

Generative Representations for the Automated Design of Modular Physical Robots

Gregory S. Hornby, Hod Lipson, *Member, IEEE*, and Jordan B. Pollack

Abstract—The field of evolutionary robotics has demonstrated the ability to automatically design the morphology and controller of simple physical robots through synthetic evolutionary processes. However, it is not clear if variation-based search processes can attain the complexity of design necessary for practical engineering of robots. Here, we demonstrate an automatic design system that produces complex robots by exploiting the principles of regularity, modularity, hierarchy, and reuse. These techniques are already established principles of scaling in engineering design and have been observed in nature, but have not been broadly used in artificial evolution. We gain these advantages through the use of a generative representation, which combines a programmatic representation with an algorithmic process that compiles the representation into a detailed construction plan. This approach is shown to have two benefits: it can reuse components in regular and hierarchical ways, providing a systematic way to create more complex modules from simpler ones; and the evolved representations can capture intrinsic properties of the design space, so that variations in the representations move through the design space more effectively than equivalent-sized changes in a nongenerative representation. Using this system, we demonstrate for the first time the evolution and construction of modular, three-dimensional, physically locomoting robots, comprising many more components than previous work on body-brain evolution.

Index Terms—Design automation, evolutionary robotics, generative representations.

I. INTRODUCTION

THE field of evolutionary robotics has demonstrated the ability to automatically design both the morphology and controller of simple physical robots through synthetic evolutionary processes (e.g., [31], [36], [44]). Despite these results, it is not clear if these genetically inspired search algorithms can attain the design complexity necessary for practical engineering. The ultimate success of these methods as tools for design automation is critically dependent on the scaling properties of the

representations. Representations in which each element of the encoded design are used at most once in translating to the design (nongenerative representations) scale linearly with the number of parts in the artifact. Consequently, search algorithms that use this style of representation will quickly become exponentially intractable, and thus will not scale to complex tasks.

In this paper, we seek ways to circumvent this fundamental restriction with the automated design of robots by using a generative representation for encoding each robot. A robot is defined by a compact programmatic form (its generative representation) and the evolutionary variation takes place on this form. Evaluation requires the representation to be compiled into a detailed construction plan for manufacturing the robot in simulation or reality. Thus, the generative representations are varied through the evolutionary search process, while the evaluation or fitness function is applied to the construction plans.

Compared to previous work in evolutionary robotics, generative representations demonstrate two advantages. First, a generative representation allows for the reuse of components in regular and hierarchical ways, providing a systematic way to create more complex modules out of simpler ones. We show this through our evolved designs. Second, the evolved generative representation may itself capture intrinsic properties of the design problem, so that variations in the representations move through the design space very effectively. We demonstrate this by showing that across many generations of machines, mutations on the generative representation, despite causing larger scale changes, are more productive than mutations on a nongenerative representation. Both of these aspects maintain the evolvability of the design while its complexity rises.

In this paper, we bring together ideas of automated design and generative representations, and apply them to the design of physical working robots for a limited physical domain. Although the robots we have built from this process are still very simple compared to human-engineered machines, their structure is more principled (regular, modular and hierarchical) compared to previously evolved machines of comparable functionality, and the virtual designs which are achieved by the system have an order of magnitude more moving parts. Moreover, we quantitatively demonstrate for this design space how the generative representation is capable of searching more efficiently than a nongenerative representation.

II. STRUCTURE OF THIS PAPER

We will begin with a brief background of evolutionary robotics and related work, and demonstrate the scaling problem with our own prior results. Next we propose the use of an

Manuscript received December 6, 2001; revised June 20, 2002. This paper was recommended for publication by Associate Editor R. Simmons and Editor I. Walker upon evaluation of the reviewers' comments. This work was supported in part by the Defense Advanced Research Projects Administration (DARPA) under Grant DASG60-99-1-0004. This paper was presented in part at the IEEE Conference on Robotics and Automation, ICRA 2001.

G. S. Hornby was with the Computer Science Department, Brandeis University, Waltham, MA 02454 USA. He is now with the QSS Group, Inc., Computational Sciences Division, NASA Ames Research Center, Moffett Field, CA 94035 USA (e-mail: hornby@email.arc.nasa.gov).

H. Lipson was with the Computer Science Department, Brandeis University, Waltham, MA 02454 USA. He is now with the Mechanical & Aerospace Engineering Department, Cornell University, Ithaca, NY 14853 USA (e-mail: Hod.Lipson@cornell.edu).

J. B. Pollack is with the Computer Science Department, Center for Complex Systems, Brandeis University, Waltham, MA 02454 USA (e-mail: pollack@cs.brandeis.edu).

Digital Object Identifier 10.1109/TRA.2003.814502




Fig. 1. Progress of a typical evolutionary design run comprising only direct mutations. The abscissa represents evolutionary time (generations) and the ordinate measures fitness. Each point in the scatter plot represents one candidate design. In general, a logarithmic slowdown in progress can be observed, characterized by longer and longer durations between successive step jumps in the fitness (a,b). Occasionally, however, progress is made after long periods of stagnation (c,d).

evolved generative representation as opposed to a nongenerative representation. We describe this representation in detail as well as the evolutionary process that uses it. We then compare progress of evolved robots with and without the generative representation and quantify the obtained advantage. Working two-dimensional (2-D) and three-dimensional (3-D) physical robots produced by the system are shown.

III. BACKGROUND

Biological evolution is characterized as a process applied to a population of individuals, which are subject to selective replication with variation [33]. Evolutionary design systems use the same principles of biological evolution to achieve machine design, yet add a target to the evolution—the functional requirements specified by the designer. Candidate designs in a population are thus still selected, replicated and varied, but selection is governed by an external design criteria. After a number of generations the selective evolutionary process may breed an acceptable design.

Genetic algorithms—a subset of evolutionary computation involving mutation and crossover in a population of fixed-length bit strings [14]—have been applied for several decades in many engineering problems as an optimization technique for a fixed set of parameters. Alternatively, more recent *open ended* evolutionary design systems, in which the process is allowed to add more and more building blocks and

parameters, seem particularly adequate to design problems requiring *synthesis*. Such open ended evolutionary design systems have been demonstrated for a variety of simple design problems, including structures, mechanisms, software, optics, robotics, control, and many others (for overviews see, for example, [4], [27]). Yet these accomplishments remain simple compared to what teams of human engineers can design and what nature has produced. The evolutionary design approach is often criticized as scaling badly when challenged with design requirements of higher complexity [32].

While there are still many poorly understood factors that determine the success of evolutionary design—such as starting conditions, variation operators, primitive building blocks and fidelity of simulation—one problem is that the design space is exponentially large, because there is an exponentially increasing number of ways a linearly increasing set of components can be assembled. Consequently, evolutionary approaches that operate on nongenerative representations quickly become intractable.

Indeed, while our own experiments in open ended evolutionary design of mechanisms and controllers for locomotion [31], have produced physically viable mechanisms, the design progress appears to eventually reach a plateau. Fig. 1(a) and Fig. 1(b) show the progress of a typical run. The abscissa represents evolutionary time (generations), the ordinate measures fitness (net movement on a horizontal plane) and each point in the scatter plot represents one candidate robot. In general, after an initial period of drift, with zero fitness, we

observe rapid growth followed by a logarithmic slowdown in progress,¹ characterized by longer and longer durations² between successive step improvements in the fitness. We thus hypothesize that one of the primary challenges in evolutionary robotics research is that of allowing the process to identify and reuse assemblies of parts, creating more complex components from simpler ones. This reuse would, in turn, lead to acceleration in the discovery process, leading in turn to higher level of search, and so forth. This scaling in knowledge and unit of construction—hierarchical modularity—is observable in the engineering [47], [48], economic organization [49], and in nature [50]. Theoretical analysis reveals that allowing an evolutionary process to repeatedly aggregate low-level building blocks into higher level groups in a hierarchical fashion, potentially transforms the design problem from exponential complexity to polynomial complexity under certain conditions [45]. The challenge then becomes the continuous identification of discovered components and the encapsulation of them as basic units of change in the variation operators. In other words, the design space in which the machines are specified and the effects of mutations and crossovers which move through the space need to evolve over time as well.

IV. GENERATIVE REPRESENTATIONS

Generative representations [15] are a class of representations in which elements in the encoded data structure of a design can be reused in creating the design. Implemented as a kind of a computer program, a generative representation can allow the definition of reusable subprocedures that can be activated in loops and recursive calls, allowing the design system to scale to more complex tasks, in fewer steps, than can be achieved with a nongenerative representation. Examples of representations that are generative are genetic programming with automatically defined functions [27] and cellular encoding [12], which has procedures and loops.

Here we use Lindenmayer systems (L-systems) as the generative representation for robot designs. L-systems are a grammatical rewriting system introduced to model the biological development of multicellular organisms [29]. Rules are applied in parallel to all characters in the string, just as cell divisions happen in parallel in multicellular organisms. Complex objects are created by successively replacing parts of a simple object by using the set of rewriting rules. A detailed specification of the L-system used in this paper follows in the next section.

L-systems and evolutionary algorithms have been used both on their own and together to create designs. L-systems have been used mainly to construct tree-like plants, [40]. However, it is difficult to manually design an L-system to produce a desired form. L-systems have been combined with evolutionary algorithms in previous work, such as the evolution of plant-like structures [22], [37], [40] and architectural floor designs [8], but

¹However, note that because of the stochastic nature of the process, it is hard to determine definitely whether progress has actually halted, and improvements may still occur after long periods of apparent stagnation (Figs. 1(c) and 1(d)).

²This real-time lingering is amplified by the fact that evaluation time or duration of a generation (in simulation or in physical reality) also increases as solutions become more complex.

only limited results have been achieved, and none have resulted in dynamic physical machines comprising any form of control.

Robot morphology and controllers have been automatically designed both separately and concurrently. Genetic algorithms have been used to evolve a variety of different control architectures for a fixed morphology, such as: stimulus-response rules [35]; neural controllers [13]; and gait parameters [18]. Again using genetic algorithms, serial manipulators have been evolved by evaluating the ability of their end manipulator to achieve a set of configurations [6], [7], [25], and [38], and tree-structured robots have been evolved that met a set of requirements on static stability, power consumption and geometry [9], [10]. Leger’s *Darwin2K* [28] uses fixed control algorithms to evaluate evolved robot morphologies. Unlike the previous systems, which did not allow for reuse or the hierarchical construction of modules, *Darwin2K* has a kind of abstraction that allows the same assembly of parts to be reused. But the ability to design controllers is necessary for evaluating robots for more complex tasks or in dynamic environments [19], [39], [41], [46].

Concurrent development of robot morphology and controllers has been achieved previously by Sims [44], Komosinski and Rotaru–Varga [26] and ourselves [31], all of which used evolutionary algorithms to simultaneously create the morphology and a neural controller in simulation. Whereas Sims and Komosinski *et al.* were not concerned with the feasibility of their creations in reality, the focus of our own line of work is to show that robots created through evolution in simulation could be successfully transferred to reality. This paper extends our initial results by investigating generative representations as a mechanism to overcome the complexity barrier.

V. METHOD

We used four levels of computation, as follows.

- 1) An *evolutionary process* that evolves generative representations of robots.
- 2) Each generative representation is an *L-System program* that, when compiled, produces a sequence of *build commands*, called an *assembly procedure*.
- 3) A *constructor* executes an assembly procedure to generate a *robot* (both morphology and control).
- 4) A physical *simulator* tests a specific robots’ performance according to a fitness criteria, to yield a figure of merit that is fed back into the evolutionary process (1).

We will describe each of the above four levels.

VI. CONSTRUCTOR AND DESIGN LANGUAGE

The design constructor builds a model from a sequence of build commands. The language of build commands is based upon instructions to a LOGO-style turtle, which direct it to move forward, backward or rotate about a coordinate axis. The commands are listed in Tables I and II. Robots (called “Genobots,” for generatively encoded robots) are constructed from rods and joints, Fig. 2, that are placed along the turtle’s path. Actuated joints are created by commands that direct the turtle to move

TABLE I
DESIGN LANGUAGE FOR 2-D ROBOT

Command	Description
[]	Push/pop orientation stack
{ block }(n)	Repeat enclosed block n times
Forward(n)	moves the turtle forward in the current direction, creating a bar if none exists or traversing to the end of the existing bar
Back(n)	Move up n levels of parents
Joint(n)	Forward, end with an actuated joint which oscillates at speed n
Clockwise(n)	Rotate heading clockwise $n \times 60^\circ$
CounterClockwise(n)	Rotate heading counterclockwise $n \times 60^\circ$
IncreaseOffset(n)	Increase current joint phase offset by $n \times 25\%$
DecreaseOffset(n)	Decrease current joint phase offset by $n \times 25\%$

TABLE II
DESIGN LANGUAGE FOR 3-D ROBOT

Command	Description
[]	Push/pop orientation stack
{ block }(n)	Repeat enclosed block n times
Forward(n)	moves the turtle forward in the current direction, creating a bar if none exists or traversing to the end of the existing bar
Back(n)	Move up n levels of parents
Revolute1(n)	Forward, end with a joint with range 0° to 90° about the current Z-axis that oscillates with speed n .
Revolute2(n)	Forward, end with a joint with range -45° to $+45^\circ$ about the current Z-axis that oscillates with speed n .
Twist90(n)	Forward, end with a joint with range 0° to 90° about the current X-axis that oscillates with speed n .
Twist180(n)	Forward, end with a joint with range -90° to $+180^\circ$ about the current X-axis that oscillates with speed n .
Up(n)	Rotate heading n times 90° about the turtle's Z axis
Down(n)	Rotate heading n times -90° about the turtle's Z axis
Left(n)	Rotate heading n times 90° about the turtle's Y axis
Right(n)	Rotate heading n times -90° about the turtle's Y axis
Clockwise(n)	Rotate heading clockwise $n \times 90^\circ$ about the turtle's Z axis
CounterClockwise(n)	Rotate heading counterclockwise $n \times -90^\circ$ about the turtle's Z axis
IncreaseOffset(n)	Increase current joint phase offset by $n \times 25\%$
DecreaseOffset(n)	Decrease current joint phase offset by $n \times 25\%$

forward and place an actuated joint at its new location, with oscillatory motion and a given offset.

The operators “[” and “]” push and pop the current state—consisting of the current rod, current orientation, and current joint oscillation offset—to and from a stack. Forward moves the turtle forward in the current direction, creating a rod if none exists or traversing to the end of the existing rod. Backward goes back up the parent of the current rod. The rotation commands turn the turtle about the Z axis in steps of 60° , for 2-D robots, and about the X, Y or Z axes, in steps of 90° , for 3-D robot. Joint commands move the turtle forward, creating a rod, and end with an actuated joint. The parameter to these commands specify the speed at which the joint oscillates, using integer values from 0 to 5, and the relative phase offset of the oscillation cycle is taken from the turtle's state. The commands “Increase-offset” and “decrease-offset” change the offset value n the turtle's state by $\pm 25\%$ of a total cycle. Command sequences enclosed by “{ }” are repeated a number of times specified by the brackets' argument.

For example, the string

$$\{\text{joint}(1) [\text{joint}(1)\text{forward}(1)] \text{clockwise}(2)\} (3)$$



Fig. 2. Basic building blocks of the system: bars of regular lengths and fixes or actuated joints.

is interpreted as

$$\begin{aligned} &\text{joint}(1) [\text{joint}(1)\text{forward}(1)] \text{clockwise}(2) \\ &\text{joint}(1) [\text{joint}(1)\text{forward}(1)] \text{clockwise}(2) \\ &\text{joint}(1) [\text{joint}(1)\text{forward}(1)] \text{clockwise}(2) \end{aligned}$$

and produces the robot in Fig. 3.

Constructed robots do not have a central controller; rather each joint oscillates independent of the others. More recent work has integrated a recurrent neural network as the robot controller [20]. In the figures, large crosses (‘x’) are used to show the location of actuated joints and small crosses show unactuated joints. The left image shows the robot with all actuated joints in their starting orientation and the image on the right shows the same robot with all actuated joints at the other extreme of their actuation cycle. In this example all actuated joints are moving in phase.

VII. PARAMETRIC OL-SYSTEMS

The class of L-systems used as the representation is a parametric, context-free L-system (POL-system). Formally, a POL-system is defined as an ordered quadruplet, $G = (V, \Sigma, \omega, P)$ where

$$\begin{aligned} V & \text{ is the alphabet of the system} \\ \Sigma & \text{ is the set of formal parameters} \\ \omega \in (V \times \mathbb{R}^*)^+ & \text{ is a nonempty parametric word called the axiom} \end{aligned}$$

and $P \subset (V \times \Sigma^*) \times C(\Sigma) \times (V \times \xi(\Sigma))^*$ is a finite set of productions.

The symbols “:” and “ \rightarrow ” are used to separate the three components of a production: the predecessor, the condition and the successor. For example, a production with predecessor $A(n0, n1)$, condition $n1 > 5$ and successor $B(n1 + 1)cD(n1 + 0.5, n0 - 2)$ is written as

$$A(n0, n1) : n1 > 5 \rightarrow B(n1 + 1)cD(n1 + 0.5, n0 - 2)$$

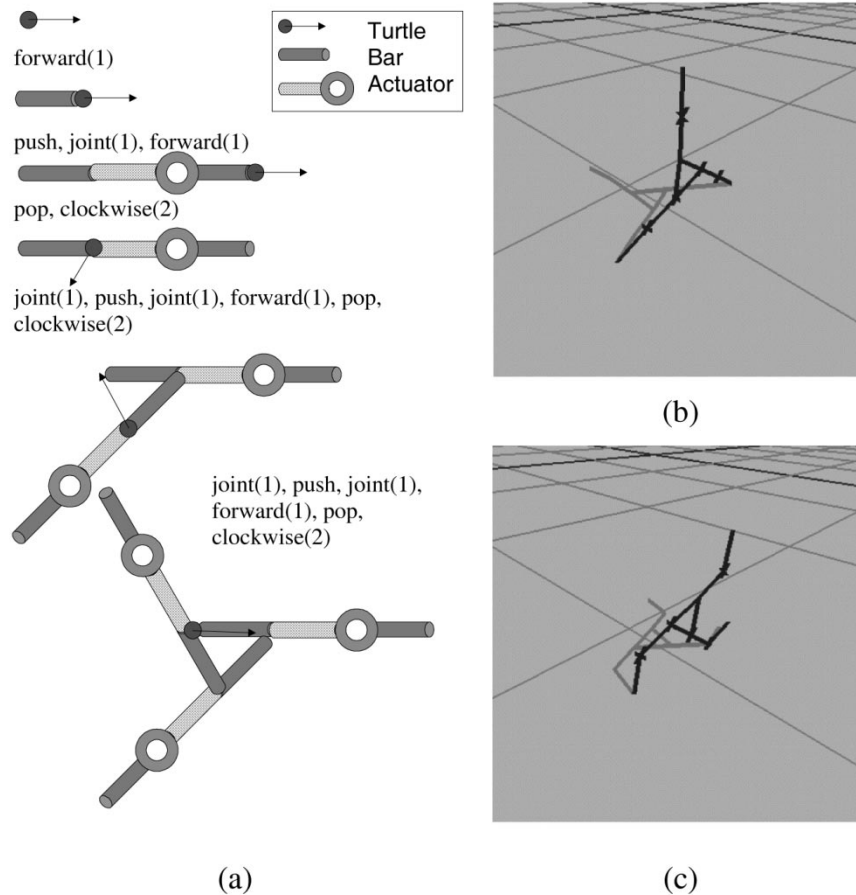


Fig. 3. Sample L-Robot. (a) Construction sequence leading to a tristar 2-D robot with three actuated joints. (b) Resulting robot, with joints at 180°. (c) With joints at 120°.

A production matches a module in a parametric word if and only if the letter in the module and the letter in the production predecessor are the same, the number of actual parameters in the module is equal to the number of formal parameters in the production predecessor, and the condition evaluates to true if the actual parameter values are substituted for the formal parameters in the production.

For ease of implementation we add constraints to our POL-system. The condition is restricted to be comparisons as to whether a production parameter is greater than a constant value. Parameters to design commands are either a constant value or a production parameter.

Parameters to productions are equations of the form

$$\frac{[\text{production parameter}|\text{constant}][+|-|\times|/]}{[\text{production parameter}|\text{constant}]}$$

The following is a POL-system using the language defined in Table I and consists of two productions with each production containing two condition-successor pairs

$$\begin{aligned} P0(n) : n > 2 &\rightarrow \{P0(n-1)\}(n) \\ n > 0 &\rightarrow \text{joint}(1)P1(n \times 2)\text{clockwise}(2) \\ P1(n) : n > 2 &\rightarrow [P1(n/4)] \\ n > 0 &\rightarrow \text{joint}(1)\text{forward}(1). \end{aligned}$$

If the POL-system is started with $P0(3)$, the resulting sequence of strings is produced

$$\begin{aligned} &P0(3) \\ &\{P0(2)\}(3) \\ &\{\text{joint}(1)P1(4)\text{clockwise}(2)\}(3) \\ &\{\text{joint}(1)[P1(1)]\text{clockwise}(2)\}(3) \\ &\{\text{joint}(1)[\text{joint}(1)\text{forward}(1)]\text{clockwise}(2)\}(3). \end{aligned}$$

This produces the robot in Fig. 3.

VIII. THE EVOLUTIONARY PROCESS

An evolutionary algorithm is used to evolve individual L-systems. Evolutionary algorithms are a stochastic search and optimization technique inspired by natural evolution [14], [1]. An evolutionary algorithm maintains a population of candidate solutions from which it performs search by iteratively replacing poor members of the population with individuals generated by applying variation to good members of the population. The initial population of L-systems is created by making random production rules. Evolution then proceeds by iteratively selecting a collection of individuals with high fitness for parents and using them to create a new population of individual L-systems through mutation and recombination. Since the initialization process and

the variation operators are customized for our representation we discuss them in greater detail.

Mutation creates a new individual by copying the parent individual and making a small change to it. Changes that can occur are: replacing one command with another; perturbing the parameter to a command by adding/subtracting a small value to it; changing the parameter equation to a production; adding/deleting a sequence of commands in a successor; or changing the condition equation.

Recombination takes two individuals, $p1$ and $p2$, as parents and creates one child individual, c , by making it a copy of $p1$ and then inserting a small part of $p2$ into it. This is done by replacing one successor of c with a successor of $p2$, inserting a subsequence of commands from a successor in $p2$ into c , or replacing a subsequence of commands in a successor of c within a subsequence of commands from a successor in $p2$. Details of the mutation and recombination operators used here, as well as other evolutionary algorithm parameters, are described in an earlier report [17].

Since we are evolving machines that can locomote, fitness was defined as the distance moved by the robot's center of mass after 10 simulated oscillation cycles. This distance is normalized by one tenth of the length of a basic bar, to produce a dimensionless figure.

The evolutionary algorithm described in this section is essentially a canonical evolutionary algorithm, differing only in the representation. First, a population size of one hundred individuals is used and this is run for five hundred generations. Once individuals are evaluated, their fitness score is adjusted to a probability of reproducing using exponential scaling with a scaling factor of 0.03 [34]. Individuals are then selected as parents using stochastic remainder selection [2]. New individuals are created through applying mutation or recombination (chosen with equal probability) to individuals selected as parents. The two best individuals from each generation were copied directly to the next generation without mutation or crossover (elitism of 2).

Initializing the population consists of creating a set of random L-systems. A random L-system has a fixed number of production rules, with each having the same number of condition successor pairs, so creating a new individual consists of creating a number of random conditions and successors. A new condition is created by selecting a random parameter and a random value to compare against. New successors are created by stringing together sequences of one to three construction symbols, with each sequence being enclosed with push/pop brackets, block replication parenthesis, or neither. Examples of these blocks of commands are

$$\begin{aligned} & \text{Forward}(n0) \text{ left}(4.0) \\ & \{ \text{up}(1.0) \} (2.0) \\ & [\text{back}(2.0) \text{ down}(1.0) \text{ joint}(n1)]. \end{aligned}$$

After a new L-system is created it is evaluated. Individuals that score below a minimum fitness value are deleted and a new one is randomly created. In this way, all robots in the initial population have a minimum degree of viability.

The mutation operator creates a new robot encoding by copying an existing robot encoding and making a random change to it. This is implemented by selecting one of the

production rules at random and changing its condition or successor. For example, if the condition P4 is selected to be mutated

$$\begin{aligned} P4(n0, n1) : & - (n1 > 6.0) [P1(n0 - 1.0, n1/3.0)] \\ & (n0 > 2.0) \{ \text{left}(1.0) \text{ forward}(2.0) \} (n0) \end{aligned}$$

then some of the possible mutations are as follows.

Mutate an argument to a construction command

$$\begin{aligned} P4(n0, n1) : & - (n1 > 6.0) [P1(n0 - 1.0, \mathbf{n1/2.0})] \\ & (n0 > 2.0) \{ \text{left}(1.0) \text{ forward}(2.0) \} (n0). \end{aligned}$$

Delete random command(s)

$$\begin{aligned} P4(n0, n1) : & - (n1 > 6.0) [P1(n0 - 1.0, n1/3.0)] \\ & (n0 > 2.0) \{ \mathbf{forward}(2.0) \} (\mathbf{n0}). \end{aligned}$$

Insert a random block of command(s)

$$\begin{aligned} P4(n0, n1) : & - (n1 > 6.0) [P1(n0 - 1.0, n1/3.0) \\ & \{ \mathbf{up}(1.0) \text{ back}(2.0) \} (\mathbf{2.0})] \\ & (n0 > 2.0) \{ \text{left}(1.0) \text{ forward}(2.0) \} (n0). \end{aligned}$$

The other method of creating new design encodings is recombination, which takes two individuals as parents and creates a new child individual by copying the first parent and inserting a small part of the second parent into it. This insertion can replace an entire condition-successor pair, just the successor, or a subsequence of commands in the successor. For example, if P4 is selected to be changed and in the first parent it is

$$\begin{aligned} P4(n0, n1) : & - (n1 > 6.0) [P1(n0 - 1.0, n1/3.0)] \\ & (n0 > 2.0) \{ \text{left}(1.0) \text{ forward}(2.0) \} (n0) \end{aligned}$$

and in the second parent it is

$$\begin{aligned} P4(n0, n1) : & - (n0 > 4.0) \text{ forward}(1.0) \text{ joint}(2.0) \\ & (n0 > 3.0) P2(n1 - 1.0, n1 - 2.0) \\ & [\text{up}(2.0) \text{ joint}(3.0)] \end{aligned}$$

then some of the possible results of recombination are as follows.

Replace an entire condition-successor pair

$$\begin{aligned} P4(n0, n1) : & - (\mathbf{n0 > 3.0}) P2(\mathbf{n1 - 1.0, n1 - 2.0}) \\ & [\mathbf{up}(2.0) \text{ joint}(3.0)] \\ & (n0 > 2.0) \{ \text{left}(1.0) \text{ forward}(2.0) \} (n0). \end{aligned}$$

Replace just a successor

$$\begin{aligned} P4(n0, n1) : & - (n1 > 6.0) [P1(n0 - 1.0, n1/3.0)] \\ & (n0 > 2.0) \mathbf{forward}(1.0) \text{ joint}(2.0). \end{aligned}$$

Replace one sequence of commands with another

$$\begin{aligned} P4(n0, n1) : & - (n1 > 6.0) [P1(n0 - 1.0, n1/3.0)] \\ & (n0 > 2.0) \{ \text{left}(1.0) [\mathbf{up}(2.0) \text{ joint}(3.0)] \} \\ & (n0). \end{aligned}$$

IX. SIMULATOR

To evaluate a robot design it is simulated in a quasistatic simulator. The simulation consists of moving the actuated joints in small angular increments of up to 0.001 radians (depending on joint speed). After each update, the robot is settled by calculating the location of its center of mass and then repeatedly rotating the robot about the edge of its footprint nearest to the center of mass until it is stable. In performing this simulation, masses of the different connectors are taken into account for

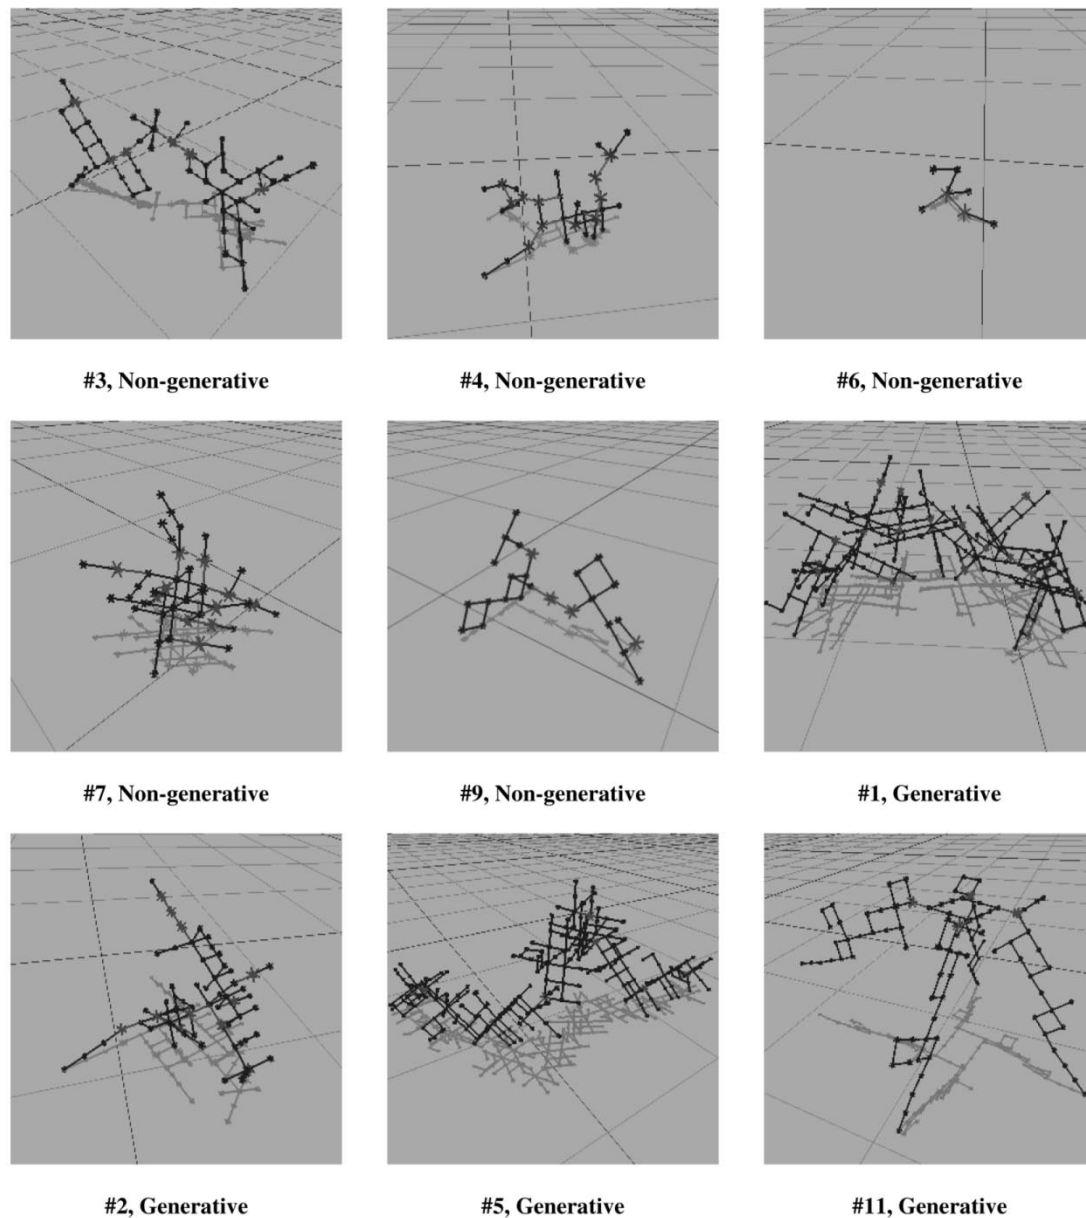


Fig. 4. Evolved robots through nongenerative and generative representations. (For full motion see videos at http://www.demo.cs.brandeis.edu/pr/evo_design/evo_design.html).

calculating the center of mass, but not in calculating inertia. Torques are not calculated, and power considerations are undefined in quasi-static motion.

To create robot designs that are robust to imperfections in real-world construction, error is added to the simulation similar to the method of Hornby *et al.* [19] and Jakobi [23]. This consists of simulating a robot design three times, once without error and twice with different error values applied to joint angles. Error consists of adding a random rotation (in the range of ± 0.1 radians about each of the three coordinate axis) to every joint that is not part of a cycle. A robot's fitness is the minimum score from these three trials.

The assumptions made in these simulations are geared toward making a simulator that is robust and fast, and sufficiently realistic so that results produced will transfer well into reality. Quasi-static simulation also eliminates the need to accurately

model masses, inertias, torques and power, and avoids real-time control issues entirely when transferring to reality. This assumption of low momentum was indeed justified in light of our results, however, inevitably more realistic and complex tasks will require more realistic and higher fidelity simulation.

X. RESULTS

We now describe how our system has been used to create modular designs that locomote in simulation and were shown to work in reality. To show that a generative representation has better scaling properties and captures intrinsic properties of the problem we ran a number of experiments with both a nongenerative representation and a generative representation. The nongenerative representation consisted of a string of up to 10 000 build commands. The generative representation

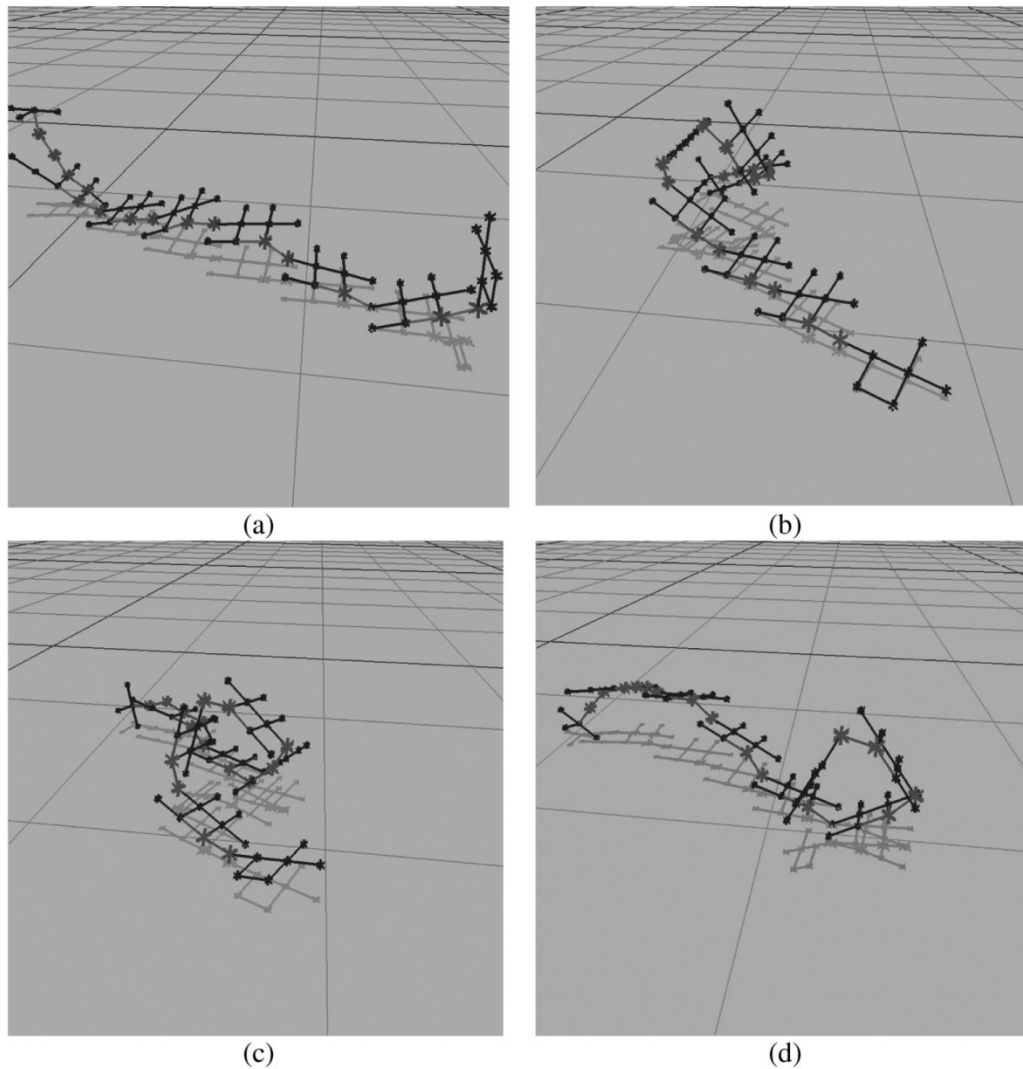


Fig. 5. The locomotion cycle of a 80-rod robot evolved using generative representation and reaching a fitness of 686. The robot moves by passing a loop from back to front. Frames (a–d) show four stages in the locomotion cycle. (For full motion, see video at http://www.demo.cs.brandeis.edu/pr/evo_design/evo_design.html).

used a POL-system with up to 15 productions, each with two parameters and three sets of condition-successor pairs. The maximum number of commands in each condition-successor pair is 15 and maximum length of a command string generated by the L-system is 10 000 build commands. In these runs fitness is defined as the distance moved by the robot's center of mass after 10 simulated oscillation cycles, with the constraint that robots could not have a sequence of more than 4 bars in a row that was not part of a structural loop, as a structural constraint.³

Evolutionary runs were similar in many ways. The first individuals started with a few rods and joints and would slowly slide on the ground. Robots produced from runs using the nongenerative representation would improve upon their sliding motion over the course of the evolutionary run.

Approximately half the runs produce “interesting” viable results. Nine selected machines out of 20 runs are shown in Fig. 4. The two main forms of locomotion found used one or more oscillating appendages to push along, or had two main body parts connected by a sequence of rods that twisted in such a way that

first one half of the robot would rotate forward, then the other. The two fastest robots evolved with the nongenerative representation, shown in Fig. 4, are #3 (fitness 1188 with 49 rods and moves by twisting) and #7 (fitness 1000 with 31 rods which moves by pushing). Half the robots had only a handful of rods, such as robot #6. Robot #4 is an example of one that uses its appendages to roll over.

Robots evolved with the generative representation not only had higher average fitness, but tended to move in a more continuous manner. Here the two fastest were #1 (a sequence of interlocking X's that rolls along with fitness 2754 and 268 rods) and #5 (whose segments are shaped like a coil and it moves by rolling sideways with fitness of 3604 and 325 rods). Not all were regular as demonstrated by robot #2, an asymmetric robot that moves by sliding along similar to many of the robots generated with a nongenerative representation (fitness 766 and 63 rods) Robot #11 is a four-legged walker with fitness 2170 and 109 rods. An example of the movement cycle of a robot produced with the generative representation is shown in Fig. 5(a)–(d). This robot has fitness 686 and 80 rods and moves by passing a loop from back to front.

³In a true dynamics simulator actual torques on joints would be calculated and then a constraint on the allowable torque could be used instead.




Fig. 6. Two parts of the locomotion cycle of an evolved 2-D robot “M”. (a, b) Simulated. (c, d) Physical. Notice regularity and symmetry.




Fig. 7. An evolved 3D robot “Kayak”. (a) Simulated. (b) Physical. Notice the reuse of a T-junction.

In general, robots evolved using the generative representation increased their speed by repeating rolling segments to smoothen out their gaits, and increasing the size of these segments/appendages to increase the distance moved in each oscillation.

Finally, we construct some of the evolved Genobots.⁴ We picked three robots that were small (in number of parts and motors) and easy to build. The “M” robot in Fig. 6 shows two parts of the locomotion cycle of the 2-D robot. With its two outer arms evolved to be 25% out of phase, it moves by bringing them together to lift its middle arm and then to shift its center of mass to the right. One modification to the constructed robot is the addition of sandpaper on the feet of the two outer arms to compensate for the friction modeled by our simulator and that of the actual surface used.

The “Kayak” robot (so called because of its form of locomotion, see video) in Fig. 7 moves by curling up in alternating directions. At the extreme stages of these oscillations, the robot contacts the ground at only three points, freeing its fourth point to move forward. Repeated oscillations gradually push the machine forward while its rear support is dragged across the ground.

⁴The “M” and “Kayak” robots were evolved in runs that did not use the structural constraint limiting the maximum number of rods without a structural loop.

The “Quatrobot” robot in Fig. 8 has one actuated joint on each of its four legs and walks by raising and lowering its legs. Each leg is 12.5% out of phase with the ones next to it and the robot moves in the direction of the lead leg. Instead of constructing this robot out of the components in Fig. 2, it was manufactured using rapid prototyping equipment in a manner similar to our original robots [31].

XI. SCALING

The results shown so far demonstrate that the generative system is capable of producing nontrivial robot designs that transfer well into reality. We now address the question of scaling—the *progress* of performance of the design process over evolutionary time. This aspect determines whether evolutionary robotics might ultimately be used as a practical engineering tool.

Fig. 9 examines the progress of average population fitness over the generations, averaged over 10 evolutionary runs. Each run used a population of 100 individuals and was evolved for 500 generations. The lower curve corresponds to the nongenerative representation, while the upper curve shows progress for generative representation over the same substrate. It is evident that the generative representation makes faster progress at the




Fig. 8. An evolved 3D robot “Quatrobot”. (a) Simulated. (b) Physical. Notice the reuse of the leg assembly. For video of the robots in motion see http://www.demo.cs.brandeis.edu/pr/evo_design/evo_design/html.

initial stage: Fitness increase per generation is more than 6 times faster in the first 100 generations, as shown by the linear dashed lines. In that short period, the generative system consistently outperforms the nongenerative representation, which does not reach an equivalent level within the entire experiment.⁵ After around 200 generations, the rate of progress with the generative representation is slowed down, yet still is approximately twice the rate of the nongenerative representation.

Fig. 10 and Fig. 11 show some properties of the design space covered by the search process. The plots show designs visited in terms of their fitness and complexity: Fig. 10 approximates complexity in terms of the physical parts count, and Fig. 11 approximates complexity in terms of length of the generating program (either the construction sequence or generative program), in number of language elements. There seems to be a correlation between the fitness and size of the robot (in physical parts); however this might be a particular of the locomotion problem. The two sets of plots show, on one hand, that the nongenerative representation searches the space more thoroughly, as it finds better solutions with low complexity. Yet it does not reach far enough into the space to explore more complex designs. Because fitness is correlated with complexity (or size) for this problem, this shortcoming prevents it from finding solutions with high fitness. Although strict or even near global optimality does not seem to be consistently attainable for this 200–600 dimension search space, the generative system is able to explore much fitter solutions through much shorter programs, as shown in Fig. 10.

One of the fundamental questions is whether the actual grammar evolved in the successful L-systems has captured some of the intrinsic properties of the design space. One way to quantify this is to measure the correlation between fitness change and a random mutation of various sizes, and compare this with the correlation observed in random mutations on the nongenerative representation as a control experiment. If the observed correlation is distinguishable and better for the generative system than it is for the blind system, then the generative system must have captured some useful properties.

The graph in Fig. 12 is a comparison of the fitness-mutation correlation between a generative representation and a random control experiment on the same substrate and on the same set of

⁵Each experiment comprises 10 runs, with run evolving a population of 100 robots over 500 generations. In all 500 000 robot simulations were run in a period of approximately six months of computation time.




Fig. 9. Fitness over time, comparing nongenerative representation with generated representation. Data averaged over 10 runs.

randomly selected individuals. For this analysis, 80 000 individuals were selected uniformly from 16 runs and over 100 generations using a generative representation. Each point represents a particular fitness change (positive or negative) associated with a particular mutation size. The points on the left plot (Fig. 12(a)) were carried out on the nongenerative representation generated by the generative representation and serve as the control set. For these points, one to six mutations were applied so as to approximate mutations of similar phenotypic size as those on the generative representation. Each mutation could modify or swap a sequence of characters. The points on the right (Fig. 12(b)) were also carried out randomly but on the generative representations of the same randomly selected individuals. Only a single mutation was applied to the generative representation, and consisted of modifying or swapping a single keyword or parameter. Mutation size was measured in both cases as the number of modified commands in the final construction sequences.

The two distributions in Fig. 12 have distinct features. The data points separate into two distinguishable clusters, with some overlap. Mutations generated on the generative representations clearly correlate with both positive fitness and negative fitness changes, whereas most mutations on the nongenerative representation result in fitness decrease. Statistics of both systems,

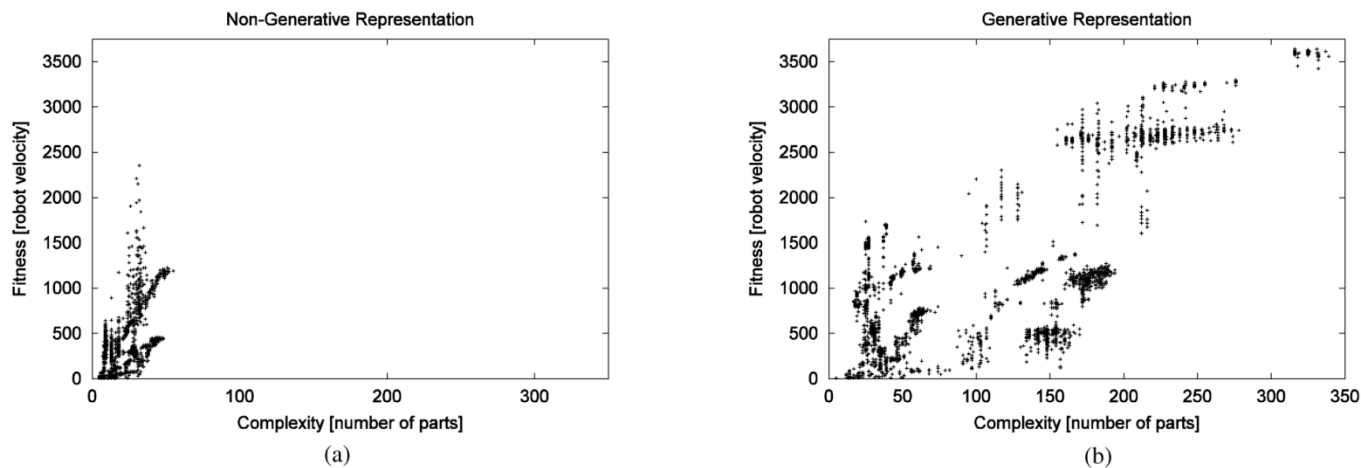


Fig. 10. Comparison of fitness versus number of parts. (a) Nongenerative representation. (b) Generative representation. One dot for the best individual of each generation for all 10 runs.

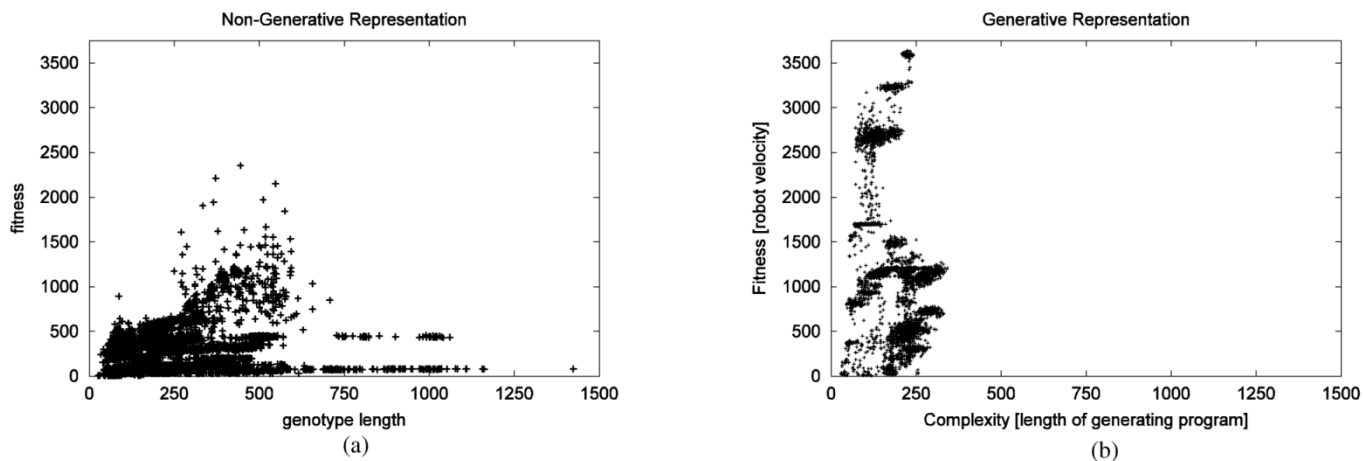


Fig. 11. Comparison of fitness versus length of generating program, measured in number of elements in encoded design. (a) Nongenerative representation. (b) Generative representation. One dot for the best individual of each generation for all ten runs.




Fig. 12. Comparison of fitness change per mutation. (a) Nongenerative representation. (b) Generative representation.

averaged over eight runs each, are summarized in Table III. A one-way ANOVA test revealed that the two means are different with at least 95% confidence. Cross correlation showed that in 40% of the instances where a nongenerative mutation was

successful, a generative mutation was also successful, whereas in only 20% of the instances where a generative mutation was successful, was a nongenerative mutation successful too. In both cases, smaller mutations are significantly more successful

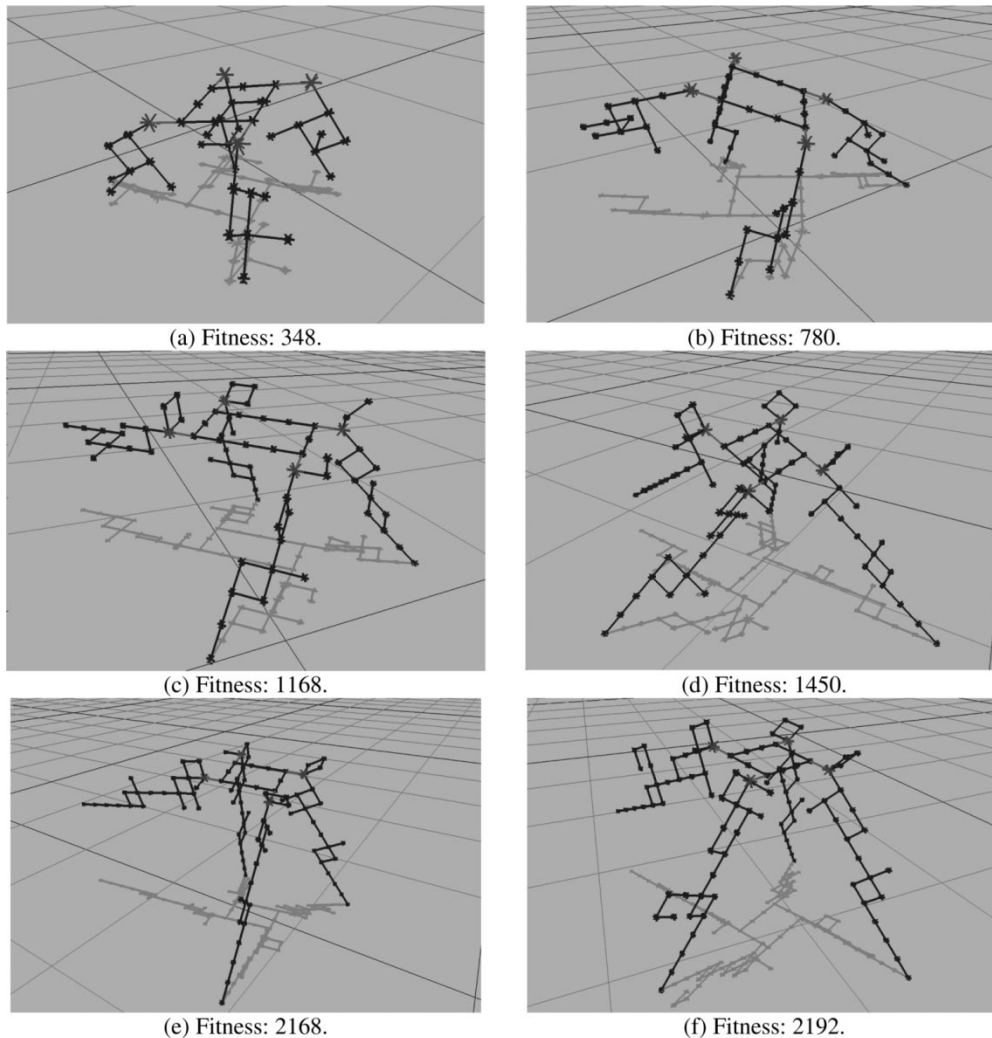


Fig. 13. Evolution of a four-legged walking robot.

TABLE III
EVOLVABILITY COMPARISON

Statistics for the experiments	Non-generative (Control)	Generative
Average fitness change	-502	-173
Standard deviation of fitness changes	± 736	± 404
% Successful mutations	11%	23%
Average fitness change of successful mutations	+104	+124
% Cross correlation of success	20%	40%
Success rate of large mutations (>100 characters)	2.3%	17%

than larger mutations. However, large mutations (>100) were an order of magnitude more often successful in the generative case than in the nongenerative case. All these measures indicate that the generative representation is more efficient in exploiting useful search paths in the design space.

An example of how the generative representation is being used to make coordinated changes can be seen by individuals taken from different generations of one of the evolutionary runs. The sequence of images in Fig. 13, which are the best individual in the population taken from different generations, show two

changes occurring. First, the rectangle that forms the body of the robot goes from being two-by-two, to three-by-three, to two-by-three. This change is possible with a single change of the generative representation but requires multiple coordinated changes on the nongenerative representation. The second change is the evolution of the robot's legs. Even though the legs are changing from image to image, all four legs are the same in each of the six images. As with the body, changing all four legs simultaneously can be done easily with the generative representation by changing the one module that is used for constructing them, but would require simultaneously making the same change to all four occurrences of the leg assembly procedure in a nongenerative representation.

XII. CONCLUSIONS

The purpose of the work reported in this paper is twofold. First, to demonstrate the ability of an evolutionary process to design both the morphology and control of a physically viable 3-D robot that, while it has the same basic locomotion functionality as previously evolved machines, demonstrates a much better form of design. In this paper, the evolved designs are an order of magnitude more complex, in terms of the number of given basic


building blocks, than previous work. Both the evolved designs, as well as the physically realized machines, show a significantly higher degree of modularity, regularity and hierarchy than previous machines whose morphology and control were generated fully automatically.

The second goal of this paper is to investigate the scaling properties of generative systems when applied to a robotic design problem. We have shown that the use of a generative representation has significantly accelerated the rate of progress at early stages of the design search, reaching levels that are not reached through direct mutation on the same design space and through the same number of evaluations. Most importantly, we have shown that at least for this design space, a generative representation is significantly more efficient in exploiting useful search paths in the design space than a nongenerative representation.

While the physically realized results shown here do not compare to design capabilities of teams of human engineers, we point out that the field is still very young. Naturally, more complex machines that can accomplish more complex tasks can be attained merely by starting with more advanced building blocks and more sophisticated genetic operators; but the eventual inventiveness of a design system is not measured by its final outcome, but rather by the distance of its outcome from its starting point. We believe that careful inclusions of fundamental design principles such as modularity, regularity, and hierarchy into self-organizing stochastic design processes like evolution can ultimately lead to powerful design automation tools for robots which can prosper in the real world.

APPENDIX

GENERATIVE REPRESENTATION for ROBOT “M”

Generative Representation	Construction Sequence
<p>This generative representation is started with P0(3.0,3.0) and run for ten iterations. (T) indicates that the condition always succeeds.</p> <p>P0: (n0>3.0) :- clockwise(1.0) cntr-clockwise(n1) [clockwise(1.0) clockwise(1.0) back(5.0)] (n0>3.0) :- clockwise(1.0) cntr-clockwise(n1) [clockwise(1.0) clockwise(1.0) back(5.0)] (T) :- [cntr-clockwise(1.0) P7(n0+1.0,n1-n0)]</p> <p>P1: (n0>3.0) :- [clockwise(1.0)] (n0>3.0) :- [clockwise(1.0)] (T) :- P0(n1,n0-1.0) clockwise(1.0) {forward(1.0) }(2.0)</p> <p>P3: (n1>5.0) :- clockwise(1.0) decrease-offset(1.0) P8(5.0,n0/5.0) decrease-offset(1.0) (n1>1.0) :- clockwise(1.0) cntr-clockwise(n1) [clockwise(1.0) clockwise(1.0) back(5.0)] (T) :- clockwise(1.0) cntr-clockwise(n1) clockwise(1.0) clockwise(1.0) back(5.0)</p> <p>P5: (n1>3.0) :- {joint(5.0) }(3.0) P6(n1-2.0,n0/n1) back(1.0) (n1>3.0) :- {joint(5.0) }(3.0) back(1.0) P6(n1-2.0,n0/n1) (T) :- P12(n1/3.0,5.0) back(1.0)</p> <p>P6: (n1>3.0) :- [clockwise(1.0) clockwise(n1)] (n1>3.0) :- [clockwise(1.0) clockwise(n1)] (T) :- forward(2.0) decrease-offset(1.0) cntr-clockwise(1.0) P13(n1+1.0,n1/n0) forward(1.0)</p> <p>P7: (n1>1.0) :- {cntr-clockwise(1.0) decrease-offset(1.0) }(1.0) (n1>1.0) :- {cntr-clockwise(1.0) decrease-offset(1.0) }(1.0) (T) :- clockwise(1.0) [cntr-clockwise(1.0)] P9(n1+1.0,n1+1.0) P3(n1-4.0,n1+1.0)</p> <p>P9: (n1>1.0) :- [P11(1.0,n0+2.0) P10(n0+5.0,n0) cntr-clockwise(1.0)] (n1>1.0) :- [P11(1.0,n0+2.0) P10(n0+5.0,n0) cntr-clockwise(1.0)] (T) :- clockwise(n0) increase-offset(1.0) cntr-clockwise(5.0) P3(3.0,3.0) P6(n1-1.0,1.0)</p> <p>P13: (n1>4.0) :- [P14(n0-n1,n1+2.0) cntr-clockwise(1.0)] (n1>4.0) :- [P14(n0-n1,n1+2.0) cntr-clockwise(1.0)] (T) :- [P5(n0,n1+4.0)][P1(n1+n0,5.0) cntr-clockwise(3.0)]</p>	 <p>[cntr-clockwise(1) clockwise(1) [cntr-clockwise(1)] clockwise(1) offset-increase(1) cntr-clockwise(5) clockwise(1) cntr-clockwise(3) [clockwise(1) clockwise(1) back(5)] forward(2) offset-decrease(1) cntr-clockwise(1) [joint(5) joint(5) joint(5) forward(2) offset-decrease(1) cntr-clockwise(1) [joint(5) joint(5) joint(5) forward(2) offset-decrease(1) cntr-clockwise(1) forward(1) back(1)] [clockwise(1) cntr-clockwise(0) [clockwise(1) clockwise(1) back(5)] clockwise(1) forward(1) forward(1) cntr-clockwise(3)] forward(1) back(1)] [clockwise(1) cntr-clockwise(2) [clockwise(1) clockwise(1) back(5)] clockwise(1) forward(1) forward(1) cntr-clockwise(3)] forward(1) clockwise(1) cntr-clockwise(1) clockwise(1) clockwise(1) back(5)]</p> <p>See Table I for meaning of commands.</p>

GENERATIVE REPRESENTATION for ROBOT "Kayak"

Generative Representation

This generative representation is started with P0(2.0, 6.0) and is run for fourteen iterations.

P0:

```
(n1>6.0) :- offset-increase(n1) P3(5.0,n1-5.0) [down(1.0)
P7(n0,n0-1.0) ]
(n0>6.0) :- offset-increase(3.0) P3(5.0,n1-5.0) [cntr-
clockwise(n0) P12(n0,n0-1.0) cntr-clockwise(n0)
P12(n0,n0-1.0) ]
(n1>0.0) :- offset-increase(3.0) P3(5.0,n1-5.0) [down(1.0)
cntr-clockwise(n0) down(1.0) cntr-clockwise(n0)
P12(n0,n0-1.0) ]
```

P1:

```
(n1>5.0) :- [back(1.0) ] P1(n1/4.0,3.0) clockwise(1.0)
(n0>4.0) :- {revolute2(n0) }(2.0) right(1.0) up(1.0) [up(1.0)
down(5.0) offset-decrease(1.0) ]
(n1>1.0) :- {revolute2(n0) }(2.0) right(1.0) up(1.0) [up(1.0)
down(5.0) down(1.0) ] [up(1.0) down(4.0) offset-
decrease(1.0) ]
```

P2:

```
(n1>0.0) :- up(1.0) P12(n1/5.0,n1/n0) revolute(1.0)
(n1>1.0) :- up(1.0) P14(n1/5.0,n1/n0) revolute(1.0)
(n1>0.0) :-
```

P3:

```
(n1>-3.0) :- left(4.0) P10(n0-3.0,n0-1.0) P2(n0/5.0,n0-5.0)
P1(n1-1.0,1.0) left(4.0)
(n0>3.0) :- left(4.0) P10(n0-3.0,n0-1.0) P2(n0/5.0,n0-5.0)
left(4.0)
(n1>4.0) :- offset-decrease(1.0) P2(n0/5.0,n0-5.0)
```

P4:

```
(n1>4.0) :- P6(n0+n1,n1=5.0) cntr-clockwise(n0) up(1.0)
P13(n1,n1/5.0) P1(4.0,5.0) right(1.0) up(n1)
(n0>4.0) :- cntr-clockwise(n0) P6(n0+n1,5.0) up(1.0)
P13(n1,n1/5.0) P1(4.0,5.0) right(1.0) up(n1) up(n1)
(n0>0.0) :- cntr-clockwise(n0) P6(n0+n1,5.0) up(1.0)
P13(n1,n1/5.0) P1(4.0,5.0) right(1.0) up(n0) up(n1)
```

P5:

```
(n0>-1.0) :- right(3.0) right(3.0) P4(n0+4.0,4.0)
(n0>0.0) :- offset-increase(3.0) right(3.0) P4(n0+4.0,4.0)
(n1>0.0) :- right(3.0) right(3.0) P4(n0+4.0,4.0)
```

P6:

```
(n0>8.0) :- [P9(2.0-3.0,n0-1.0) P10(n0+n1,n1-2.0) ]
(n0>-3.0) :- [P9(2.0-3.0,n0-1.0) P10(n0+n1,n1-2.0) ]
(n1>0.0) :- clockwise(n0) cntr-clockwise(n1)
```

P7:


```
(n1>0.0) :- left(1.0)
(n0>2.0) :-
(n0>0.0) :- {{down(1.0) }(n0) }(4.0)
```

P8:

```
(n0>5.0) :- cntr-clockwise(1.0) cntr-clockwise(3.0) up(1.0)
(n0>1.0) :- clockwise(2.0) P14(2.0,n0+1.0) cntr-
clockwise(1.0) cntr-clockwise(1.0) up(1.0)
(n1>0.0) :- [clockwise(2.0) P14(2.0,n0+1.0) ]
```

P9:

```
(n1>5.0) :- up(5.0)
```


Construction Sequence

```
offset-increase(3) left(4) right(1) offset-decrease(5) cntr-
clockwise(1) cntr-clockwise(1) right(3) right(3) cntr-
clockwise(5) [ up(5) right(1) offset-decrease(5) [ up(5) [
offset-decrease(3) forward(2) ] [ clockwise(1) ] ] cntr-
clockwise(1) cntr-clockwise(1) right(3) right(3) left(1) [
back(1) right(1) ] [ offset-decrease(1) clockwise(2) cntr-
clockwise(1) cntr-clockwise(1) up(1) ] left(4) offset-
decrease(5) [ up(5) [ offset-decrease(3) forward(2) ] [
clockwise(1) ] ] cntr-clockwise(1) cntr-clockwise(1)
right(3) right(3) left(1) [ back(1) right(1) ] [ offset-
decrease(1) clockwise(2) cntr-clockwise(1) cntr-
clockwise(1) up(1) ] left(4) offset-decrease(5) ] up(1) [ [
offset-decrease(1) ] clockwise(4) offset-decrease(1)
left(1) ] left(1) revolute2(4) revolute2(4) right(1) up(1) [
up(1) down(5) down(1) ] [ up(1) down(4) offset-
decrease(1) ] right(1) up(4) up(4) left(1) [ back(1)
right(1) ] right(3) right(3) cntr-clockwise(4) [ up(5)
right(1) offset-decrease(5) [ up(5) [ offset-decrease(3)
forward(2) ] [ clockwise(1) ] ] cntr-clockwise(1) cntr-
clockwise(1) left(1) [ back(1) right(1) ] [ offset-
decrease(1) ] left(4) offset-decrease(5) [ up(5) [ offset-
decrease(3) forward(2) ] [ clockwise(1) ] ] cntr-
clockwise(1) cntr-clockwise(1) left(1) [ back(1) right(1)
] [ offset-decrease(1) ] left(4) offset-decrease(5) ] up(1)
[ [ offset-decrease(1) ] clockwise(4) offset-decrease(1)
left(1) ] left(1) revolute2(4) revolute2(4) right(1) up(1) [
up(1) down(5) down(1) ] [ up(1) down(4) offset-
decrease(1) ] right(1) up(4) up(4) left(1) left(1) [ offset-
decrease(1) clockwise(2) [ clockwise(1) ] clockwise(1)
cntr-clockwise(1) cntr-clockwise(1) up(1) ] left(1)
offset-decrease(5) cntr-clockwise(1) cntr-clockwise(1)
right(3) right(3) cntr-clockwise(5) [ up(5) right(1)
offset-decrease(5) [ up(5) [ offset-decrease(3)
forward(2) ] [ clockwise(1) ] ] cntr-clockwise(1) cntr-
clockwise(1) right(3) right(3) left(1) [ back(1) right(1) ]
[ offset-decrease(1) clockwise(2) cntr-clockwise(1) cntr-
clockwise(1) up(1) ] left(4) offset-decrease(5) [ up(5) [
offset-decrease(3) forward(2) ] [ clockwise(1) ] ] cntr-
```


GENERATIVE REPRESENTATION for ROBOT “Kayak”
(Continued)

<p>(n1>1.0) :- P5(n0-1.0,n0/5.0) P7(n1/n0,5.0/4.0) P7(n1/n0,n1/1.0) (n0>0.0) :- P5(n0-1.0,n0/5.0)</p> <p>P10: (n1>4.0) :- [offset-decrease(3.0) forward(2.0)] [clockwise(1.0)] (n1>3.0) :- right(1.0) {offset-decrease(5.0) P11(3.0,n1+n0) P6(n0-5.0,n0-4.0) }(2.0) (n1>0.0) :- right(1.0) {offset-decrease(5.0) P6(n0-5.0,n0-4.0) P11(4.0,n1+n0) }(2.0) offset-decrease(5.0)</p> <p>P11: (n0>1.0) :- cntr-clockwise(1.0) P11(n0/4.0,n0+n1) (n0>0.0) :- cntr-clockwise(1.0) P9(n1/4.0,n0-1.0) P12(n1/5.0,5.0) (n0>0.0) :- cntr-clockwise(1.0) P12(n1/5.0,5.0)</p> <p>P12: (n1>1.0) :- left(1.0) [back(1.0) right(1.0)] P9(n1-4.0,3.0) [offset-decrease(1.0) P8(n1,n1+n0)] left(n0) (n1>-1.0) :- left(1.0) [back(1.0) down(n0)] P9(n1-4.0,3.0) left(n1) (n1>0.0) :- [back(1.0) down(n0)] [P8(n0,n0-n1)]</p> <p>P13: (n0>4.0) :- left(1.0) P12(4.0,1.0-2.0) (n0>3.0) :- [[offset-decrease(1.0)] clockwise(4.0) offset- decrease(1.0) left(1.0)] left(1.0) P12(4.0,1.0-2.0) (n1>0.0) :- left(1.0) clockwise(4.0)</p> <p>P14: (n0>5.0) :- (n0>5.0) :- clockwise(1.0) P5(n1+n0,n1/5.0) (n0>0.0) :- [clockwise(1.0)] clockwise(1.0)</p>	<p>clockwise(1) cntr-clockwise(1) right(3) right(3) left(1) [back(1) right(1)] [offset-decrease(1) clockwise(2) cntr-clockwise(1) cntr-clockwise(1) up(1)] left(4) offset-decrease(5)] up(1) [[offset-decrease(1)] clockwise(4) offset-decrease(1) left(1)] left(1) revolute2(4) revolute2(4) right(1) up(1) [up(1) down(5) down(1)] [up(1) down(4) offset-decrease(1)] right(1) up(4) up(4) left(1) [back(1) right(1)] right(3) right(3) cntr-clockwise(4) [up(5) right(1) offset-decrease(5) [up(5) [offset-decrease(3) forward(2)] [clockwise(1)]] cntr-clockwise(1) cntr-clockwise(1) left(1) [back(1) right(1)] [offset-decrease(1)] left(4) offset-decrease(5) [up(5) [offset-decrease(3) forward(2)] [clockwise(1)]] cntr-clockwise(1) cntr-clockwise(1) left(1) [back(1) right(1)] [[offset-decrease(1)] left(4) offset-decrease(5)] up(1) [[offset-decrease(1)] clockwise(4) offset-decrease(1) left(1)] left(1) revolute2(4) revolute2(4) right(1) up(1) [up(1) down(5) down(1)] [up(1) down(4) offset-decrease(1)] right(1) up(4) up(4) left(1) left(1) [offset-decrease(1) clockwise(2) [clockwise(1)] clockwise(1) cntr-clockwise(1) cntr-clockwise(1) up(1)] left(1) left(4) [down(1) cntr-clockwise(2) down(1) cntr-clockwise(2) left(1) [back(1) down(2)] left(1) left(1) left(1)]</p> <p><i>See Table II for meaning of commands.</i></p>
--	--

GENERATIVE REPRESENTATION for ROBOT “Quatrobot”

Generative Representation	Construction Sequence
<p>This generative representation is started with P0(5.0,4.0) and run thirteen iterations. (T) indicates that the condition always succeeds. The assembly procedure is not included because it consists of almost 900 commands.</p> <p>P0: (n1>1.0) :- [P12(n0-0.0,n1-0.0)] [P3(4.0,n1/1.0) forward(1.0) P0(n1/2.0,5.0) decrease-offset(n0)] (n1>0.0) :- [decrease-offset(1.0) cntr-clockwise(1.0) left(1.0) forward(1.0)] increase-offset(1.0)</p> <p>P1: (n1>1.0) :- increase-offset(1.0) P8(n1,n0+5.0) [P7(n1- n0,n1+n0) right(5.0) P10(n1/n0,n0-4.0) increase-offset(1.0)] left(1.0) forward(1.0) clockwise(1.0) (n0>0.0) :- decrease-offset(1.0)</p> <p>P3: (n1>2.0) :- [forward(1.0) revolute(n1) down(1.0) forward(1.0) P6(n0+2.0,3.0) increase-offset(1.0) forward(1.0)] forward(1.0) right(1.0) increase-offset(1.0) (n1>0.0) :- [forward(1.0) revolute(n1) forward(1.0) down(1.0) P6(n0+2.0,3.0) increase-offset(1.0) forward(1.0)] forward(1.0) right(1.0) increase-offset(1.0)</p>	 <p><i>(Construction sequence has not been included do to its length)</i></p>

<p>P4: $(n1 > 6.0) :-$ [forward(n0) down(1.0)] forward(n0) [down(1.0)] [left(1.0) forward(1.0)] [P7(n1-n0,n1/5.0) cntr-clockwise(3.0)] revolute(1.0) back(5.0) P1(2.0,n1/2.0) $(n0 > 0.0) :-$ [forward(n0) down(1.0)] [forward(n0) down(1.0)] [left(1.0) forward(1.0)] [P7(n1-n0,n1/5.0) cntr-clockwise(3.0)] revolute(1.0) back(5.0) P1(2.0,n1/2.0)</p> <p>P6: $(n0 > 2.0) :-$ right(1.0) down(1.0) [forward(1.0) forward(1.0)] [right(5.0) cntr-clockwise(2.0)] increase-offset(n1) P4(n0- 2.0,n1+3.0) back(n1) left(1.0) $(n0 > 0.0) :-$ [right(1.0) revolute2(1.0) revolute2(1.0) left(1.0)] [cntr-clockwise(n1) left(n1) up(n0) left(4.0)]</p> <p>P7: $(n1 > 2.0) :-$ clockwise(n0) clockwise(n0) $(n1 > 0.0) :-$ [decrease-offset(1.0) forward(1.0) increase- offset(1.0)] clockwise(n0) clockwise(n0)</p> <p>P8: $(n1 > 5.0) :-$ forward(5.0) left(1.0) cntr-clockwise(1.0) clockwise(1.0) decrease-offset(1.0) forward(5.0) cntr- clockwise(1.0) [clockwise(3.0) revolute2(1.0) left(1.0)] $(n1 > 0.0) :-$ back(1.0)</p>	<p>P10: $(n1 > -2.0) :-$ forward(n1) P14(n1-5.0,n0/3.0) [increase-offset(3.0) increase-offset(1.0) right(1.0)] [up(1.0)] {right(1.0) }(2.0) $(n0 > 0.0) :-$ P14(n1-5.0,n0/3.0) forward(n1) [increase-offset(3.0) increase-offset(1.0) up(1.0)] [up(1.0)] {right(1.0) }(2.0)</p> <p>P12: $(n1 > 3.0) :-$ clockwise(2.0) up(5.0) up(4.0) $(n1 > 0.0) :-$ clockwise(2.0) forward(5.0) up(4.0)</p> <p>P14: $(n1 > 2.0) :-$ [P0(n0+1.0,n0/3.0) down(5.0) increase-offset(1.0) clockwise(3.0)] down(n0) down(4.0) back(1.0) down(1.0) down(1.0) decrease-offset(1.0) $(n0 > 0.0) :-$ up(1.0) right(1.0) forward(4.0) P14(2.0,n1-2.0) P8(3.0,n0+2.0) forward(1.0) clockwise(1.0) decrease- offset(1.0) back(1.0) forward(1.0) increase-offset(1.0)</p>
--	--

ACKNOWLEDGMENT

The authors thank J. Rieffel for construction help with the Quatrobot.

REFERENCES

- [1] T. Bäck, F. Hoffmeister, and H.-P. Schwefel, "A survey of evolution strategies," in *Proc. 4th Int. Conf. Genetic Algorithms*, R. K. Belew and L. B. Booker, Eds., 1991, pp. 2–9.
- [2] T. Back, *Evolutionary Algorithms in Theory and Practice*. New York: Oxford Univ. Press, 1996.
- [3] P. J. Bentley, "Generic evolutionary design of solid objects using a genetic algorithm," Ph.D. dissertation, Univ. of Huddersfield, Huddersfield, U. K., 1996.
- [4] —, *Evolutionary Design by Computers*. San Mateo, CA: Morgan Kaufman, 1999.
- [5] P. Bentley and S. Kumar, "Three ways to grow designs: a comparison of embryogenies of an evolutionary design problem," in *Proc. Genetic and Evolutionary Computation Conf.*, 1999, pp. 35–43.
- [6] I. Chen and J. Burdick, "Determining task optimal robot assembly configurations," in *Proc. IEEE Int. Conf. Robotics and Automation*, 1995, pp. 132–137.
- [7] O. Chocron and P. Bidaud, "Genetic design of 3-D modular manipulators," in *Proc. IEEE Int. Conf. Robotics and Automation*, 1997, pp. 223–228.
- [8] P. Coates, T. Broughton, and H. Jackson, "Exploring three-dimensional design worlds using Lindenmayer systems and genetic programming," in *Evolutionary Design by Computers*, P. J. Bentley, Ed. San Mateo, CA: Morgan Kaufmann, 1999.
- [9] S. Farritor, S. Dubowsky, N. Rutman, and J. Cole, "A systems-level modular design approach to field robotics," in *Proc. IEEE Int. Conf. Robotics and Automation*, 1996, pp. 2890–2895.
- [10] S. Farritor and S. Dubowsky, "On modular design of field robotic systems," *Auton. Robots*, vol. 10, no. 1, pp. 57–65, 2001.
- [11] P. Funes and J. B. Pollack, "Computer evolution of buildable objects," in *Proc. 4th Eur. Conf. Artificial Life*, P. Husbands and I. Harvey, Eds. Cambridge, MA, 1997, pp. 358–367.
- [12] F. Gruau, "Neural network synthesis using cellular encoding and the genetic algorithm," Ph.D. dissertation, Ecole Normale Supérieure de Lyon, Lyon, France, 1994.
- [13] R. Grzeszczuk and D. Terzopoulos, "Automated learning of muscle-actuated locomotion through control abstraction," in *Proc. SIGGRAPH '95 Conf.*, Los Angeles, CA, 1995, pp. 63–70.
- [14] J. H. Holland, *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: Univ. Michigan Press, 1975.
- [15] G. S. Hornby and J. B. Pollack, "The advantages of generative grammatical encodings for physical design," in *Proc. Congress on Evolutionary Computation*, 2001, pp. 600–607.
- [16] G. S. Hornby and J. B. Pollack, "Body-brain coevolution using L-systems as a generative encoding," in *Proc. Genetic and Evolutionary Computation Conf.*, San Francisco, CA, 2001, pp. 868–875.
- [17] G. S. Hornby and J. B. Pollack, "Evolving L-systems to generate virtual creatures," *Computers and Graphics*, vol. 25, no. 6, pp. 1041–1048, 2000.
- [18] G. S. Hornby, M. Fujita, S. Takamura, and T. Yamamoto, "Autonomous evolution of gaits with the Sony quadruped robot," in *Proc. Genetic and Evolutionary Computation Conf.*, W. Banzhaf, J. Daida, A. E. Eiben, M. Garzon, V. Honavar, M. Jakiela, and R. E. Smith, Eds., 1999.
- [19] G. S. Hornby, S. Takamura, O. Hanagata, and M. Fujita, "Evolution of controllers from a high-level simulator to a high-DOF robot," in *Evolvable Systems: From Biology to Hardware*, J. Miller, Ed. New York: Springer, 2000.
- [20] G. S. Hornby and J. B. Pollack, "Creating high-level components with a generative representation for body-brain evolution," in *Artificial Life*, to be published.
- [21] P. Husbands, G. Gerny, M. McIlhagga, and R. Ives, "Two applications of genetic algorithms to component design," in *Evolutionary Computing*, T. Fogarty, Ed. New York: Springer-Verlag, 1996, pp. 50–61.
- [22] C. Jacob, "Genetic L-system programming," in *Parallel Problem Solving from Nature, Lecture Notes in Computer Science*, Y. Davidor and P. Schwefel, Eds. New York: Springer-Verlag, 1994, vol. 866, pp. 334–343.
- [23] N. Jakobi, "Minimal simulations for evolutionary robotics," Ph.D. dissertation, School of Cognitive and Comput. Sci., Univ. of Sussex, Sussex, U. K., 1998.
- [24] C. Kane and M. Schoenauer, "Genetic operators for two-dimensional shape optimization," in *Artificiale Evolution—EA95*, J.-M. Alliot, E. Lutton, E. Ronald, M. Schoenauer, and D. Snyers, Eds. New York: Springer-Verlag, 1995.
- [25] J.-O. Kim and P. K. Khosla, "Design of space shuttle tile servicing robot: an application of task based kinematic design," in *Proc. IEEE Int. Conf. Robotics and Automation*, 1993, pp. 867–874.
- [26] M. Komosinski and A. Rotaru-Varga, "From directed to open-ended evolution in a complex simulation model," in *Artificial Life*. San Mateo, CA: Morgan Kaufmann, 2000, vol. 7, pp. 293–299.
- [27] J. R. Koza, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press, 1992.
- [28] P. C. Leger, "Automated synthesis and optimization of robot configurations: An evolutionary approach," Ph.D. dissertation, Carnegie Mellon Univ., Pittsburgh, PA, 1999.
- [29] A. Lindenmayer, "Mathematical models for cellular interaction in development, parts I and II," *J. Theoretical Biol.*, vol. 18, pp. 280–299, 1968.
- [30] A. Lindenmayer, "Adding continuous components to L-systems," in *L-Systems, Lecture Notes in Computer Science*, G. Rozenberg and A. Salomaa, Eds. New York: Springer-Verlag, 1974, vol. 15, pp. 53–68.
- [31] H. Lipson and J. B. Pollack, "Automatic design and manufacture of robotic lifeforms," *Nature*, vol. 406, pp. 974–978, 2000.
- [32] M. Mataric and D. Cliff, "Challenges in evolving controllers for physical robots," *Robot. Auton. Syst.*, vol. 19, no. 1, pp. 67–83, 1996.

- [33] J. Maynard-Smith and E. Szathmary, *The Major Transitions in Evolution*. Oxford, U. K.: Oxford Univ. Press, 1995.
- [34] Z. Michalewicz, *Genetic Algorithms + Data Structures = Evolution Programs*. New York: Springer-Verlag, 1992.
- [35] J. T. Ngo and J. Marks, "Space-time constraints revisited," in *Proc. SIGGRAPH '93 Conf.*, 1993, pp. 343–350.
- [36] S. Nolfi and D. Floreano, *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press, 2000.
- [37] G. Ochoa, "On genetic algorithms and Lindenmayer systems," in *Parallel Problem Solving from Nature*, A. Eiben, T. Baeck, M. Schoenauer, and H. P. Schwefel, Eds. New York: Springer-Verlag, 1998, vol. V, pp. 335–344.
- [38] C. Paredis, "An agent-based approach to the design of rapidly deployable fault tolerant manipulators," Ph.D. dissertation, Dept. Elect. Comput. Eng., Carnegie Mellon Univ., Pittsburgh, PA, 1996.
- [39] J. B. Pollack, H. Lipson, S. Ficici, P. Funes, G. Hornby, and R. Watson, "Evolutionary techniques in physical robotics," in *Evolvable Systems: From Biology to Hardware*, J. Miller, Ed. New York: Springer-Verlag, 2000.
- [40] P. Prusinkiewicz and A. Lindenmayer, *The Algorithmic Beauty of Plants*. New York: Springer-Verlag, 1990.
- [41] G. P. Roston, "A genetic methodology for configuration design," Ph.D. dissertation, Dept. Mech. Eng., Carnegie Mellon Univ., Pittsburgh, PA, 1994.
- [42] M. Schoenauer, "Shape representations and evolution schemes," in *Evolutionary Programming*, L. J. Fogel, P. J. Angeline, and T. Back, Eds. Cambridge, MA: MIT Press, 1996, vol. 5.
- [43] H. Simon, *The Sciences of the Artificial*. Cambridge, MA: MIT Press, 1969.
- [44] K. Sims, "Evolving virtual creatures," in *Proc. SIGGRAPH '94 Conf.*, 1994, pp. 15–22.
- [45] R. A. Watson and J. B. Pollack, "A Computational Model of Symbiotic Composition in Evolutionary Transitions," *Biosystems*, to be published.
- [46] J. B. Pollack, H. Lipson, S. Ficici, P. Funes, G. Hornby, and R. Watson, "Evolutionary techniques in physical robotics," in *Lecture Notes In Computer Science*. New York: Springer, 2000, vol. 1801.
- [47] K. T. Ulrich and K. Tung, "Fundamentals of product modularity," in *Issues in Design/Manufacture Integration*, A. Sharon, Ed. New York: ASME, 1991, pp. 73–79.
- [48] C. C. Huang and A. Kusiak, "Modularity in design of products and systems," *IEEE Trans. Syst., Man, Cybern. A*, vol. 28, pp. 66–77, Jan. 1998.
- [49] R. N. Langlois, "Modularity in technology and organization," *J. Econ. Behavior, Org.*, to be published.
- [50] L. H. Hartwell, J. J. Hopfield, S. Leibler, and A. W. Murray, "From molecular to modular cell biology," *Nature*, vol. 402, pp. C47–C52, 1999.



Gregory S. Hornby received the Ph.D. degree in computer science from Brandeis University, Waltham, MA, in 2002.

He is a Computer Scientist with QSS Group Inc., working in the Computational Sciences Division, NASA Ames Research Center, Moffett Field, CA. Previously, he was a Visiting Researcher at Sony's Digital Creatures Laboratory, where he evolved the dynamic gait used on the consumer version of AIBO. His current work consists of using evolutionary algorithms to assist NASA missions,

such as designing better antennas for satellites.



Hod Lipson is an Assistant Professor with the Department of Mechanical and Aerospace Engineering and Computing and Information Science, Cornell University, Ithaca, NY. Prior to this appointment, he was a Postdoctoral Researcher at Brandeis University, Waltham, MA, working on evolutionary computation and evolutionary robotics. He was also a Lecturer at the Massachusetts Institute of Technology, Cambridge, MA, where he taught design and conducted research in design automation. His current research interests are in the area of automatic synthesis, both physical and computational.



Jordan B. Pollack received the Ph.D. degree in computer science from the University of Illinois, Urbana, in 1987.

He has worked on artificial intelligence since 1975. He is now an Associate Professor with Brandeis University, Waltham, MA. He directs the Dynamical and Evolutionary Machine Organization, known as the DEMO Laboratory. At DEMO, he and his colleagues have applied evolutionary learning to significant problems in game playing, problem solving, search, language induction, robotics, and even educational learning across the Internet.