

# The Utility of Evolving Simulated Robot Morphology Increases with Task Complexity for Object Manipulation

---

Josh Bongard\*  
University of Vermont

---

## Keywords

Evolutionary robotics, embodied artificial intelligence, embodied cognition, robot shaping, multi-objective optimization

**Abstract** Embodied artificial intelligence argues that the body and brain play equally important roles in the generation of adaptive behavior. An increasingly common approach therefore is to evolve an agent's morphology along with its control in the hope that evolution will find a good coupled system. In order for embodied artificial intelligence to gain credibility within the robotics and cognitive science communities, however, it is necessary to amass evidence not only for *how* to co-optimize morphology and control of adaptive machines, but *why*. This work provides two new lines of evidence for why this co-optimization is useful: Here we show that for an object manipulation task in which a simulated robot must accomplish one, two, or three objectives simultaneously, subjugating more aspects of the robot's morphology to selective pressure allows for the evolution of better robots as the number of objectives increases. In addition, for robots that successfully evolved to accomplish all of their objectives, those composed of evolved rather than fixed morphologies generalized better to previously unseen environmental conditions.

---

## I Introduction

Embodied cognitive science [20, 22, 15] and its sister discipline embodied artificial intelligence [13, 33, 32] stress the importance of the body and its interaction with the environment for achieving adaptive behavior. This suggests that the body plan for an adaptive machine should be chosen carefully to increase the probability of finding an optimized controller such that they together produce useful behavior. It has been argued that designing the mechanical structure for a machine is more intuitive and therefore better left to human designers than designing the controller, this latter being a non-intuitive process and therefore better shaped by automated optimization [30, p. 22]. However, here we show that even a seemingly appropriate configuration for a robotic arm—an anthropomorphic design—is increasingly improved by placing more aspects of the body plan under evolutionary control as the task that the robot must accomplish becomes more complex.

One approach to morphological design is to adopt body plans inspired by biological organisms [6, 44], but often the specific details of the robot's niche are subtly (or grossly) different from the niche

---

\* Department of Computer Science, Vermont Advanced Computing Center, 33 Colchester Ave., Burlington VT 05408. E-mail: josh.bongard@uvm.edu

of the organism and therefore not an ideal design choice. An alternative approach is to automate the optimization of both morphology and controller simultaneously (e.g., [39, 25, 11]), thereby searching for a coupled dynamical system (between the moving body parts and the updating controller) that matches the dynamics of the environment well to produce useful behavior [32, 5].

There is a growing body of work demonstrating different ways for *how* to optimize morphology and control: direct methods in which each part of the genome corresponds to only one part of the body plan or neural controller (e.g., [25, 8]), recursive methods in which each part of the genome may correspond to one or more phenotypic components (e.g., [39, 43, 1, 23]), and more biologically inspired methods in which morphology and control are shaped by simulations of biological growth processes (e.g., [17, 19, 7, 11, 35, 41, 28]). Further, such approaches can be divided into those that optimize parameters of a fixed body plan and those that optimize the body plan itself. The current work falls into the former category.

However, there is relatively little evidence about *why* morphology and control should be co-optimized, beyond the common, rather vague argument that evolution can find a better fit between morphology and control if both are optimized together. It has been demonstrated that indirect methods may produce phenotypes that grow in complexity without requiring a corresponding increase in genotypic complexity [9], and that control can be simplified if an appropriate morphology is found [29, 24]. More recently it has been shown that a careful choice of morphology allows parts of the body to functionally specialize, rather than its being assumed a priori which parts of the body will perform which functions [2].

Here we contribute two other lines of evidence about why coevolving morphology and control is important: In the object manipulation task investigated here, evolving more aspects of the robot's body plan along with control increased the probability of discovering a successful solution as the task complexity increased. Also, among those robots that evolved to successfully perform the task, those with more evolved morphological aspects were more robust than those with less.

The next section introduces the methods used to investigate how evolvability and robustness are affected by evolving morphological aspects of a simulated robot. Section 3 reports results for various experimental regimes in which an increasing amount of the robot's body plan is evolved. The final section provides some discussion, concluding remarks, and directions for future research.

## 2 Methods

This section introduces the simulated robot and its task environment, which aspects of its morphology and controller can be optimized and which are fixed, and finally the algorithm used to optimize the robots' object manipulation abilities.

### 2.1 Task Environment

The class of robot investigated here is an articulated arm and hand equipped with either three, four, or five fingers (e.g., Figure 1).

The task is to manipulate a series of objects that are placed within the robot's reach. Object manipulation is here defined as a general behavioral competence that may be composed of one or more of the following objectives: grasping, lifting, and active categorical perception [3, 31, 42]. Task complexity in this work is therefore defined as the number of objectives that a robot must accomplish simultaneously. Grasping requires the robot to minimize the distance between its fingers and the object, lifting requires it to maximize the vertical displacement of the object, and active perception requires it to interact with objects of different shapes in order to distinguish between them. Object manipulation has been previously studied in robotics (e.g., [14, 27, 21, 46, 42]), and methods for enabling a robot to perform multiple behaviors simultaneously and in sequence constitute a popular area of study (e.g., [46]). Here, however, the goal is not to demonstrate that a robot can accomplish these tasks, but rather the task domain is used to show that there is a positive

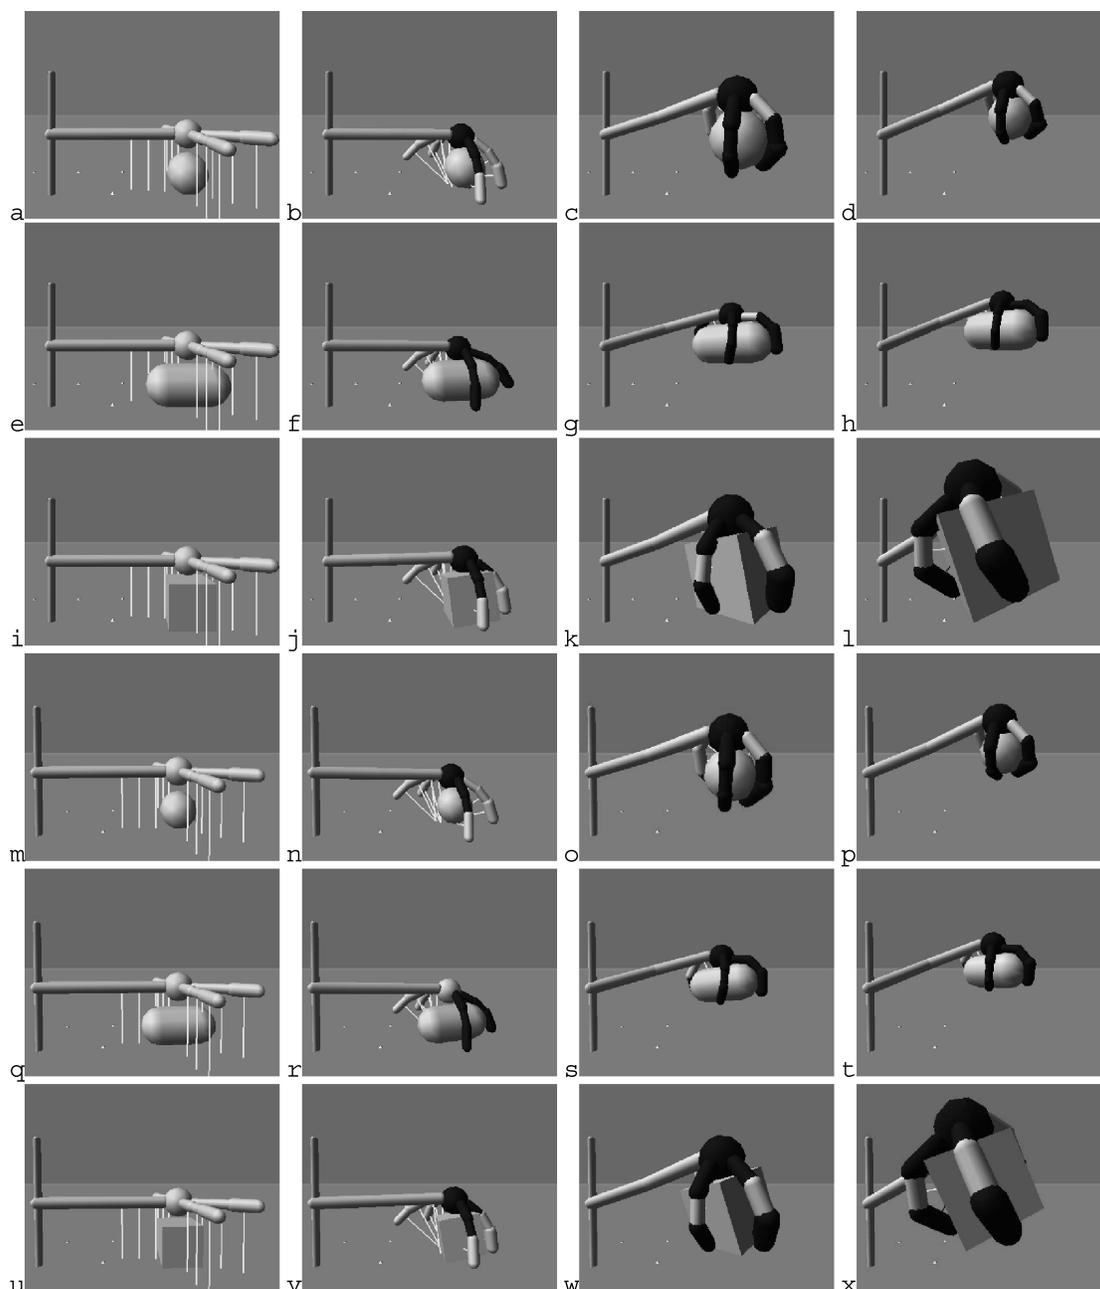


Figure 1. A sample evolved robot. This robot can successfully grasp, lift, and actively distinguish between spheres of different sizes (a–d, radius  $r = 35$  cm; e–h,  $r = 31.5$  cm), cylinders of difference sizes (i–l,  $r = 35$  cm, length  $l = 70$  cm; m–p,  $r = 31.5$  cm,  $l = 63$  cm), and cubes of different sizes (q–t,  $l, w, h = 70$  cm; u–x,  $l, w, h = 63$  cm). Black and white indicate finger segments in which the tactile sensor is on or off, respectively. White lines denote the range sensors. Videos of this and other robots can be found at [www.cs.uvm.edu/~jbongard](http://www.cs.uvm.edu/~jbongard).

correlation between the likelihood of successfully performing complex tasks and the amount of the robot's morphology placed under evolutionary control.

In the experiments described here, populations of robots are evolved to achieve one or more of these objectives simultaneously when presented with different objects. Some of the objectives are synergistic in the sense that an evolutionary improvement of one increases the likelihood of subsequent evolutionary improvement of another: For example, if the robot is evolved to grasp and lift objects, evolutionary improvement in grasping will likely lead to better lifting in future generations. Alternatively, some of the objectives are antagonistic: Improvement in one reduces the likelihood of

success at the other. For example, better grasping of different-shaped objects requires a uniformly tight grip on all of them, which reduces sensory differences between the actions. This reduction in sensory difference makes active perception of differences between objects more difficult.

Nevertheless, it was found that in these experiments robots sometimes evolve to simultaneously achieve all three objectives. For example, the robot shown in Figure 1 evolved to successfully grasp, lift, and actively distinguish between six different objects. The different trajectories of the arm (away from the camera for the spheres and cylinders, and toward the camera for the cubes) and the differential firing of the tactile sensors indicates visually that it is actively distinguishing between objects of different shapes. Similar actions for objects of the same shape but different size indicate that it has not learned to recognize specific objects, but rather to classify them based on one of their aspects (shape) rather than another (size). Section 2.4 describes in more detail how these different objectives were selected for.

## 2.2 Robot Morphology

The three-fingered robot is composed of 13 body parts with 15 degrees of freedom (DOFs). Each DOF is composed of a rotational joint and an actuator: The actuator causes the two body parts to rotate relative to one another through some predefined plane; the joint acts as the fulcrum. Two exceptions are the shoulder joint (a two-DOF actuated rotational joint) and the wrist joint (a three-DOF actuated rotational joint). If we define the robot's long axis to be the vector that passes from its shoulder through its hand, then the shoulder becomes its posterior point and its hand its anterior point. The sagittal plane is then the vertical plane that cuts through the length of the arm, the coronal plane the horizontal plane that cuts through the length of the arm, and the transverse plane a horizontal plane perpendicular to the sagittal and coronal planes. Details regarding the physical characteristics of the body parts and joints are provided in Table 1. The robots with four fingers have  $4 + 4 \times 3 = 16$  body parts and  $6 + 4 \times 3 = 18$  DOFs; the robots with five fingers have  $4 + 5 \times 3 = 19$  body parts and  $6 + 5 \times 3 = 21$  DOFs.

In the experiments reported on here, zero to three aspects of the robot's body plan may be evolved along with its controller. The first aspect is the length of the phalanges (finger segments): Each phalange may be evolved away from the default length of 0.3 m to any value in the range [0.01 m, 0.6 m]. The second aspect is the phalange radii: Each phalange may be evolved away from its default radius of 0.075 m to any value in the range [0.015 m, 0.3 m]. The third aspect affects the spacing between the fingers: At the outset of each trial, all fingers are arrayed around the hand with uniform spacing; if this aspect is evolved, the relative spacings between any pair of fingers  $i$  and  $i + 1$  (where  $i \in [0, 1]$  for the three-fingered hand,  $i \in [0, 2]$  for the four-fingered hand, and  $i \in [0, 3]$  for the five-fingered hand) may change to any angle in  $[-\pi, \pi]$ .

Each robot is equipped with three binary tactile sensors per finger (one for each phalange), and an additional one in the hand. Each phalange is also equipped with a range sensor (indicated by the white lines in Figure 1). A range sensor returns a value between zero and one commensurate with the length of the ray emitted by the sensor. Clearly these range sensors detract from the anthropomorphic aspect of the robot, but as the goal of the current work is to demonstrate the reasons for co-optimization of morphology and control rather than evolve a realistic robot, this is a minor detail: As has been shown previously [14, 27, 42], it is possible to evolve object manipulation for anthropomorphic robot arms without range sensors. The shoulder contains a proprioceptive sensor that measures the sagittal rotation of the arm: High positive values indicate the arm is raised, values near zero indicate the arm is horizontal, and high negative values indicate the arm is rotated downward. In initial experiments more proprioceptive sensors were also used, but were not found to increase the evolvability of the robot (data not shown). In the present experiments no noise was added to the sensors, motors, or task environment.

## 2.3 Robot Control

A continuous-time recurrent neural network [4] is used to control the robot. For the robot with three fingers, there are  $4 + 3 \times 2 = 10$  motor neurons controlling the 10 independent motors;

Table I. Physical parameters of the robot and the target objects.

Part	Length (m)	Height (m)	Width (m)	Mass (kg)
Large sphere [S]	0.35			1
Large cylinder [Y]	0.35	0.7		1
Large cube [C]	0.7	0.7	0.7	1
Small sphere [s]	0.315			1
Small cylinder [y]	0.315	0.63		1
Small cube [c]	0.63	0.63	0.63	1
Trunk [Tr]	0.05	2.0		1
Upper arm [Ua]	0.1	1.0		1
Fore arm [Fa]	0.1	1.0		1
Hand [Ha]	0.25			1
Proximal phalanges [Pp]	0.075	0.3		1
Intermediate phalanges [Ip]	0.075	0.3		1
Distal phalanges [Dp]	0.075	0.3		1

Joint	Rotation angle (deg)		Plane of rotation
	min	max	
Shoulder [Tr][Ua]	-30	30	Sagittal and coronal
Elbow [Ua][Fa]	-30	30	Coronal
Wrist [Fa][Ha]	-30	30	Sagittal, coronal, and transverse
Metacarpophalangeal [Ha][Pp]	-90	90	Relative to finger
Proximal interphalangeal [Pp][Ip]	-60	60	Relative to finger
Distal interphalangeal [Ip][Dp]	-60	60	Relative to finger

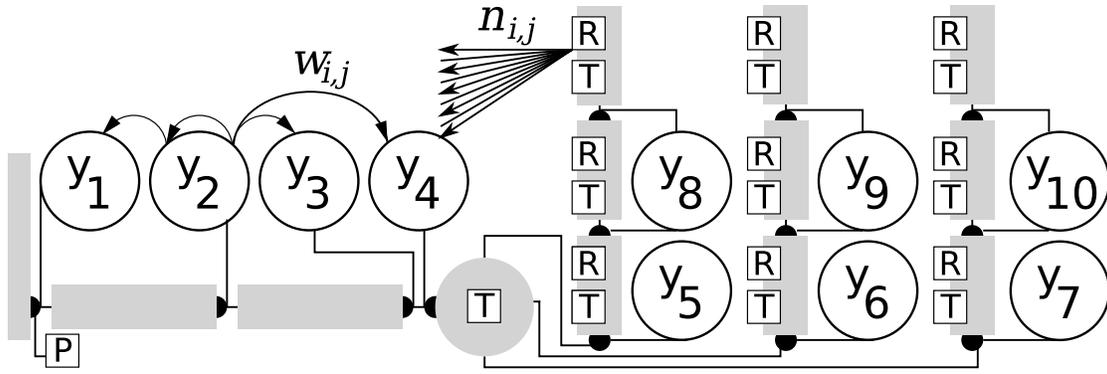


Figure 2. Architecture of the robot’s controller. Gray bars indicate body parts; black circles indicate actuated rotational joints connecting them. For the three-fingered hand, 10 motor neurons ( $y_1, \dots, y_{10}$ ) actuate the joints. Each motor neuron is fully connected to every other with a weighted connection ( $w_{i,j}$ ). Each of the 20 sensors is also connected to every motor neuron with a weighted connection ( $n_{i,j}$ ; R = range sensor; T = tactile sensor; P = proprioceptive sensor). Each distal and intermediate phalange is co-actuated by the same motor neuron ( $y_8, y_9, y_{10}$ ).

for each finger, the intermediate and distal phalanges move synchronously (as they do in the human hand). This is achieved by feeding the output from a single motor neuron to both phalanges. Otherwise, all DOFs move independently. Controllers of the four-fingered and five-fingered robots contain  $4 + 4 \times 2 = 12$  and  $4 + 5 \times 2 = 14$  motor neurons, respectively.

The value of each motor neuron is updated according to

$$\tau_i y_i' = -y_i + \sum_{j=1}^{m^{(3)}+2} w_{ji} \sigma(y_j - \theta_j) + \sum_{j=1}^{s_r^{(3)}+s_t^{(3)}} n_{ji} r_j, \quad (1)$$

where  $\tau_i$  is the time constant associated with neuron  $i$ ,  $y_i$  is the value of neuron  $i$ ,  $m^{(3)}$  is the number of motor neurons for the three-fingered robot (likewise  $m^{(4)}$ ,  $m^{(5)}$  for the four- and five-fingered hands), the additional two hidden neurons are used for active perception as explained in the Section 2,  $w_{ji}$  is the weight of the synapse connecting neuron  $j$  to neuron  $i$ ,  $\sigma(x) = 1/(1 + e^{-x})$  is an activation function that brings the value of neuron  $i$  back into  $[0, 1]$ ,  $\theta_j$  is the bias of neuron  $j$ ,  $s_r^{(3)}$  is the number of range sensors in the three-fingered hand,  $s_t^{(3)}$  is the number of tactile sensors in the three-fingered hand,  $n_{ji}$  is the weight of the synapse connecting sensor  $j$  to neuron  $i$ , and  $r_j$  is the value of sensor  $j$ . In this formulation, each sensor may have a direct effect on every motor neuron. However, this effect may be minimized or eliminated by low values for  $r_j$ , or by behaviors that cause a motor neuron to saturate to extremal values.  $\tau_i$  can range in  $[0.0001, 1.0]$ ,  $w_{ji}$  in  $[-16, 16]$ ,  $\theta_j$  in  $[-4, 4]$ , and  $n_{ji}$  in  $[-16, 16]$ . The continuous-time recurrent neural network (CTRNN) architecture is reported in Figure 2.

The virtual robot is equipped with a CTRNN; genetic modifications of its morphology may be introduced, and it is then evaluated over a set number of simulation steps in a physical simulator.<sup>1</sup> At the outset of each step, the sensor values are retrieved from the physical simulator, one update of the CTRNN is made, and the resulting values of the motor neurons are calculated. The values are scaled to the minimum and maximum rotation angles of the corresponding joint (Table 1), forming the desired angle for that joint. Torque is then applied to the joint, commensurate with the difference between the joint’s current angle and the desired angle. The positions and velocities of the objects in the simulation are then updated using a step size of 0.003.

<sup>1</sup> Open Dynamics Engine, [www.opende.com](http://www.opende.com).

## 2.4 The Optimization Process

In the following experiments robots are evolved to successfully perform one or more of the three object manipulation objectives. This is accomplished using multi-objective optimization [16], in which non-dominated solutions in the populations are retained from one generation to the next, and dominated solutions are overwritten by mutated copies of the non-dominated solutions. A solution is defined to be dominated if the value it obtains for each of the objectives is equal to or less than the value of the corresponding objective attained by another solution in the population.

Each evolutionary trial begins with 200 randomly generated genomes. At each generation the genomes that have not yet been evaluated are translated into robots and simulated, and a fitness value is calculated for each of the objectives that the robot must accomplish. The population is then scanned to find the non-dominated solutions; the dominated solutions are deleted; for each empty slot in the population a non-dominated solution is chosen at random, copied, mutated, and stored in the slot. The process is repeated until every empty slot is filled. This process continues for 300 generations, or until a robot evolves that can successfully manipulate all six objects.

The genome is defined as the tuple  $g = [\tau, w, \theta, n, LN, RD, SP]$  where the first four elements define the robot's controller,  $LN$  is a matrix that specifies the phalanges' lengths,  $RD$  is a matrix that specifies the phalanges' radii, and  $SP$  is a vector that specifies the  $k - 1$  spacings between the  $k$  fingers. If any of these morphological aspects is not evolved, it is not included in the genome. For the three-, four-, and five-fingered hand robots in which all morphological parameters are evolved along with the control parameters, there are 340, 507, and 706 free parameters, respectively.

When a genome is mutated, each value in each element of  $g$  has a probability 0.05 of being changed. The value is replaced with a new random value chosen with a Gaussian distribution such that the mean is equal to the original value and the variance is equal to the legal range for that parameter. If the new value falls outside the legal range, it is thresholded to its minimum or maximum value.

### 2.4.1 The Fitness Functions

When a robot has been evaluated, its ability to grasp, lift, and/or actively distinguish between objects is measured as a function of its recorded sensor values. These abilities are calculated as follows. Grasping is evolved by selecting for robots that minimize the distance between their fingers and the object. The fitness of grasping is thus defined as

$$f_g^{(k)'} = \frac{\sum_{i=1}^t \sum_{j=1}^{s_r} x_{ij}^{(k)}}{ts_r},$$

$$f_g^{(k)} = \begin{cases} 0.95, & f_g^{(k)'} > 0.95, \\ f_g^{(k)'}, & \text{otherwise,} \end{cases}$$

where  $k$  labels the  $k$ th object that the robot is attempting to grasp,  $t$  is the number of simulation time steps, and  $x_{ij}^{(k)}$  is the sensor value obtained during time step  $i$  from sensor  $j$  when manipulating object  $k$ . This measure therefore selects for robots that grasp objects as rapidly as possible and hold them firmly. A robot is considered to have successfully grasped an object if  $f_g^{(k)}$  is above a specific threshold. For robots that were only selected for grasping (i.e., runs in which  $F = f_g$ ), robots were routinely evolved that could achieve  $f_g^{(k)} > 0.95$  for all six objects. (The rationale for this thresholding will be explained in the next section.) Therefore the fitness of a robot evolved to grasp  $o$  objects is defined as

$$f_g = \frac{\sum_{k=1}^o f_g^{(k)}}{o}, \quad (2)$$

where  $o$  is the current number of objects that the robot must grasp.

The fitness of lifting is defined as

$$f_l^{(\kappa),t} = \frac{\sum_{i=1}^t \left( \left( \sum_{j=1}^{s_t} x_{ij}^{(\kappa)} \right) > 0 \right) x_{i s_p}^{(\kappa)}}{t},$$

$$f_l^{(\kappa)} = \begin{cases} 0.69, & f_l^{(\kappa),t} > 0.69, \\ f_l^{(\kappa),t} & \text{otherwise,} \end{cases}$$

where  $s_p$  is the value of the proprioceptive sensor in the shoulder. This function accumulates values at each time step if one or more of the touch sensors is in contact with the object (determined by the first term in the numerator). Robots were routinely found that could lift all six objects such that  $f_l > 0.69$  for each object. The fitness of a robot evolved to lift  $o$  objects is defined as

$$f_l = \frac{\sum_{\kappa=1}^o f_l^{(\kappa)}}{o}. \quad (3)$$

The robot must also learn to classify objects by physically interacting with them rather than passively observing them, a competence known as active categorical perception [3, 31, 42]. In this work, the robot evolves to consider objects of the same shape as belonging to the same class, and objects of different shape as belonging to different classes. This is achieved using a technique adapted from [42], in which values of hidden neurons are used to indicate the category of the current object being manipulated. This is selected for using the fitness function

$$f_a^{(\kappa,l),t} = \begin{cases} \sqrt{\left( b_{0,t}^{(\kappa)} - b_{0,t}^{(l)} \right)^2 + \left( b_{1,t}^{(\kappa)} - b_{1,t}^{(l)} \right)^2}, & l = \kappa + 3, \\ 1 - \sqrt{\left( b_{0,t}^{(\kappa)} - b_{0,t}^{(l)} \right)^2 + \left( b_{1,t}^{(\kappa)} - b_{1,t}^{(l)} \right)^2} & \text{otherwise,} \end{cases}$$

$$f_a^{(\kappa,l)} = \begin{cases} 0.8, & f_a^{(\kappa,l),t} > 0.8, \\ f_a^{(\kappa,l),t} & \text{otherwise,} \end{cases}$$

where  $b_{i,t}^{(\kappa)}$  is the value of the  $i$ th hidden neuron during the final simulation time step  $t$  when manipulating object  $\kappa$ . Since objects are shown to each robot in the sequence shown in Table 1, object  $l$  has the same shape as object  $\kappa + 3$ . This function therefore rewards the robot for attaining similar hidden neuron values when presented with objects of the same shape, but different hidden neuron values when presented with objects of different shape. Robots were routinely found that could actively distinguish between all six objects in that  $f_a > 0.8$  for each pair of objects. The fitness of a robot evolved to actively categorize  $o$  objects is therefore defined as

$$f_a = \frac{\sum_{i=1}^{o-1} \sum_{j=i+1}^o f_a^{(i,j)}}{o(o-1)/2}. \quad (4)$$

Section 3 reports seven experimental regimes: In the first, second, and third regimes, the population is only evolved to succeed at  $f_g$ ,  $f_a$ , or  $f_l$ , respectively; in the fourth regime, the population is evolved to succeed at both  $f_a$  and  $f_l$ ; in the fifth regime, it is evolved to succeed at both  $f_g$  and  $f_l$ ; in the sixth regime, both  $f_g$  and  $f_a$ ; and in the final regime the population is evolved to succeed at all

three objectives  $f_g$ ,  $f_a$ , and  $f_l$ . It should be noted that multi-objective optimization is used to evolve populations within each of these regimes, even though the first three contain only a single objective. In these cases, there is only ever one non-dominated solution that produces mutated copies that replace the other  $n - 1$  solutions, thereby collapsing to a  $1 + \lambda$  evolution strategy [36] without an adaptive mutation rate. In the fourth, fifth, and sixth regimes the non-dominated front is a line, and in the seventh regime the non-dominated front is a two-dimensional surface.

### 2.4.2 Shaping

A common technique in AI, robot shaping [40, 18, 38, 34, 37, 26, 47], was originally adapted from the concept of scaffolding in developmental psychology (e.g., [45]). Typically, the population of robots learn or evolve in a simplified environment. The environment is then gradually changed to continuously challenge the robots. If done correctly, this provides a smoother learning gradient and/or fitness landscape for the population to climb, and therefore more often leads to a satisfactory solution. Shaping is employed here by initially exposing the robot population to the first object (large sphere; see Table 1). Once one member of the population learns to manipulate the object successfully, the population is reevaluated against the first object and the second one (large cylinder), and evolution continues until a robot evolves that can successfully manipulate both objects. This process is repeated until either 300 generations elapse or a robot evolves that can successfully manipulate all six objects. Success is defined here as the ability of a robot to obtain the maximum value for all of the objectives  $x$  that it must achieve ( $\forall x \in \{g, a, l\} : f_x = [0.9 | 0.69 | 0.8]$ ).

### 2.4.3 Early Stopping

The main limitation of shaping methods is that they do not scale well: As the evolving population accumulates many training or fitness instances, each solution takes longer to evaluate. Here we introduce a modification to shaping that removes this scalability problem; we refer to this modification as early stopping.

We begin by assuming that a robot population is currently evolving to manipulate  $k > 1$  objects. Each robot is evaluated against the training set of objects in reverse order: If  $k = 2$ , a new robot would be evaluated against the large cylinder first and then the large sphere; if  $k = 3$ , a new robot would first be evaluated against the large cube, and so on. After evaluating a new robot with its first object, its  $f_x$  values for each objective  $x$  that it must achieve are calculated using  $f_x^{(1)}$ , and  $f_x^{(\max)}$  for the additional  $k - 1$  objects. Its  $f_x$  values are then compared with the corresponding  $f_x$  values for all

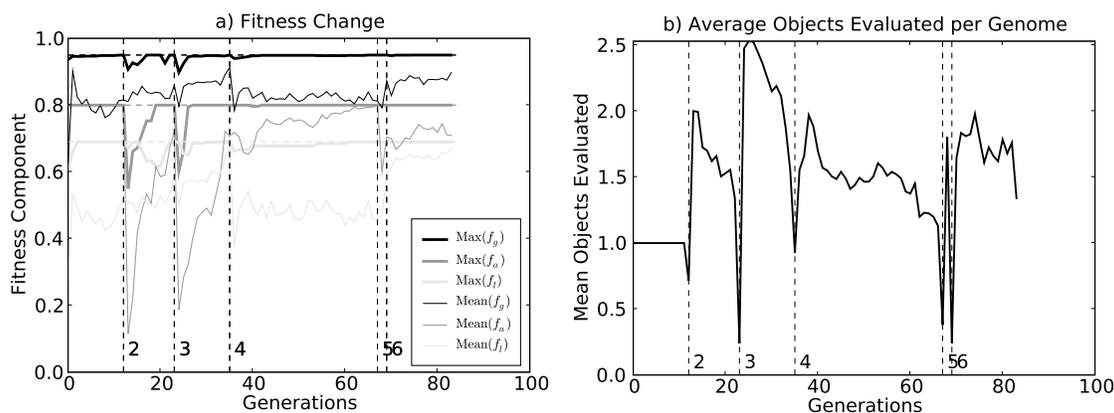
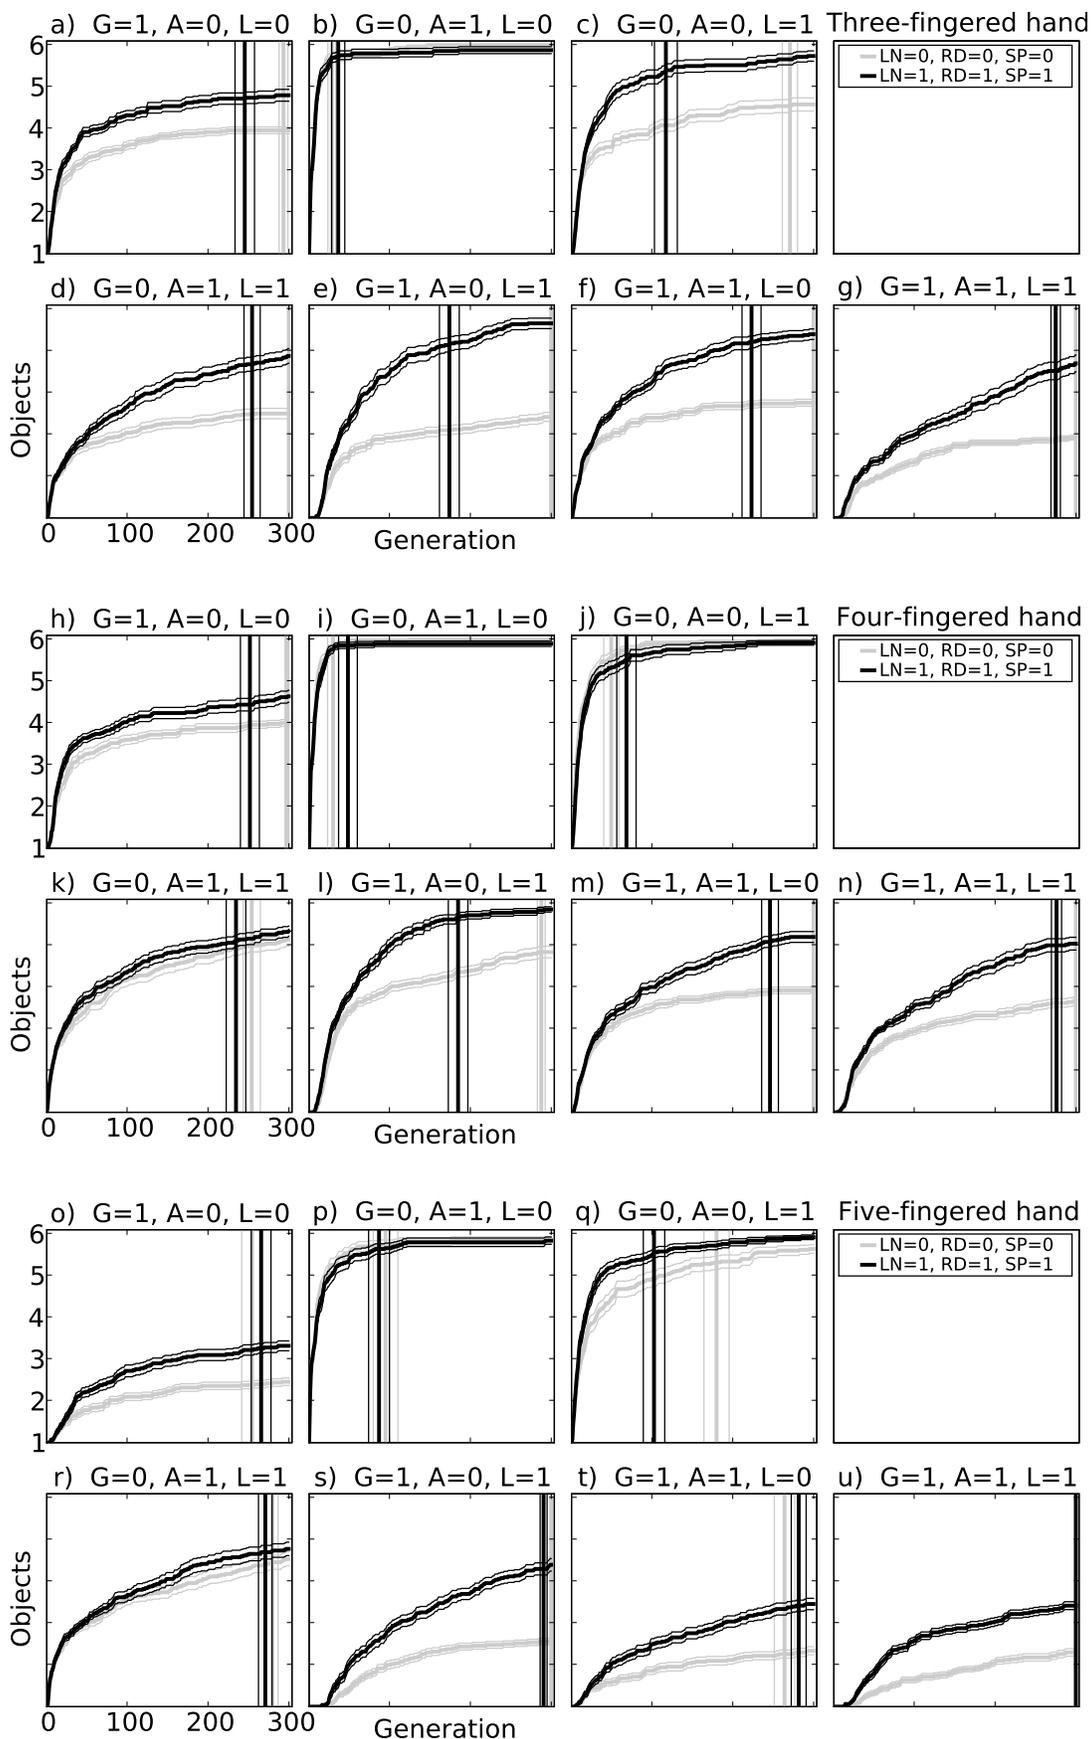


Figure 3. Evolutionary trajectory that produced the robot in Figure 1. (a) Black, dark gray, and light gray lines indicate ability to grasp, actively distinguish between, and lift objects, respectively. Thick lines indicate ability of the fittest robot; thin lines indicate average ability across the population. Vertical dashed lines indicate successful manipulation of the current set of objects; numbers next to lines indicate the new number of objects that must be manipulated. The robot successfully evolved to manipulate all six objects at generation 84. (b) The mean number of objects evaluated per robot during the same evolutionary run.



the previously evaluated solutions using their  $f_x^{(1)}, \dots, f_x^{(\ell)}$  values. If the current robot is dominated by one other solution, it will remain dominated by that solution even if evaluated against the remaining objects, because its  $f_x$  values can only decrease, as they are currently set to  $f_x^{(\max)}$ . Therefore, the current genome can be tagged as dominated, and the remaining  $\ell - 1$  evaluations canceled. If on the other hand it is non-dominated, it is evaluated against the second object, and its  $f_x$  values are recalculated using its  $f_x^{(1)}$  and  $f_x^{(2)}$  values, and the remaining  $f_x^{(3)}, \dots, f_x^{(\ell)}$  values maximized. It is again tested against the rest of the population for dominance, and if it is dominated, evaluation of the current robot stops. This process continues, and is repeated up to  $\ell$  times for each new robot.

Using this approach, we can terminate the evaluation of a robot early if further evaluation cannot allow it to escape dominance by another robot. Further, by evaluating a new robot against the training set of objects in reverse order, we increase the probability of early stopping, because a new robot is more likely to manipulate objects worse than its parents did if the objects are newly added to the training set than if they have been there for many generations.

The cost of early stopping is having to determine what the maximum value is for each objective. Here this was done by performing runs in which only one objective was selected for, and shaping was not used: Robots were evolved against all six objects. Because these task environments are simpler, populations rapidly evolved robots capable of achieving high values for the objectives: We performed three sets of runs with each of the single objectives, took the population with the lowest mean  $f_x$  from each, and set  $f_x^{(\max)}$  to the value obtained by the best robot in that population. In future work we will investigate methods for overcoming this initial cost of early stopping.

### 3 Results

As mentioned above, seven experimental regimes were tested: In each regime, robots are evolved to meet one, two, or all three of the object manipulation objectives. Within each regime,  $8 \times 50 = 400$  evolutionary trials were performed: in the first set of 50 independent trials, only the controller was evolved; in the second set, the controller and phalange lengths were evolved; in the third set, the controller and phalange radii were evolved; in the fourth set, the controller and finger spacings were evolved; in the fifth, sixth, and seventh sets, two of the three morphological aspects were evolved along with the controller; and in the eighth set, the controller, phalange length, radii, and finger spacings were evolved together. These regimes and sets of trials were replicated on robots with three-, four-, and five-fingered hands, resulting in a total of  $3 \times 7 \times 8 \times 50 = 8400$  trials.

Figure 1 shows the final successful robot from a trial conducted in the seventh regime and the eighth set: The controller and all three morphological aspects of each robot were optimized so that the robots evolved to successfully grasp, lift, and actively distinguish between all six objects. The evolutionary trajectory for this trial is shown in Figure 3. During each generation of a trial, the best robot in the population is found by collapsing the objective values for each robot to a single number using  $f = f_g f_a f_l$  (or a subset of these if not all three objectives are selected for), and finding  $\max(f)$  in the population. Figure 3a reports the change in the best robot's ability to achieve all three objectives, as well as the mean ability of the population to do so. It can also be observed that periodically a robot in the population successfully manipulates the current set of objects (vertical dashed lines), and a new object is added to the training set (the size of the set is indicated by the numbers next to the dashed lines). The robot shown in Figure 1 was evolved in the 84th generation, causing the trial to terminate.

---

Figure 4. Mean evolutionary change in manipulation ability for the three-fingered (a–g), four-fingered (h–n), and five-fingered hands (o–u). G,A,L = 0,1 indicate regimes in which grasping, active categorical perception, or lifting was not or was selected for, respectively. LN,RD,SP = 0,1 indicate trial sets in which phalanges' lengths or radii or finger spacings were not or were evolved, respectively. Thick lines indicate the mean number of objects in the training set; thin lines indicate one standard error of the mean. Gray and black curves indicate the manipulation abilities of hands with fixed morphologies and those in which all three morphological aspects were evolved, respectively. Vertical lines indicate the mean number of generations until robots evolved to successfully manipulate all six objects.

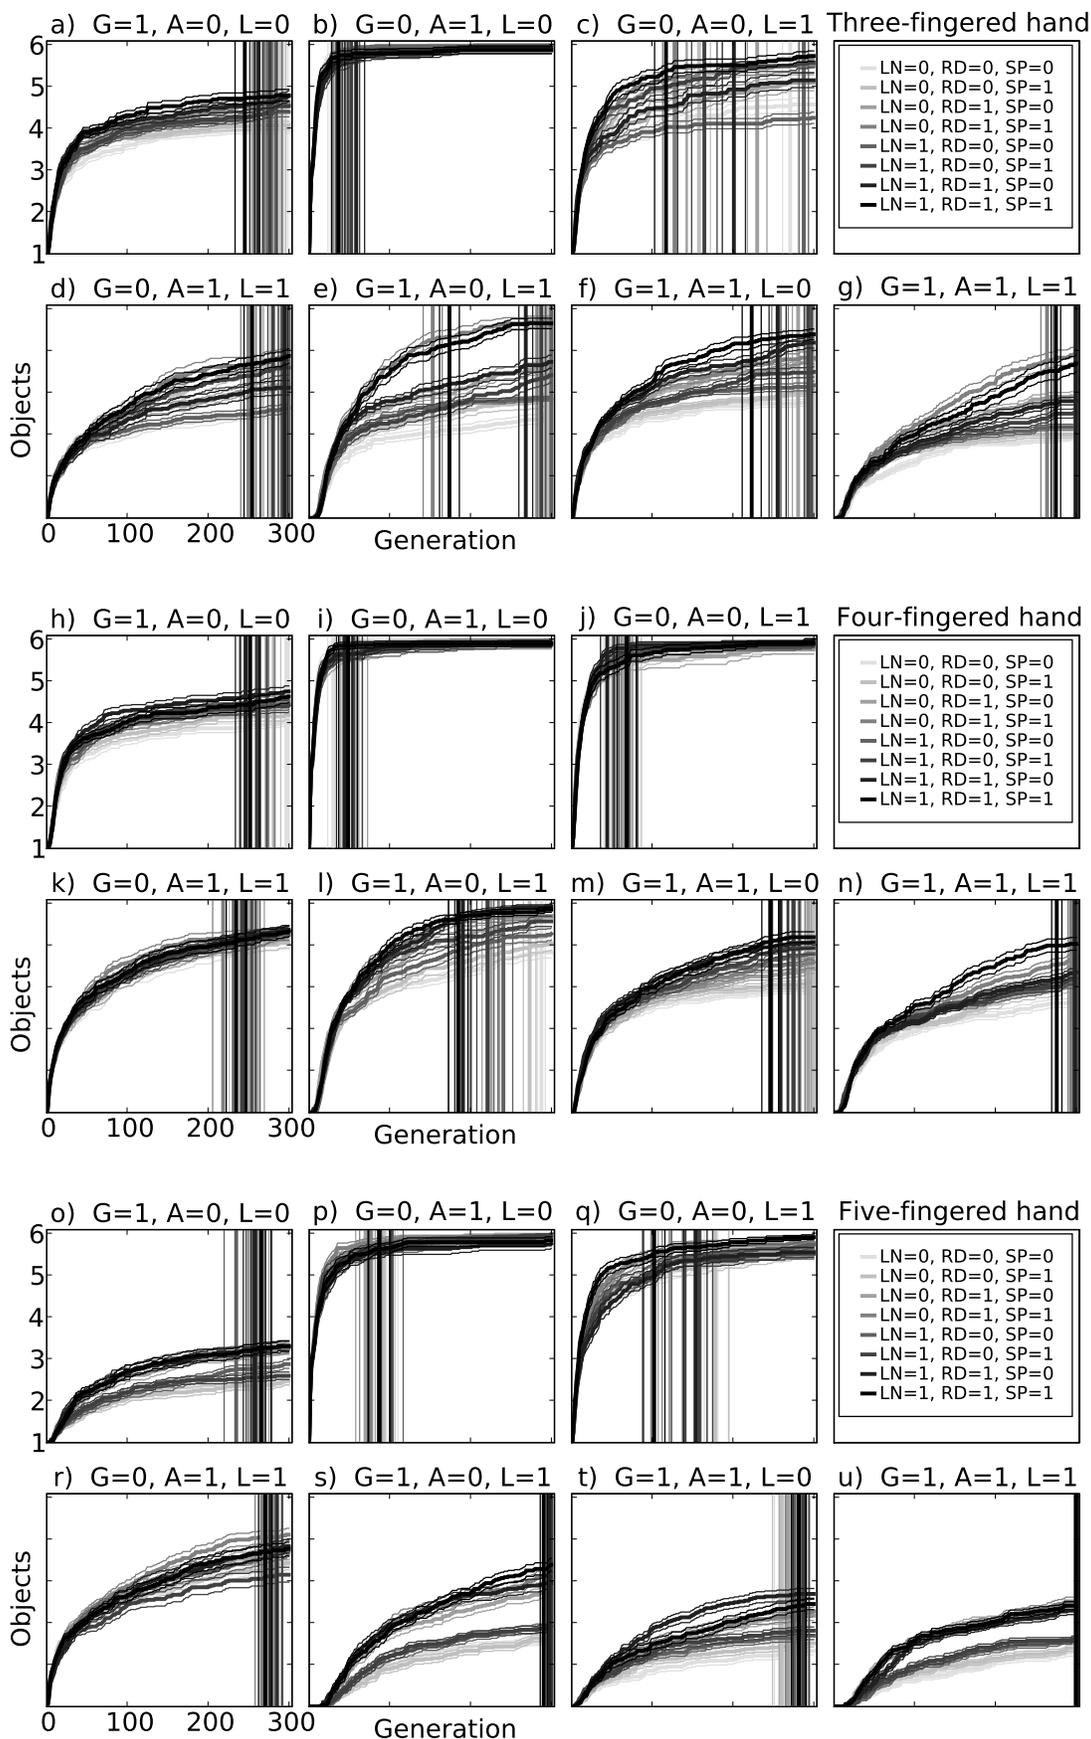


Figure 3b reports the mean number of objects manipulated by robots in the population during that generation. As can be seen, when a new object is added to the training set, there is an increase in the number of objects that are evaluated. This is expected, because most robots in the population experience a drop in fitness across the objectives, and it is easier for a new robot to escape dominance by those solutions on the non-dominated front. As the population evolves to better manipulate the new object, the mean number of objects manipulated by the population drops, and the simulation runs faster. This is due to the fact that the non-dominated solutions are now adept at manipulating the objects in the set, and the probability that a new robot will be dominated early rises. It can also be noted that as the number of objects in the training set rises, the mean number of object manipulations does not: This indicates that the evaluation time per robot does not increase with the size of the training set. This was found to be typical of the other runs (data not shown).

### 3.1 Task Complexity

Figure 4 reports the performance difference between sets of trials in which only the controller was evolved (gray lines), and in which the controller and the three morphological aspects were evolved for robots with three-, four-, and five-fingered hands (black lines). It is observed that when only one objective is selected for, all three robot classes fail to perform significantly better when morphology is evolved than when it is fixed (Figure 4a–c, h–j, o–q). Indeed, there is no benefit when only active perception is selected for (Figure 4b, i, p): This is because these robots typically evolve a classification strategy in which the objects are passively scanned by the range sensors without requiring physical contact between the robot and the object; such strategies are not improved by allowing evolution to shape hand morphology.

For robots evolved to achieve two objectives (Figure 4d–f, k–m, r–t), there is a larger improvement in those regimes in which morphology was evolved. Two exceptions are when active perception and lifting are selected for the four- and five-fingered hands (Figure 4k, r). Finally, the largest relative improvement between robots with fixed and evolved morphologies is observed in cases where all three objectives are selected for, compared to the other seven cases in which one or two objectives are selected for (Figure 4g, n, u). This indicates that as task complexity increases (where task complexity is defined as the number of objectives that the robot must accomplish simultaneously), the benefit of evolving morphology becomes increasingly important for discovering successful solutions.

Figure 5 reports the relative performance of all regimes and sets of runs for the three robot classes. As can be seen, there is relatively little difference between the sets when two or fewer objectives are evolved. Exceptions are seen in Figure 5e, r, in which evolving the phalange radii and finger spacings provides significant improvement; Figure 5f, in which evolving all three morphological aspects helps considerably; and Figure 5t, in which evolving the phalange length and radii provides significant improvement. Finally, it can be observed that for the three-fingered hand, when robots must accomplish all three objectives, it is significantly helpful to evolve either the phalange radii and finger spacings, or all three aspects. For the four-fingered hand, significant improvement is obtained when all three morphological aspects are evolved. For the five-fingered hand, evolving the phalange radii provides significant improvement. This indicates that when selecting for three simultaneous objectives, performance is improved by placing more aspects of the robot's morphology under evolutionary control. More specifically, placing any additional morphological aspect under evolutionary control never significantly degrades performance, and often significantly improves it.

These observations are further supported by Table 2, in which the relative performance of the sets of runs is compared for the regime in which all three objectives are selected for ( $G, A, L = 1$ ). It can be seen that runs in which phalange radius and finger spacings are evolved ( $LN = *$ ,  $RD, SP = 1$ ) significantly outperform all other sets of runs. This pattern is also observed for the robots with

---

Figure 5. Mean evolutionary change in manipulation ability for robots in which different aspects of morphology are evolved. For details regarding this figure, refer to Figure 4.

four-fingered (Table 3) and five-fingered (Table 4) hands. However, on comparing the mean number of generations until run termination for robots with four-fingered hands, it is observed that the runs finished significantly earlier when all three morphological aspects were evolved than when only the phalange radii and finger spacings were evolved (Table 5). For the robots with three- and five-fingered hands, there is no significant difference between ending times for the two conditions (data not shown). This provides additional evidence that for the regime in which all three objectives must be met, placing all three morphological aspects under evolutionary control significantly improves performance.

It seems possible that the superiority of evolving morphology is due to a poor choice of the default morphological settings in the cases in which one or more of the morphological aspects is fixed. If this were so, one would expect to observe evolution consistently drag the morphological parameters away from these default settings. However, Figure 6 reports the mean evolved values for all morphological parameters in which they were evolved. It can be seen that the mean values are all close to the original default settings (indicated by the dashed lines). The one exception is that the phalange radii for the fourth fingers typically are larger than the default settings (right-hand bars in Figure 6e–h). The fourth finger is typically the one closest to the forearm, and supplies the most pressure when the object is lifted; we hypothesize that this broadening helps with lifting.

The proximity of the evolved settings to the default settings indicates that there are no universally superior parameter settings for any of the morphological aspects far from the default settings. Rather, evolution benefits by being able to tune the *relative* dimensions of the morphological parameters within a single robot. For example, the intermediate phalanges of the robot in Figure 1 are narrower than its proximal and distal phalanges, allowing the tactile sensors on the intermediate phalanges to fire when its grasps spheres or cylinders (Figure 1d,h,l,p), but not when it grasps cubes (Figure 1t,x). This presumably aids in active categorical perception for this robot.

### 3.2 Robustness

The ability of robots to continue operation in unexpected situations through adaptation remains one of the major challenges for artificial intelligence, but progress is being made [10, 12]. To ascertain the effect of evolving morphology on robustness, robots that evolved to successfully manipulate all six objects were extracted and subjected to unseen task environments.

First, each robot was exposed to spheres, cylinders, and cubes with sizes different than those experienced during evolution. 400 spheres with radii in [0.75 m, 0.76 m, ..., 1.15 m]; 400 cylinders with radii in [0.75 m, 0.76 m, ..., 1.15 m] and lengths in [1.5 m, 1.52 m, ..., 2.3 m]; and 400 cubes with side lengths in [0.75 m, 0.76 m, ..., 1.15 m] were presented to each robot in triplets: The robot was evaluated against a single sphere, cylinder, and cube of the same size index, and its ability to grasp, lift, and/or actively distinguish between them was calculated, based on the regime from which the robot was drawn. The percentage drop in performance (performance impact) for each triple evaluation was then calculated as

$$p = \frac{\sum_{i=1}^{400} \sum_{j=1}^3 (f_j^{(\max)} - f_j^{(i)})}{1200 f_j^{(\max)}} \times 100\%,$$

where  $f_j^{(\max)}$  is the maximum attainable value for objective  $j$ , and  $f_j^{(i)}$  is the performance of a robot when confronted with triplets of objects with size index  $i$  and required to achieve objective  $j$ . For robots drawn from regimes in which two or fewer objectives were selected for, the summation over  $j$  is reduced.  $p$  was then calculated for all successful robots within each regime, averaged across those robots with zero, one, two, or three evolved morphological aspects, and is reported in Figure 7a. It can be seen that as the number of evolved morphological aspects increases, the robots become increasingly robust when confronted with objects of previously unseen size.



Table 2. (continued)

I	0	0	2.94 ± 0.05	3.0 ± 0.07	3.94 ± 0.1	4.92 ± 0.18	3.51 ± 0.1	3.79 ± 0.08	4.69 ± 0.2
			3.2 ± 0.06**	3.2 ± 0.06	3.2 ± 0.06	3.2 ± 0.06	3.2 ± 0.06	3.2 ± 0.06	3.2 ± 0.06
I	0	0	2.94 ± 0.05	3.0 ± 0.07	3.94 ± 0.1	4.92 ± 0.18	3.2 ± 0.06	3.79 ± 0.08	4.69 ± 0.2
			3.51 ± 0.1***	3.51 ± 0.1***	3.51 ± 0.1	3.51 ± 0.1	3.51 ± 0.1*	3.51 ± 0.1	3.51 ± 0.1
I	0	0	2.94 ± 0.05	3.0 ± 0.07	3.94 ± 0.1	4.92 ± 0.18	3.2 ± 0.06	3.79 ± 0.08	4.69 ± 0.2
			3.79 ± 0.08***	3.79 ± 0.08***	3.79 ± 0.08	3.79 ± 0.08	3.79 ± 0.08***	3.79 ± 0.08*	3.79 ± 0.08
I	0	0	2.94 ± 0.05	3.0 ± 0.07	3.94 ± 0.1	4.92 ± 0.18	3.2 ± 0.06	3.79 ± 0.08	4.69 ± 0.2
			4.69 ± 0.2***	4.69 ± 0.2***	4.69 ± 0.2**	4.69 ± 0.2	4.69 ± 0.2***	4.69 ± 0.2***	4.69 ± 0.2***











A similar procedure was then repeated with the same robots, in which a single range or tactile sensor in each robot was systematically disabled:

$$p = \frac{\sum_{k=1}^{s_r+s_t} \sum_{i=1}^3 \sum_{j=1}^3 \left( f_j^{(\max)} - f_j^{(k,i)} \right)}{9(s_r + s_t)f_j^{(\max)}} \times 100\%$$

where  $f_j^{(k,i)}$  is now the performance of the robot for objective  $j$  when manipulating object  $i$  with sensor  $k$  disabled.  $p$  was again averaged across the robots and grouped into regimes in which zero or

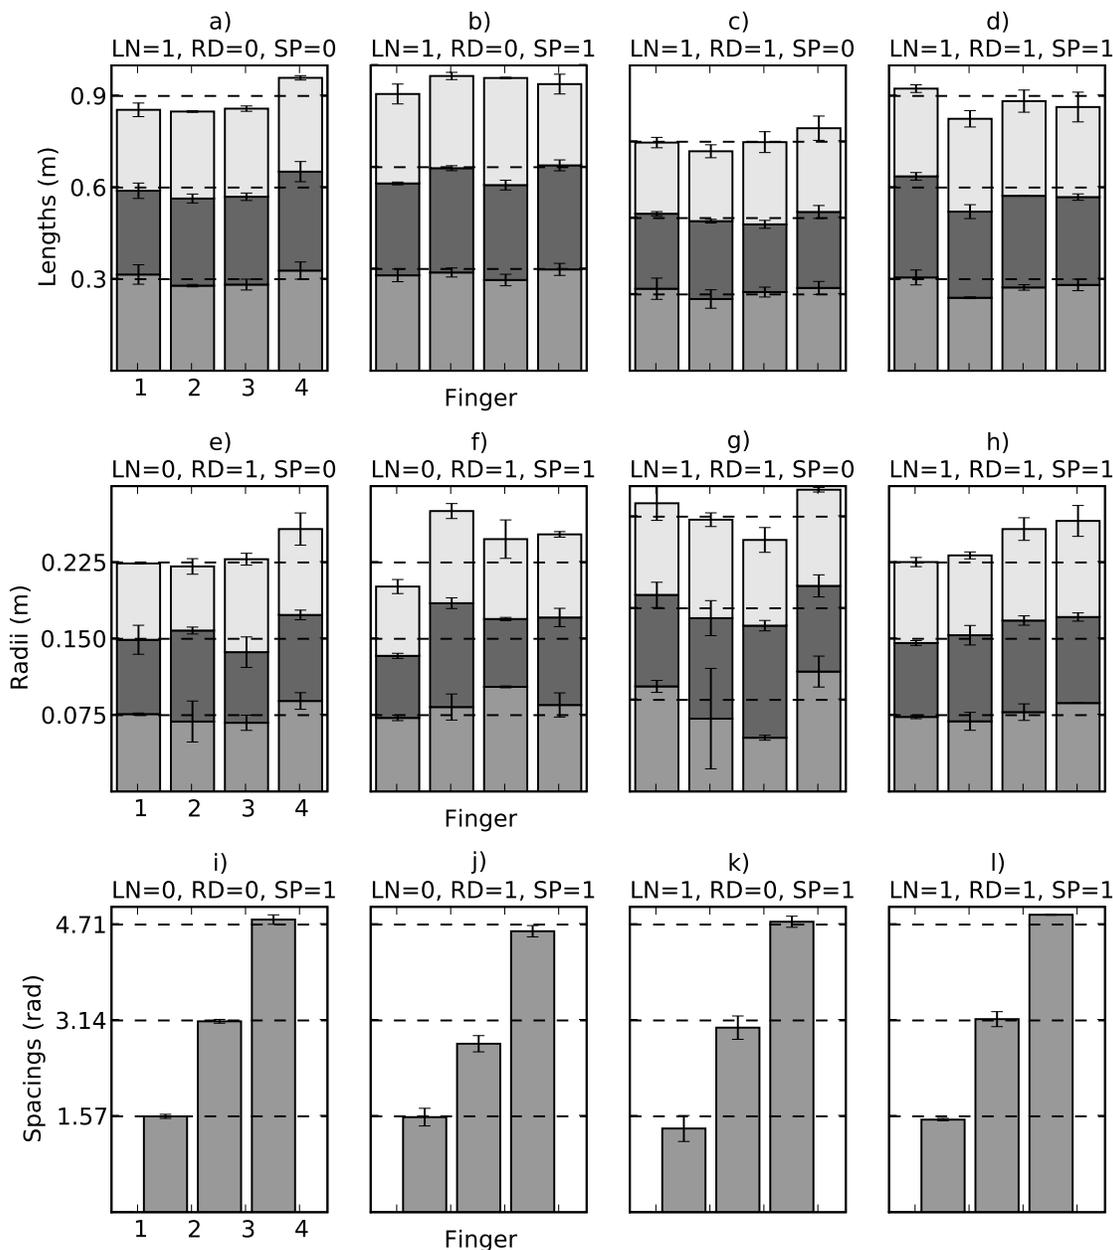


Figure 6. Mean evolved values for morphological parameters of the four-fingered hand. (a–d) Evolved phalange lengths for the regimes in which phalange lengths were evolved. (e–h) Evolved phalange radii for the regimes in which phalange radii were evolved. (i–l) Evolved finger spacings for the regimes in which finger spacings were evolved.

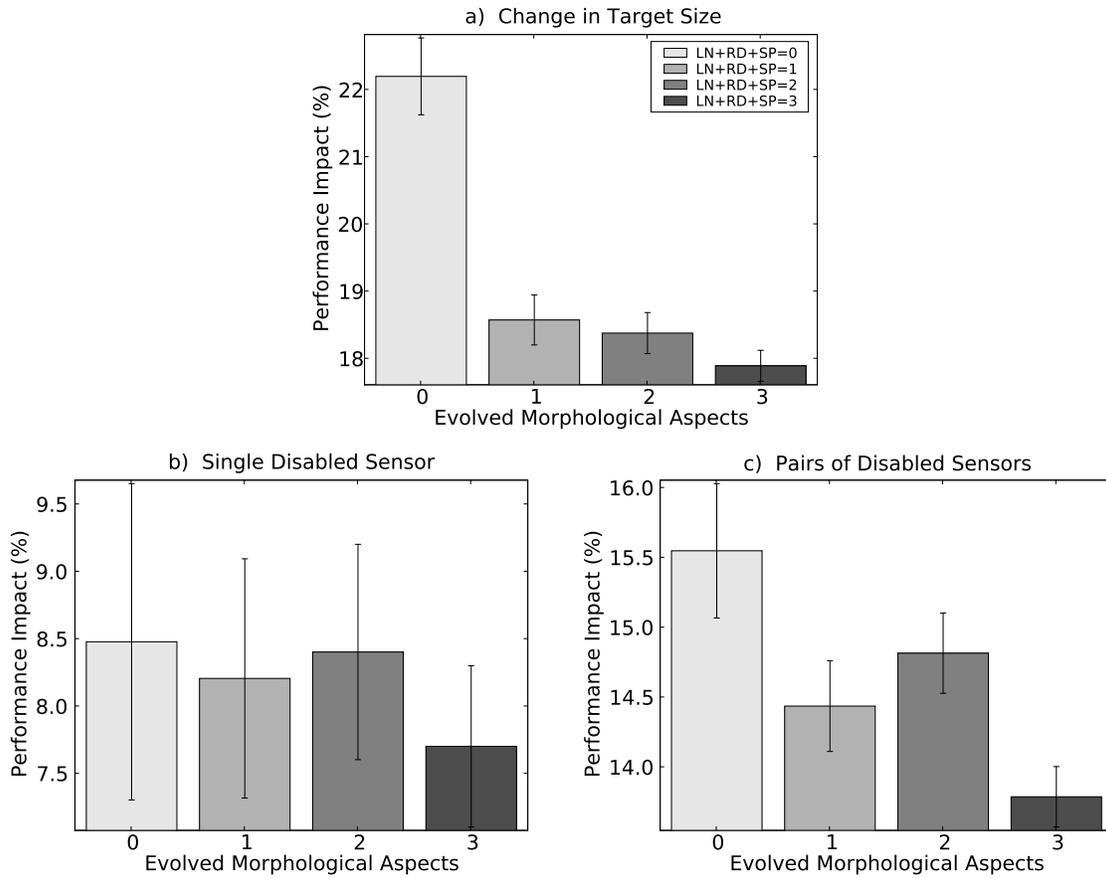


Figure 7. Performance impacts for the evolved four-fingered hands when confronted with novel environments. (a) Mean performance impacts when confronted with target objects of previously unseen sizes, grouped into regimes in which zero, one, two, or all three morphological aspects were evolved. (b) Mean performance impacts when a single sensor fails. (c) Mean performance impacts when pairs of sensors fail.

more morphological aspects were evolved, and is reported in Figure 7b. Here there is nonsignificant increase in robustness as the number of evolved morphological aspects increases. The procedure was again repeated with the same robots, but each pair of sensors was systematically disabled, and the drop in performance measured using

$$p = \frac{\sum_{l=1}^{s_r+s_i-1} \sum_{k=l+1}^{s_r+s_i} \sum_{i=1}^3 \sum_{j=1}^3 \left( f_j^{(\max)} - f_j^{(i,k,l)} \right)}{4.5(s_r + s_i)(s_r + s_i - 1)f_j^{(\max)}} \times 100\%,$$

where  $f_j^{(i,k,l)}$  is now the performance attained by the robot on objective  $j$  when manipulating object  $i$  with its  $k$ th and  $l$ th sensors disabled.  $p$  was again averaged across regimes in which zero or more morphological aspects were evolved, and is reported in Figure 7c. It can be seen that there is a greater impact on performance than that due to disabling a single sensor, and that now the robustness increases significantly with the number of evolved morphological aspects.

#### 4 Conclusions

Here we have evolved simulated robots to meet one or more object manipulation objectives, and measured whether and how much evolving morphology along with control improves the evolvability

of the population as the number of objectives that the robots must accomplish increases. It was found that performance generally degraded as the number of objectives increased, as was to be expected: Evolution must carefully tune the robot's phenotype so that an increase in performance for one objective does not overly degrade performance on another.

It was found that for one (or in some cases two) objectives, there was little or mild improvement in performance as more aspects of the robot's body plan were placed under evolutionary control. However, when all three objectives were selected for, there was a distinct advantage among the robots in which all morphological aspects were evolved. This demonstrates that choosing the appropriate morphology for an adaptive machine is not always intuitive. Even for seemingly appropriate designs, such as the anthropomorphic arm-and-hand system used here, there were an increasing number of morphological improvements that could be made as the task became more challenging. In addition it was found that among robots that evolved to successfully meet all objectives for all objects, those with more evolved morphological aspects were more robust in the face of two unanticipated conditions: objects of novel size, and sensor failure. Both of these observations provide two additional lines of evidence for why morphology should be optimized along with control when evolving adaptive machines: Evolving more aspects of the machine's body plan becomes more useful as task complexity increases, and the resulting evolved body plans are more robust than those with fixed morphologies. In future work we plan to investigate whether this correlation between evolved morphology, task complexity, and robustness holds in other task environments, and methods for formulating which aspects of morphology should be evolved based on the task at hand.

In order for embodied artificial intelligence to gain credibility within the robotics and cognitive science communities, it is necessary to amass evidence not only about how to co-optimize the physical as well as neural components of adaptive machines, but also about why. This work has provided two more such reasons.

## References

1. Adamatzky, A., Komosinski, M., & Ulatowski, S. (2000). Software review: Framsticks. *Kybernetes: The International Journal of Systems & Cybernetics*, 29, 1344–1351.
2. Auerbach, J., & Bongard, J. C. (2009). Evolution of functional specialization in a morphologically homogeneous robot. In *Genetic and Evolutionary Computation Conference (GECCO 2009)*.
3. Beer, R. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4), 209.
4. Beer, R. (2006). Parameter space structure of continuous-time recurrent neural networks. *Neural Computation*, 18, 3009–3051.
5. Beer, R. D. (2008). The dynamics of brain-body-environment systems: A status report. In P. Calvo & A. Gomila (Eds.), *Handbook of Cognitive Science: An Embodied Approach* (pp. 99–120). Elsevier.
6. Beer, R. D., Quinn, R. D., Chiel, H. J., & Ritzmann, R. E. (1997). Biologically inspired approaches to robotics. *Communications of the ACM*, 40(3), 30–38.
7. Bentley, P. J., & Kumar, S. (1999). Three ways to grow designs: A comparison of embryogenies for an evolutionary design problem. In W. Banzhaf, J. M. Daida, Á. E. Eiben, M. H. Garzon, V. Honavar, M. J. Jakiela, & R. E. Smith (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 1999)* (pp. 35–43).
8. Bongard, J., & Paul, C. (2000). Investigating morphological symmetry and locomotive efficiency using virtual embodied evolution. In *Proceedings of the Sixth International Conference on Simulation of Adaptive Behaviour* (pp. 420–429).
9. Bongard, J., & Pfeifer, R. (2001). Repeated structure and dissociation of genotypic and phenotypic complexity in artificial ontogeny. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2001)* (pp. 829–836).
10. Bongard, J., Zykov, V., & Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science*, 314, 1118–1121.

11. Bongard, J. C. (2002). Evolving modular genetic regulatory networks. In *Proceedings of the IEEE 2002 Congress on Evolutionary Computation (CEC2002)* (pp. 1872–1877).
12. Bongard, J. C. (2009). Accelerating self-modeling in cooperative robot teams. *IEEE Transactions on Evolutionary Computation*, 13(2), 321–332.
13. Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–160.
14. Buehrmann, T., & Di Paolo, E. (2004). Closing the loop: Evolving a model-free visually guided robot arm. In *Artificial life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Artificial Life* (pp. 63–72). MIT Press.
15. Clark, A. (1998). *Being there: Putting brain, body, and world together again*. Cambridge, MA: Bradford Books.
16. Deb, K. (2001). *Multi-objective optimization using evolutionary algorithms*. Wiley.
17. Dellaert, F., & Beer, R. (1994). *Co-evolving body and brain in autonomous agents using a developmental model* (Tech. Rep. CES-94-16). Cleveland, OH: Dept. of Computer Engineering and Science, Case Western Reserve University.
18. Dorigo, M., & Colombetti, M. (1994). Robot shaping: Developing autonomous agents through learning. *Artificial Intelligence*, 71(2), 321–370.
19. Eggenberger, P. (1997). Evolving morphologies of simulated 3D organisms based on differential gene expression. In *Proceedings of the Fourth European Conference on Artificial Life* (pp. 205–213).
20. Gibson, J. (1986). *The ecological approach to visual perception*. Lawrence Erlbaum.
21. Gomez, G., & Eggenberger, P. (2007). Evolutionary synthesis of grasping through self-exploratory movements of a robotic hand. In *IEEE Congress on Evolutionary Computation, 2007. CEC 2007* (pp. 3418–3425).
22. Hendriks-Jansen, H. (1996). *Catching ourselves in the act: Situated activity, interactive emergence, evolution, and human thought*. Boston: MIT Press.
23. Hornby, G., & Pollack, J. (2001). Evolving L-systems to generate virtual creatures. *Computers & Graphics*, 25(6), 1041–1048.
24. Lichtensteiger, L., & Eggenberger, P. (1999). Evolving the morphology of a compound eye on a robot. In *Proceedings of the Third European Workshop on Advanced Mobile Robots* (pp. 127–134).
25. Lipson, H., & Pollack, J. B. (2000). Automatic design and manufacture of artificial lifeforms. *Nature*, 406, 974–978.
26. Lungarella, M., Metta, G., Pfeifer, R., & Sandini, G. (2003). Developmental robotics: A survey. *Connection Science*, 15(4), 151–190.
27. Massera, G., Cangelosi, A., & Nolfi, S. (2007). Evolution of prehension ability in an anthropomorphic neurobotic arm. *Frontiers in Neurobotics*, 1(4), 1–9.
28. Mazzapioda, M., Cangelosi, A., & Nolfi, S. (2009). Evolving morphology and control: A distributed approach. In *2009 IEEE Congress on Evolutionary Computation (CEC 2009)* (pp. 2217–2224).
29. McGeer, T. (1990). Passive dynamic walking. *International Journal of Robotics Research*, 9(2), 62–82.
30. Nelson, A. L., Barlow, G. J., & Doitsidis, L. (2009). Fitness functions in evolutionary robotics: A survey and analysis. *Robotics and Autonomous Systems*, 57(4), 345–370.
31. Noë, A. (2004). *Action in perception*. MIT Press.
32. Pfeifer, R., & Bongard, J. (2006). *How the body shapes the way we think: A new view of intelligence*. MIT Press.
33. Pfeifer, R., & Scheier, C. (1999). *Understanding intelligence*. Cambridge, MA: MIT Press.
34. Pratt, J., Chew, C., Torres, A., Dilworth, P., & Pratt, G. (2001). Virtual model control: An intuitive approach for bipedal locomotion. *The International Journal of Robotics Research*, 20(2), 129–143.
35. Quick, T., Nehaniv, C., Dautenhahn, K., & Roberts, G. (2003). Evolving embodied genetic regulatory network-driven control systems. In *Advances in artificial life, ECAL'03, Lecture notes in computer science*, Vol. 2801 (pp. 266–277).
36. Rechenberg, I. (1994). *Evolutionsstrategie: Optimierung technischer System nach Prinzipien der biologischen Evolution*. Stuttgart: Frommann-Holzboog.

37. Reil, T., & Husbands, P. (2002). Evolution of central pattern generators for bipedal walking in a realtime physics environment. *IEEE Transactions on Evolutionary Computation*, 6(2), 159–168.
38. Saksida, L., Raymond, S., & Touretzky, D. (1997). Shaping robot behavior using principles from instrumental conditioning. *Robotics and Autonomous Systems*, 22, 231–250.
39. Sims, K. (1994). Evolving 3D morphology and behaviour by competition. In *Artificial Life IV* (pp. 28–39).
40. Singh, S. P. (1992). Transfer of learning across sequential tasks. *Machine Learning*, 8(3), 323–339.
41. Taylor, T. (2004). A genetic regulatory network-inspired real-time controller for a group of underwater robots. In *Proceedings of the Eighth Conference on Intelligent Autonomous Systems (IAS-8)* (pp. 403–412).
42. Tuci, E., Massera, G., & Nolfi, S. (2009). Active categorical perception in an evolved anthropomorphic robotic arm. In *2009 IEEE Congress on Evolutionary Computation (CEC 2009)* (pp. 31–38).
43. Ventrella, J. (1994). Explorations of morphology and locomotion behaviour in animated characters. In *Artificial Life IV* (pp. 436–441).
44. Webb, B., & Thomas, R. (2001). *Biorobotics: Methods and applications*. AAAI Press/MIT Press.
45. Wood, D., Bruner, J., & Ross, G. (1976). The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry*, 17(2), 89–100.
46. Yamashita, Y., & Tani, J. (2008). Emergence of functional hierarchy in a multiple timescale neural network model: A humanoid robot experiment. *PLoS Computational Biology*, 4, 11.
47. Ziemke, T., Bergfeldt, N., Buason, G., Susi, T., & Svensson, H. (2004). Evolving cognitive scaffolding and environment adaptation: A new research direction for evolutionary robotics. *Connection Science*, 16(4), 339–350.