

### Ballistic development.

Ballistic development  $\beta(t) = a + t(b - a)/\tau$ , is simply a linear function from  $a$  to  $b$  in ontogenetic time  $t \in (0, \tau)$ . For Evo robots,  $a = b$ , and hence  $\beta(t) = a$ , which is just a constant value in ontogenetic time.

### Current length.

For smaller voxels, it is necessary to implement damping into their actuation to avoid simulation instability. Actuation amplitude is limited by a linear damping factor  $\xi(x) = \min\{1, (4x - 1)/3\}$ , which only affects voxels with resting length less than one centimeter, and approaches zero (no actuation) as the resting length goes to its lower bound of 0.25 cm.

Actuation of the  $k^{\text{th}}$  voxel  $\psi_k(t)$  therefore depends on the starting and final phase offsets  $(\phi_k, \phi_k^*)$  for relative displacement, and on the starting and final resting lengths  $(\ell_k, \ell_k^*)$  for amplitude. With maximum amplitude  $A = 0.14$  cm and a fixed frequency  $f = 4$  Hz, we have

$$\psi_k(t) = A \cdot \sin \left[ 2\pi f t + \phi_k + \frac{t(\phi_k^* - \phi_k)}{\tau} \right] \cdot \xi \left[ \ell_k + \frac{t(\ell_k^* - \ell_k)}{\tau} \right] \quad (1)$$

The current length of the  $k^{\text{th}}$  voxel at time  $t$ , which we denote by  $\mathcal{L}_k(t)$ , is the resting length plus the offset and damped signal  $\psi_k(t)$ .

$$\mathcal{L}_k(t) = \ell_k + \frac{t(\ell_k^* - \ell_k)}{\tau} + \psi_k(t) \quad (2)$$

Current length is broken down into its constituent parts for a single voxel, under each treatment, in Supplementary Fig. S2.

### Evolutionary algorithm.

We employ a standard evolutionary algorithm, Age-Fitness-Pareto Optimization<sup>37</sup>, which uses the concept of Pareto dominance and an objective of age (in addition to fitness) intended to promote diversity among candidate designs. For 30 runs, a population of 30 robots is evolved for 10000 generations. Every generation, the population is first doubled by creating modified copies of each individual in the population. The age of each individual is then incremented by one. Next, an additional random individual (with age zero) is injected into the population. Finally, selection reduces the population down to its original size according to the two objectives of fitness (maximized) and age (minimized).

The mutation rate is also evolved for each voxel, by maintaining a separate vector of mutation rates which are slightly modified every time a genotype is copied from parent to child. These 48 independent mutation rates are initialized such that only a single voxel is mutated on average. Mutations follow a normal distribution ( $\sigma_\ell = 0.75$  cm,  $\sigma_\phi = \pi/2$ ) and are applied by first selecting what parameter types to mutate  $(\phi_k, \phi_k^*, \ell_k, \ell_k^*)$ , and then choosing, for each parameter separately, which voxels to mutate. In Supplemental Materials we provide exact derivations of the expected genotypic impact of mutations in terms of the proportions of voxels and parameters modified for a given fixed mutation rate  $\lambda$ . There is a negligible difference between Evo and Evo-Devo in terms of the expected number of parent voxels modified during mutation (Supplementary Fig. S4).

### Developmental windows.

The amount of development in a particular voxel can range from zero (in the case that starting and final values are equal) to 1.5 cm for the morphology (which ranges from 0.25 cm to 1.75 cm) and  $\pi$  for the controller (which ranges from  $-\pi/2$  to  $\pi/2$ ). The morphological development window,  $W_L$ , is the sum of the absolute difference of starting and final resting lengths across the robot's 48 voxels, divided by the total amount of possible morphological development.

$$W_L = \frac{1}{48(1.5)} \sum_{k=1}^{48} |\ell_k^* - \ell_k| \quad (3)$$

The controller development window,  $W_\Phi$ , is the sum of the absolute difference of starting and final phase offsets across the robot's 48 voxels, divided by the total amount of possible controller development.

$$W_\Phi = \frac{1}{48\pi} \sum_{k=1}^{48} |\phi_k^* - \phi_k| \quad (4)$$

### Statistical analysis.

We use the Mann-Whitney  $U$  test to assess statistical significance.

### Source code.

<https://github.com/skriegman/how-devo-can-guide-evo> contains the source code necessary for reproducing the results reported in this paper.

### References

1. Baldwin, J. M. A new factor in evolution. *The american naturalist* **30**, 441–451 (1896).
2. Morgan, C. L. On modification and variation. *Sci.* **4**, 733–740 (1896).
3. Waddington, C. H. Canalization of development and the inheritance of acquired characters. *Nat.* **150**, 563–565 (1942).
4. Hinton, G. E. & Nowlan, S. J. How learning can guide evolution. *Complex systems* **1**, 495–502 (1987).
5. Floreano, D. & Mondada, F. Evolution of plastic neurocontrollers for situated agents. In *From Animals to Animats 4, Proceedings of the 4th International Conference on Simulation of Adaptive Behavior (SAB 1996)*, LIS-CONF-1996-001, 402–410 (MIT Press, 1996).
6. Dellaert, F. & Beer, R. D. A developmental model for the evolution of complete autonomous agents. In *Proceedings of the fourth international conference on simulation of adaptive behavior* (1996).
7. Eggenberger, P. Evolving morphologies of simulated 3D organisms based on differential gene expression. *Procs. Fourth Eur. Conf. on Artif. Life* 205–213 (1997).
8. Bongard, J. C. & Pfeifer, R. Repeated structure and dissociation of genotypic and phenotypic complexity in Artificial Ontogeny. *Proc. The Genet. Evol. Comput. Conf. (GECCO 2001)* 829–836 (2001).
9. Doursat, R. Organically grown architectures: Creating decentralized, autonomous systems by embryomorphic engineering. In *Organic computing*, 167–199 (Springer, 2009).
10. Ventrella, J. Designing emergence in animated artificial life worlds. In *Virtual Worlds*, 143–155 (Springer, 1998).
11. Komosinski, M. The framsticks system: versatile simulator of 3d agents and their evolution. *Kybernetes* **32**, 156–173 (2003).
12. Bongard, J. C. Morphological change in machines accelerates the evolution of robust behavior. *Proc. Natl. Acad. Sci.* **108**, 1234–1239 (2011). DOI 10.1073/pnas.1015390108.
13. Kriegman, S., Cheney, N., Corucci, F. & Bongard, J. C. A minimal developmental model can increase evolvability in soft robots. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 131–138 (ACM, 2017).
14. Cheney, N., Bongard, J. C. & Lipson, H. Evolving soft robots in tight spaces. In *Proceedings of the 2015 annual conference on Genetic and Evolutionary Computation*, 935–942 (ACM, 2015).
15. Miriyev, A., Stack, K. & Lipson, H. Soft material for soft actuators. *Nat. Commun.* **8** (2017).
16. Wehner, M. *et al.* An integrated design and fabrication strategy for entirely soft, autonomous robots. *Nat.* **536**, 451–455 (2016).

17. Lipson, H. Challenges and opportunities for design, simulation, and fabrication of soft robots. *Soft Robotics* **1**, 21–27 (2014). DOI 10.1089/soro.2013.0007.
18. Hiller, J. & Lipson, H. Dynamic simulation of soft multimaterial 3d-printed objects. *Soft Robotics* **1**, 88–101 (2014).
19. Wagner, G. P. & Altenberg, L. Perspective: complex adaptations and the evolution of evolvability. *Evol.* **50**, 967–976 (1996).
20. Lipson, H. Principles of modularity, regularity, and hierarchy for scalable systems. *The J. Biol. Phys. Chem.* **7**, 125–128 (2007).
21. Cheney, N., Bongard, J., SunSpiral, V. & Lipson, H. Scalable co-optimization of morphology and control in embodied machines. *arXiv preprint arXiv:1706.06133* (2017).
22. French, R. M. Catastrophic forgetting in connectionist networks. *Trends cognitive sciences* **3**, 128–135 (1999).
23. Szegedy, C. *et al.* Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199* (2013).
24. Carlson, J. & Murphy, R. R. How ugvs physically fail in the field. *IEEE Transactions on robotics* **21**, 423–437 (2005).
25. Nguyen, A., Yosinski, J. & Clune, J. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 427–436 (2015).
26. Pfeifer, R., Lungarella, M. & Iida, F. The challenges ahead for bio-inspired ‘soft’ robotics. *Commun. ACM* **55**, 76–87 (2012).
27. Shepherd, R. F. *et al.* Multigait soft robot. *Proc. Natl. Acad. Sci.* **108**, 20400–20403 (2011).
28. Fisher, R. A. *The genetical theory of natural selection* (Oxford University Press, 1930).
29. Lande, R. Adaptation to an extraordinary environment by evolution of phenotypic plasticity and genetic assimilation. *J. evolutionary biology* **22**, 1435–1446 (2009).
30. Bateson, P. How do sensitive periods arise and what are they for? *Animal Behav.* **27**, 470–486 (1979).
31. Fawcett, T. W. & Frankenhuis, W. E. Adaptive explanations for sensitive windows in development. *Front. Zool.* **12**, S3 (2015). DOI 10.1186/1742-9994-12-S1-S3.
32. Moczek, A. P. *et al.* The role of developmental plasticity in evolutionary innovation. *Proc. Royal Soc. Lond. B: Biol. Sci.* **278**, 2705–2713 (2011).
33. Snell-Rood, E. C. An overview of the evolutionary causes and consequences of behavioural plasticity. *Animal Behav.* **85**, 1004–1011 (2013).
34. Snell-Rood, E. C. Selective processes in development: Implications for the costs and benefits of phenotypic plasticity. *Integr. & Comp. Biol.* **52** (2012).
35. Murren, C. J. *et al.* Constraints on the evolution of phenotypic plasticity: limits and costs of phenotype and plasticity. *Hered.* **115**, 293–301 (2015).
36. Cheney, N., MacCurdy, R., Clune, J. & Lipson, H. Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, 167–174 (ACM, 2013).
37. Schmidt, M. & Lipson, H. Age-fitness pareto optimization. In *Genetic Programming Theory and Practice VIII*, 129–146 (Springer, 2011).

## Acknowledgments

This work was supported by NSF awards PECASE-0953837 and INSPIRE-1344227, as well as ARO contract W911NF-16-1-0304. We thank Shawn Beaulieu, Francesco Corucci, Chris Fusting and Frank Veenstra for useful discussions.

## Supplementary Video

### Supplementary Video S1.

<https://youtu.be/nWbpegOCeQY> provides a brief overview of the results reported in this paper (2:22 minutes).

### Supplementary Video S2.

<https://youtu.be/N8AThp4sZZ0> presents the paper in a nutshell (18 seconds): An unfit robot (1) evolves into a walker (2). Its descendant ‘discovers’ rolling late in life (3). A mutation causes rolling to show up earlier in its descendant (4). The final descendant rolls throughout its life (5).

## Supplementary Discussion

### Embodiment.

We consider an agent to be embodied if its output affects its input. This relationship between an agent and its environment may be modeled as a partially observable Markov decision process (POMDP), where:

- $S$  is a set of states,
- $A$  is a set of actions, and
- $T$  is a set of state transition probabilities.

At each time period, the environment is in some state  $s \in S$ . The agent takes an action  $a \in A$ , which causes the environment to transition to state  $s'$  with probability  $T(s' | s, a)$ . A control policy  $\pi$  specifies the action  $a = \pi(b)$  a robot will take next given its current belief  $b$  about the new state of the environment. Usually, the goal is to find, given  $T$ , the optimal policy  $\pi^*$  which yields the highest overall performance.

The morphology of an agent can be considered as part of  $T$  since it mediates the transition between  $s$  and  $s'$ , yet it is not part of  $\pi$ . By changing morphology we were able to modify a subset of  $T$  in such a way that facilitated the search for  $\pi^*$ . This might also be possible in disembodied agents if other dimensions of  $T$  can be changed by some search process such as to facilitate the search for  $\pi^*$ .

## Supplementary Methods

### Mutations.

Mutations are applied by first choosing what material properties (i.e. voxel-level parameters such as resting length and phase offset) to mutate, and then choosing, separately for each property, which voxels to modify them in. For each of the  $n$  material properties, we select it with independent probability  $p = 1/n$ . If none are selected, we randomly choose one. This occurs with probability  $(1 - p)^n$ .

Hence the number of selected material properties for mutation is a random variable  $S$  which follows a binomial distribution truncated on  $S \geq 1$  such that the entire untruncated probability mass at  $S = 0$  is placed on top of  $S = 1$ .

$$\Pr(S = s | n) = \begin{cases} 0 & \text{for } s = 0 \\ np(1-p)^{n-1} + (1-p)^n & \text{for } s = 1 \\ \binom{n}{s} p^s (1-p)^{n-s} & \text{for } s > 1 \end{cases} \quad (5)$$

The expected number of selected material properties is then

$$\mathbb{E}(S) = np(1-p)^{n-1} + (1-p)^n + \sum_{s=2}^n s \binom{n}{s} p^s (1-p)^{n-s} \quad (6)$$

$$= (1-p)^n + \sum_{s=1}^n s \binom{n}{s} p^s (1-p)^{n-s} \quad (7)$$

$$= (1-p)^n + np \quad (8)$$

$$= (1-p)^n + 1 \quad (9)$$

The selected material property of each voxel is mutated independently with probability  $\lambda$ , and thus the expected *number* of genotype elements mutated given  $K$  total voxels is

$$\delta_{\text{gene}} = \lambda K \cdot \mathbb{E}(S) \quad (10)$$

Dividing by the length of the genome,  $nK$ , the expected *proportion* of genotype elements mutated is

$$\pi_{\text{gene}} = \lambda/n \cdot \mathbb{E}(S) \quad (11)$$

Note that bilaterally symmetrical properties have the same expected values since they are half the size but a mutation effects two voxels.

We have  $K = 48$  total voxels, and  $n = \{2, 4\}$  material properties for our two main experimental treatments {Evo, Evo-Devo}, respectively:

	$n = 2$	$n = 4$
$\delta_{\text{gene}}$	$60\lambda$	$63.8175\lambda$
$\pi_{\text{gene}}$	$0.625\lambda$	$0.3291\lambda$

The expected number of *voxels* mutated is lower than the average number of genotype elements mutated because there can be overlap/redundancies among the voxels selected between the material properties. To calculate the average number of voxels mutated we need to consider a hierarchy of binomial distributions.

The number of material properties mutated within the  $k^{\text{th}}$  voxel given  $S$  selected,  $M|S$ , follows  $\text{binom}(S, \lambda) \forall k$ .

$$\Pr(M = m | S, \lambda) = \binom{S}{m} \lambda^m (1 - \lambda)^{S-m} \quad (12)$$

For convenience, let's denote the probability that at least one mutation occurs to the  $k^{\text{th}}$  voxel as  $\theta$ .

$$\theta = \Pr(M > 0 | S, \lambda) = 1 - (1 - \lambda)^S \quad (13)$$

Then the number of voxels mutated,  $V$ , across a total of  $K$  voxels and  $S$  selected material properties, follows  $\text{binom}(K, \theta)$ .

$$\Pr(V = v | S, K, \lambda, n) = \binom{K}{v} \theta^v (1 - \theta)^{K-v} \quad (14)$$

And the expected number of voxels mutated (out of  $K$  total) is

$$\delta_{\text{vox}} = \mathbb{E}(V | K, \lambda, n) \quad (15)$$

$$= \mathbb{E}_S \mathbb{E}_V(V | S, K, \lambda, n) \quad (16)$$

$$= \mathbb{E}_S(K\theta | S, K, \lambda, n) \quad (17)$$

$$= K \{1 - \mathbb{E}_S[(1 - \lambda)^S | \lambda, n]\} \quad (18)$$

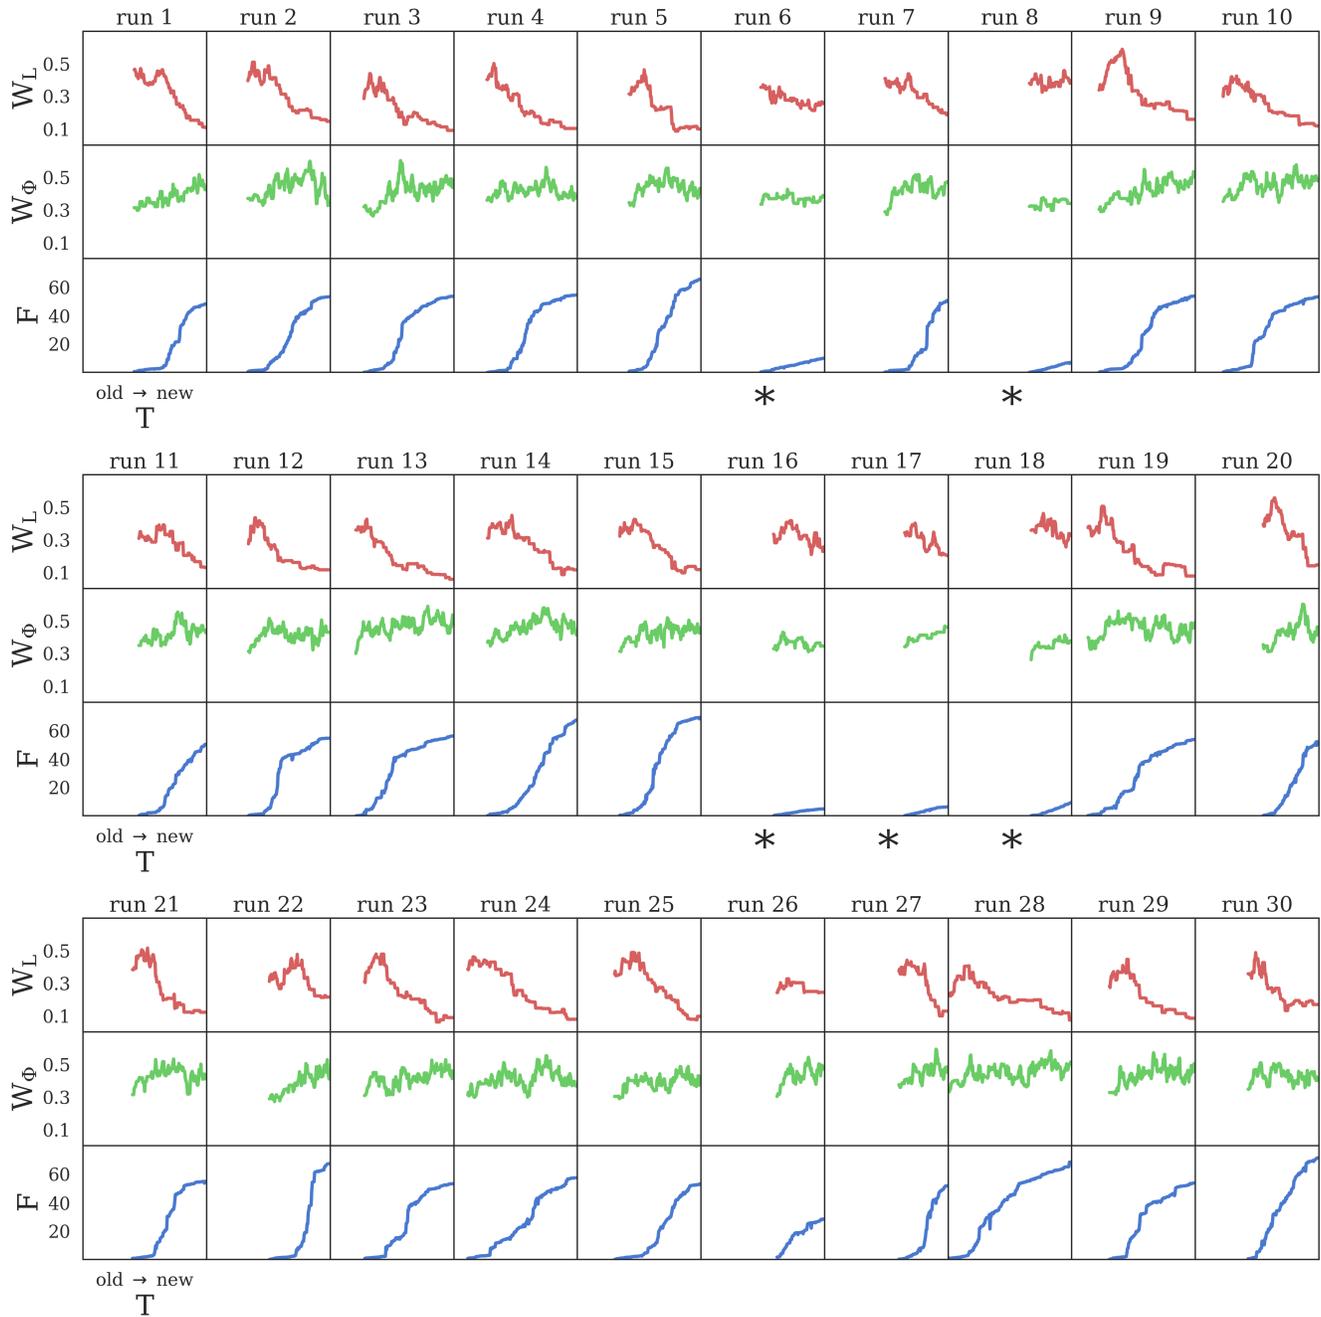
$$= K \left\{ 1 - \left[ (1 - \lambda)(1 - p)^n + \sum_{s=1}^n (1 - \lambda)^s \binom{n}{s} p^s (1 - p)^{n-s} \right] \right\} \quad (19)$$

$$= K \left\{ 1 - (1 - p)^n \left[ \left( \frac{\lambda p - 1}{p - 1} \right)^n - \lambda \right] \right\} \quad (20)$$

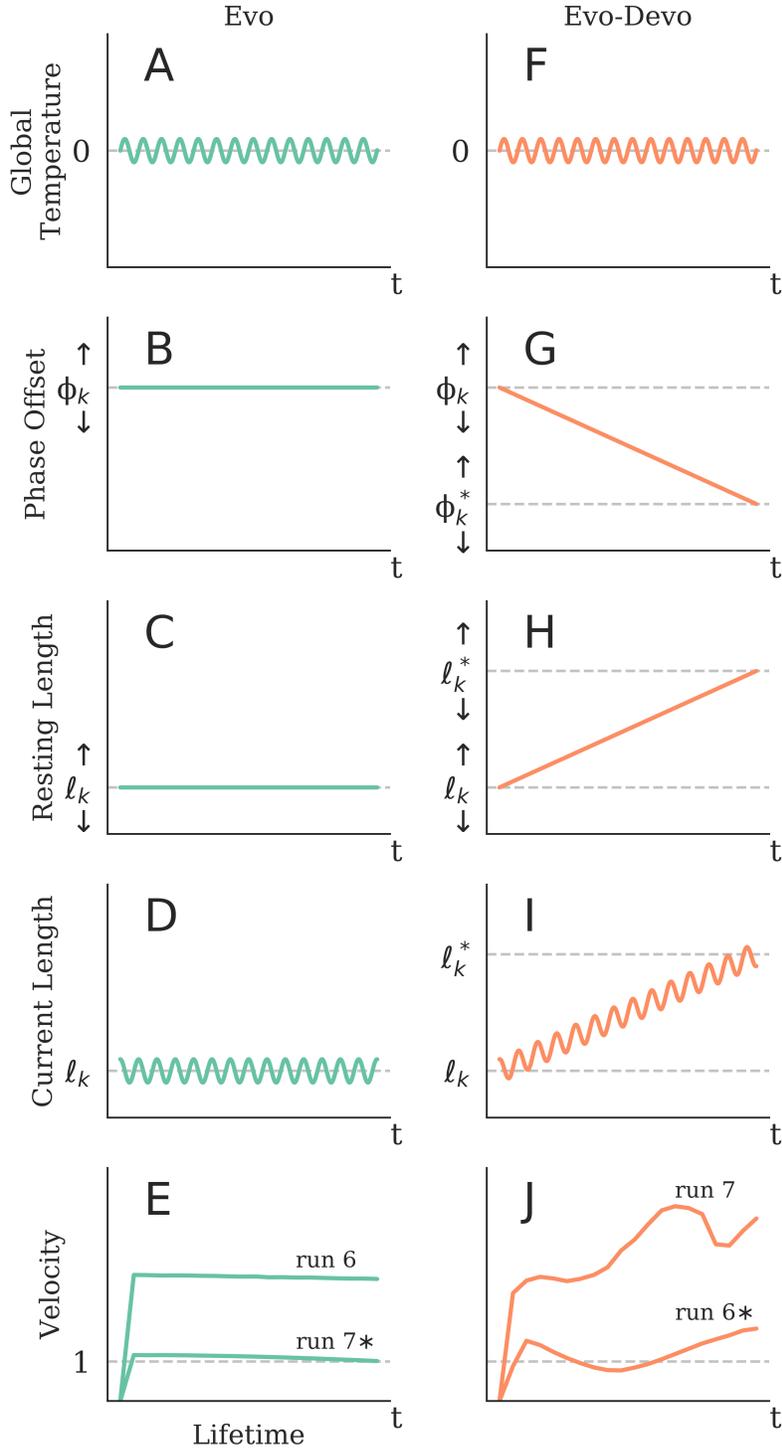
There is an extremely tight bound on the proportion of voxels mutated,  $\pi_{\text{vox}} = \delta_{\text{vox}}/K$ , for any  $n > 1$  (Supplementary Fig. S4). Thus, mutations in Evo ( $n = 2$ ) and Evo-Devo ( $n = 4$ ) have practically the same genotypic impact in terms of the number of voxels modified. For completeness, the following table displays  $\delta_{\text{vox}}$  for the specific values of  $\lambda$  considered by our hyperparameter sweep ( $K = 48$ ) (Supplementary Fig. S3).

		$\lambda$							
		1/48	2/48	4/48	8/48	16/48	24/48	32/48	48/48
$n$	2	1.25	2.48	4.92	9.67	18.67	27	34.67	48
	4	1.3	2.6	5.14	10.05	19.17	27.46	34.98	48

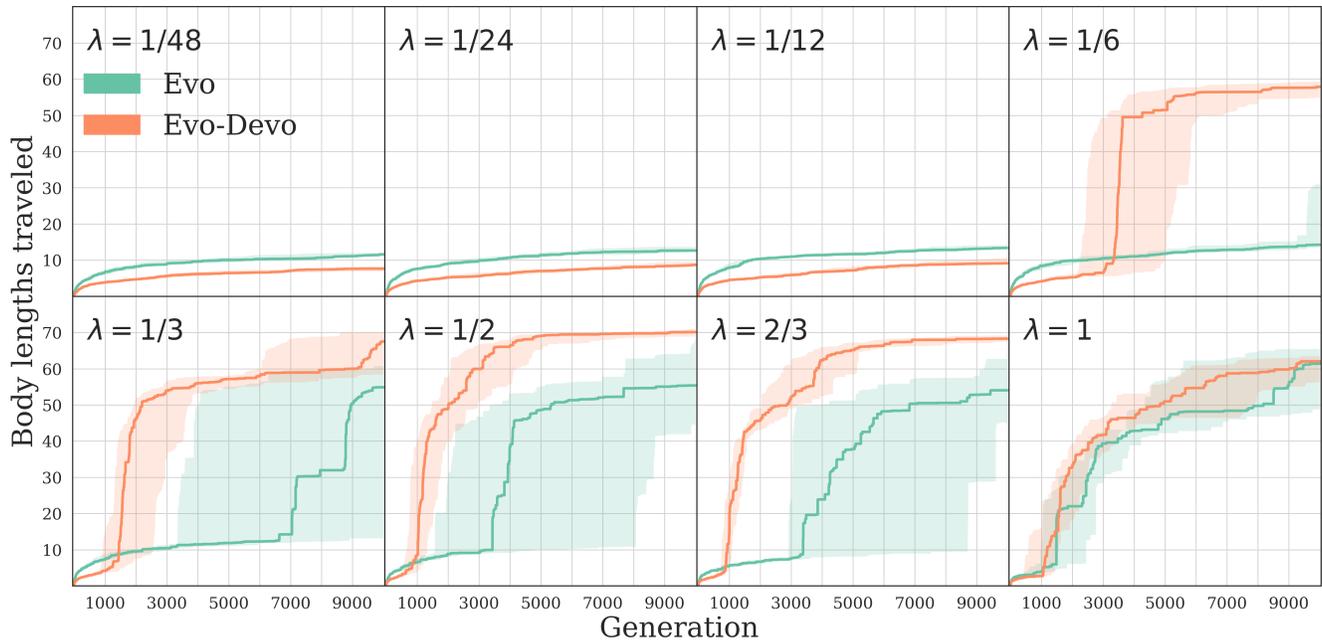
## Developmental Windows



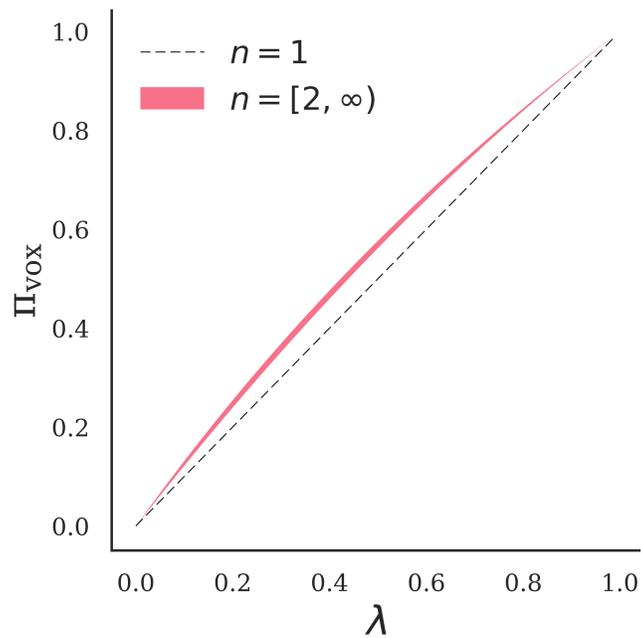
**Figure S1. Evolutionary change during 30 Evo-Devo trials.** The amount of morphological development,  $W_L$  (see Equation 3), controller development,  $W_\Phi$  (see Equation 4), and fitness,  $F$ , for the lineages of the 30 Evo-Devo run champions. Evolutionary time,  $T$ , moves from the oldest ancestor (left) to the run champion (right). A general trend emerges wherein lineages initially increase their morphological development in  $T$  (rising red curves) and subsequently decrease morphological development to almost zero (falling red curves). Five of the 30 evolutionary trials, annotated by \*, fell into a local optima.



**Figure S2. Experimental treatments.** The phase of an oscillating global temperature (A, F) is offset in the  $k^{\text{th}}$  voxel by a linear function from  $\phi_k$  to  $\phi_k^*$  (B, G). The resting length of the  $k^{\text{th}}$  voxel is a linear function from  $l_k$  to  $l_k^*$  (C, H). For Evo, there is no development, so  $\phi_k = \phi_k^*$  and  $l_k = l_k^*$ . The offset actuation is added on top of the resting length to give the current length of the  $k^{\text{th}}$  voxel (D, I). These example voxel-level changes occur across ontogenetic time ( $t$ ), independently in each of the 48 voxels, and together interact with the environment to generate robot-level velocity (E, J). Velocities are averaged across intervals of two actuation cycles.



**Figure S3. Mutation rate sweep.** Median fitness (with 95% bootstrapped confidence intervals) under various mutation rates,  $\lambda$ , a hyperparameter defined in Supplemental Materials which affects the probability a voxel is mutated. In the main experiment of this paper, the mutation rate is evolved for each voxel independently, and is constantly changing. In this mutation rate sweep,  $\lambda$  is held uniform across all voxels. The fastest walking and rolling behaviors are produced with the lowest mutation rate we tested ( $\lambda = 1/48$ ), although higher mutation rates facilitate the discovery of rolling which is much faster than walking. Without development, the search space has a single spike of high fitness corresponding to this rolling behavior. One can not do better than random search in such a space. When  $\lambda = 1$ , optimizing Evo morphologies is random search, and this is the only mutation rate where Evo does not require significantly more generations than Evo-Devo to find the faster design type. This can be observed for  $\lambda \in \{1/6, 1/3, 1/2, 2/3, 1\}$  at the generations where the slopes of the fitness curves tend to increase dramatically. However, the best two treatments (Evo-Devo at  $\lambda = 1/2$  and  $\lambda = 2/3$ ), as measured by the highest median speed at the end of optimization, have development, and they are significantly faster than random search (with and without development) ( $p < 0.01$ ).



**Figure S4. Mutational impact.** The expected proportion of voxels modified,  $\pi_{\text{vox}}$ , where  $n$  is the number of material properties per voxel (voxel-level parameters such as  $\ell_k$  and  $\phi_k$ ) and  $\lambda$  is the mutation rate (the probability of mutating a voxel for a selected material property). A derivation is provided in Supplemental Materials. In this paper, we have treatments Evo, with  $n = 2$ , and Evo-Devo, with  $n = 4$ . There is an extremely tight bound on the proportion of voxels mutated for any  $n > 1$ . At  $\lambda = 1$  every voxel must be mutated while at  $\lambda = 0$  there can be no voxels mutated. The arch between these two points is limited by the possibility of overlap (selecting the same voxel more than once).