Adaptation and Learning in Animated Creatures

Jiming Liu **Hong Qin**

Y. Y. Tang **Y. T. Wu**

Dept. of Computing Studies, Baptist University
224 Waterloo Road, Kowloon Tong, Hong Kong
Phone: 852-2339-7088 Fax: 852-2339-7892
EMail: *jiming@comp.hkbu.edu.hk*

Abstract

This paper is concerned with synthetic agents interacting with virtual environments, called *animated creatures*. The animated creatures are articulated graphical figures that are equipped with a set of primitive behavioral patterns. These patterns qualitatively specify which body modules will move concurrently, hence forming a motion group, and which group will move prior to another. The parameterization of these patterns is carried out by the creatures given certain external stimuli. The key to such behavioral adaptation lies in an embedded evolution strategy based selection mechanism. Two examples will be given where this selection mechanism enables a bipedal creature and a six-legged creature to dynamically search for the exact positions as well as duration of body joints as constrained by the qualitatively defined gait patterns. The acquired new stimulus-response pairs are recorded and inserted into a behavioral conditioning network which can be reused and refined during future movements.

Introduction

Modern computer graphics technology has enjoyed rapid development in recent years, and hence has attracted researchers and practitioners to explore a wide spectrum of applications, ranging from computer-aided graphical designs to artificial life and virtual-reality (Maes 1995). This paper is concerned with the animation-based entertainment use of computer graphics, and in particular, it describes our present research work on creating articulated creatures, i.e., animated robots that can adapt to their virtual environments, and learn new behaviors to attain some specific goals.

Autonomous Agents 97, Marina Del Rey, California USA

These goals may be given in terms of high-level behavioral commands composed of expressions that indicate the situated actions of the creatures in response to both global task constraints and local conditions, such as walking through a terrain without collision.

Many researchers have investigated the problem of articulated creature (or figure) generation. Among them, Ching and Badler (Ching & Badler 1992) have studied the issue of generating realistic motions by treating it as a robot global motion planning problem in which a configuration-space collision-free path planning algorithm can be applied. Vasilonikolidakis and Clapworthy (Vasilonikolidakis & Clapworthy 1991) have used an inverse Lagrangian dynamics algorithm to compute inertial motions based on carefully studied gait determinants. Loizidou and Clapworthy (Loizidou & Clapworthy 1993) have explored the use of dynamic analysis in the manipulation and control of articulated figures by using a hybrid direct-inverse dynamics method. Their work applied the findings from physiological and photographical studies to decide gait determinants, as well as those from legged robot research to account for ground reaction forces. Considering the issue of computational complexity, Arai (Arai 1993) has proposed to use partial dynamics only for some specific parts of the body such as legs and arms. Green and Halliday (Green & Halliday 1996) have developed a system in which both geometrical and behavioral descriptions of an object are allowed.

As may be noted, all the above mentioned studies have, to some extent, shared one thing in common; namely, the realistic motion is achieved by solving either complete or partial kinematic and dynamic equations. Two questions that remain are (1) how a realistic movement can be most *efficiently* generated, and (2) how the articulated figure will select an *appropriate* behavior in response to not only the given goal but also some *unpredictable* conditions in its environment. This issue is particularly relevant if we are to develop synthetic agents that can "survive", autonomously, and acquire previously undefined behaviors in their virtual environments. It would be impossible to write complete animation code for each single agent that we create.

In this paper, we address this issue by building and investigating animated creatures with embedded behavioral adaptation and emergence mechanisms. In order to limit our scope, here we shall not focus on the precise dynamics of the creature once a behavior has been either emerged or selected.

Partially related to our work are some of the previous studies on behavior selection and emergence. For example, Maes (Maes 1991) has developed a selection mechanism that emerges an action by spreading activation energy over a behavior network. In our current implementation, the reflexes of a creature are linked among each other in the form of stimulus-response pairs, somewhat similar to Maes' predecessor and successor links. In relation to search based behavior selection, Auslander et al (Auslander *et al.* 1995) have developed a system that contains banked stimulus-response controllers dynamically selected through an optimization algorithm. However, unlike our qualitatively bounded search space, their search space is composed of all the possible motion controllers. As reported, such a search can be a very slow process especially when 3D animation is concerned. Furthermore, Ventrella (Ventrella 1994) has studied the possibility of emerging the structure and locomotion behaviors of an animate using genetic algorithms (Goldberg 1989) (Holland 1975), his system used a model of specifically tailored qualitative forward dynamics to generate gravitational, inertial, momentum, frictional, and dampening effects. Sims (Sims 1994) has developed a system in which both the animated 3D creaturess bodies (i.e., morphology) and their neural control systems (i.e., virtual "brains) were genetically evolved and/or co-evolved. His system can produce realistic dynamics simulations of gravity, collisions, and friction.

The Animated Creatures

The animated creatures to be presented in this paper can be characterized as follows: They are goal-attaining creatures capable of dynamically adjusting their behaviors in reaction to their environment. Two kinds of information will be made available for the creatures at all time; namely, (1) a set of isolated primitive behavioral patterns in terms how various body modules can move concurrently and what kind of temporal ordering they may have, and (2) the aim of their movement (or the commanded goal) and the sensory feedback about their surrounding environments. An evolution strategy based optimum seeking mechanism (Schwefel 1995) will be embedded inside the creatures that governs the behavioral adaptation and learning. The objectives are, by way of evolution, the creatures can find the precise body coordination strategies pertaining to a selected specific behavioral pattern, and furthermore, gradually emerge a more complex behavior through the chaining of basic stimulus-response pairs.

Organization of the Paper

This paper is organized as follows: Section 2 provides an overview of the creatures' evolution strategy as used in their behavioral adaptation and learning. Section 3 provides two creature examples as produced by our implemented system. Section 4 provides some further details on the implemented system. Finally, section 5 concludes the paper and points out some issues for future investigation.

Evolution Strategy Based Behavioral Adaptation in Creatures

Given a high-level behavioral command, an animated creature will first select a primitive behavioral pattern from a user-defined library[1]. This pattern provides information about the synchronized motions about different modules of the creature, but not the exact motion parameters such as initial positions, velocities, and final positions. In order to take into account both the immediate environment constraints as sensed by the creature using its virtual proximity sensors, and the aim/goal of its movement as implied in the high-level command, the creature undertakes a moment of behavioral *adaptation* based on an embedded evolution strategy, i.e., parameterization of the chosen behavioral pattern. The adaptation step will result in a series of internal configuration changes in the creature. This series is regarded as the optimal response by the creature, as conditioned by the original stimulus. As soon as it finishs executing such a conditioned response, the creature may again face some new local constraints and/or updated goals. Hence, the above process repeats itself.

It should be pointed out that all the conditioned pairs of stimulus-response as found during the adaptation will be recorded (i.e., *learned*) by the creature into its behavioral conditioning network, which would allow for future instant reflexes and/or refinement, should same stimuli are revisited.

Figure 1 presents a schematic diagram that outlines the creature's mechanism for behavioral adaptation and learning. In what follows, we shall provide some details on the implementation of this mechanism.

Primitive Behavioral Patterns

A primitive behavior has no goals and objectives. It is simply a primitive coordinated motion template.

Definition 1. *A primitive behavioral pattern is defined as a tuple:* $< C, S_c >$*, where* $C = \{C\}$ *and configuration* $C =< \theta_1, \theta_2, \ldots, \theta_p >$ *denotes a set of grouped body modules that are considered to undertake synchronized motions.* $S_c = [C_1, C_2, \ldots, C_q]$ *denote a sequence of temporally ordered configurations.*

[1]The selection of a primitive behavioral template could also be made adaptive.

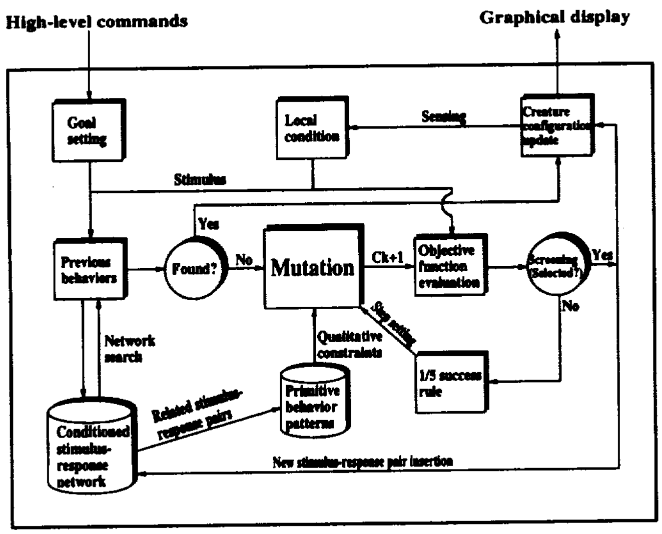

Figure 1: An animated goal-attaining creature adapts to its virtual environment based on *primitive behavioral patterns* and *evolution strategies*.

It should be mentioned that a primitive behavioral description is qualitative in nature which only gives the valid ranges of the joint angles rather than exact values.

Evolution Strategy

To search a particular configuration given a certain stimulus in the environment, the animated creature will undergo an evolution strategy based search process to find a set of values that would best suit the requirements. Here, the stimulus is termed in a broader sense in that it encompasses not only goal specification, such as "move straight and turn left once an obstacle is encountered", but at the same time also the local surrounding information, such as sensory data based perception.

The evolution strategy based optimum search technique applied in this work was inspired by the earlier work of Schwefel (Schwefel 1995). More specifically, the evolution strategy can be expressed as follows:

$$\begin{array}{lcl} C_1^{k+1} & = & C_1^k + M_1^k(0,\sigma) \\ C_2^{k+1} & = & C_2^k + M_2^k(0,\sigma) \\ \vdots & \vdots & \vdots \\ C_q^{k+1} & = & C_q^k + M_q^k(0,\sigma) \end{array} \tag{1}$$

where C^{k+1} and C^k are the creature's current configuration (offspring) and previous configuration (parent), respectively. M^k denotes a mutation term, which can be interpreted as the sum of many individual events and hence is implemented in our system to satisfy a Gaussian probability distribution with zero mean and σ standard deviation (Schwefel 1995):

$$p(M^k) = p(m_1, m_2, \cdots, m_p) = \prod_{i=1}^{p} p(m_i) \tag{2}$$

$$p(m_i) = \frac{1}{\sqrt{2\pi}\sigma_i} exp\left(-\frac{(m_i - \xi_i)^2}{2\sigma_i^2}\right) \tag{3}$$

where ξ_i and σ_i correspond to the mean and standard deviation of the distribution, respectively.

The evolution strategy works as follows: Given a certain external stimulus, the creature forms a cost function F to evaluate whether its current configuration meets the requirements as induced from the stimulus. If not, the creature's mutation mechanism generates an offspring configuration which can replace the parent (i.e., old configuration) *if and only if* F satisfies the following:

$$F(C_1^{k+1}, C_2^{k+1}, \ldots, C_q^{k+1}) < F(C_1^k, C_2^k, \ldots, C_q^k) \tag{4}$$

If the above is not satisfied, the offspring configuration will not be selected. While responding to a particular stimulus, a series of configurations C may be generated and parameterized based on the triggered behavioral pattern.

Unlike genetic algorithms, the evolution strategies operate on floating point vectors, directly, with a dynamically changing reproduction rate. And furthermore, the relative order of selection and recombination procedures is different between evolution strategies and genetic algorithms. Michalewica (Michalewicz 1992) has presented a more detailed comparison between evolution strategies and genetic algorithms.

Control of Mutation Steps

Back et al (Back, Hoffmeister, & Schwefel 1991) have proven that with the above-mentioned evolution strategy, a global optimum can always be found with probability one for sufficiently long search time, that is, $P\{\lim_{k\to\infty} F(C^k) = F_{opt}\} = 1$. Furthermore, Rechenberg (Rechenberg 1973) has proposed a so-called "$\frac{1}{5}$ success rule", which states that from time to time during the evolution strategy based optimum seeking, the ratio φ of successful mutations to all mutations should be $\frac{1}{5}$. If the ratio is greater than $\frac{1}{5}$, increase the variance, if it is less than $\frac{1}{5}$, decrease the variance.

In our present implementation, the animated creature adjusts its mutation steps using a more precise formulation of the "$\frac{1}{5}$ success rule", which was originally provided by Schwefel (Schwefel 1995). This rule can be stated as follows: For every n mutations, the creature checks how many successes have occurred over the preceding $10n$ mutations. If the number of successes is less than $2n$, it multiplies the step sizes by a factor of 0.85; divides them by 0.85 if more than $2n$ successes have occurred. This mechanism of changing mutation steps can reasonably maintain the found local optima and hence prevent the search from becoming completely random[2].

Behavioral Chaining

Once a complete response, as expressed in terms of a series of configuration changes, has been selected, a stimulus-response behavioral pair will be inserted into a behavioral conditioning network where each arc connects a pair of stimulus and response, i.e., $S_1 \leadsto R_1$. At the end of one response, if a new stimulus is present (e.g., when the previous stimulus-response pair has resulted in either inefficient or unstable configurations with respect to the new local environment), the creature will continue to select a corresponding response either from the conditioning network or through the evolution strategy based search. In such a case, *behavioral chaining* would become possible. This may be best illustrated in the following diagram:

$$\begin{array}{ccccc} S_1 & \leadsto & R_1 & & \\ & & \Downarrow & & \\ & & S_2 & \leadsto & R_2 \\ & & & & \Downarrow \\ & & & & \cdots \quad \cdots \end{array} \tag{5}$$

where "$\Downarrow$" implies that a specific response has led to a new stimulus.

Such a behavioral chaining technique allows for the emergence of *complex behaviors* on the basis of the primitive simpler ones. A similar idea of using local primitive behaviors has been proposed before in the context of learning high-level global control strategies for collective robot tasks (Mataric 1992).

Examples of Animated Adaptive Creatures

This section presents two examples of animated creatures incorporating the above-mentioned behavioral adaptation mechanism. In both examples, the creatures are provided with a series of primitive behavioral patterns. The stimuli are recognized by the creatures at run-time, taking into account (1) the aim of the movement, (2) the current sensory inputs about their immediate geometric environments, and (3) the balance constraints that bound the relative movements among the legs.

Bipedal Movements

Our first example is concerned with an animated creature in a virtual staircase environment, as shown in Figure 2. The primitive template available for this creature is *bipedal gait*. The creature is commanded to **walk straight** without any further instruction on how to move up and down the staircases, and is required to figure out itself the appropriate configurations for different bipedal movements while trying to keep its entire body in good balance.

Figure 2: An animated bipedal creature commanded to walk straight forward in a virtual staircase environment.

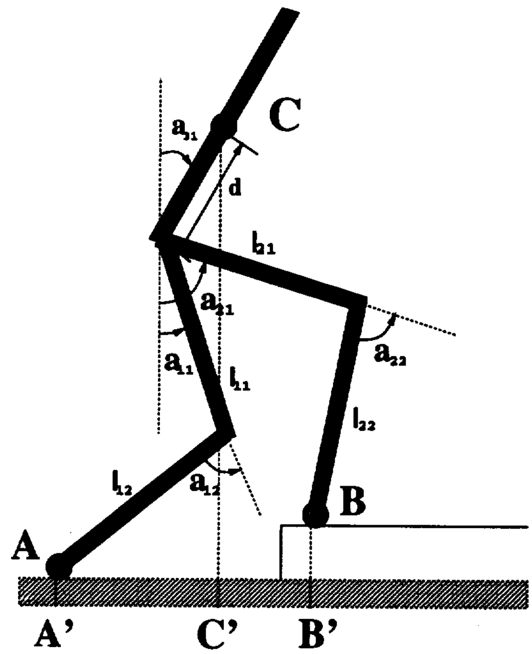

Figure 3: An illustration of the coupled joints in a bipedal creature whose exact values are to be searched with respect to both the balance constraints and the movement plan.

Figure 3 shows the articulated joints of the bipedal creature. One example of the balance constraints in this case would be to maintain point C' as the centroid of $A'B'$ as closely as possible by selecting the coupled leg joint angles, a_{11}, a_{12}, a_{21}, and a_{22}, and the body orientation angle, a_{31}, within their respective ranges. This constraint can be written as follows:

$$\mid \frac{1}{2}A'B' - A'C' \mid \leq \epsilon \tag{6}$$

where ϵ denotes a small positive error bound. $A'B'$ and $A'C'$ are determined by the joint angles as in the

[2] Another mechanism for achieving this balanced mutation has been proposed by Harvey (Harvey 1992), which is called Species Adaptation Genetic Algorithms (SAGA).

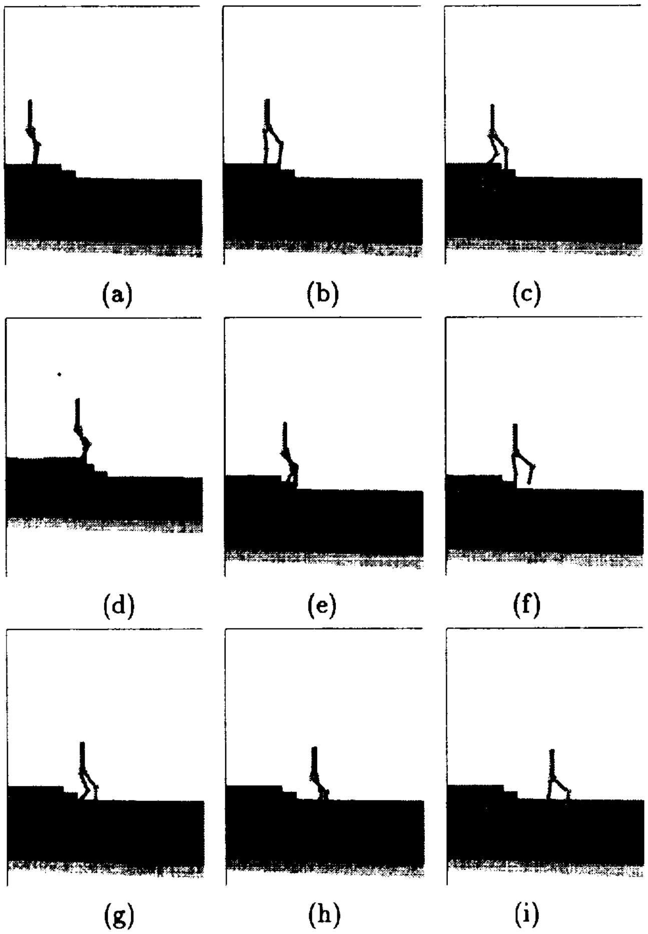

Figure 4: The animated creature performs a sequence of bipedal movements. Its behaviors at the local environments as well as their transitions, such as *walking_on_flat_surface* $\Rightarrow$ *stepping_downwards* $\Rightarrow$ *walking_on_flat_surface*, must be autonomously evolved.

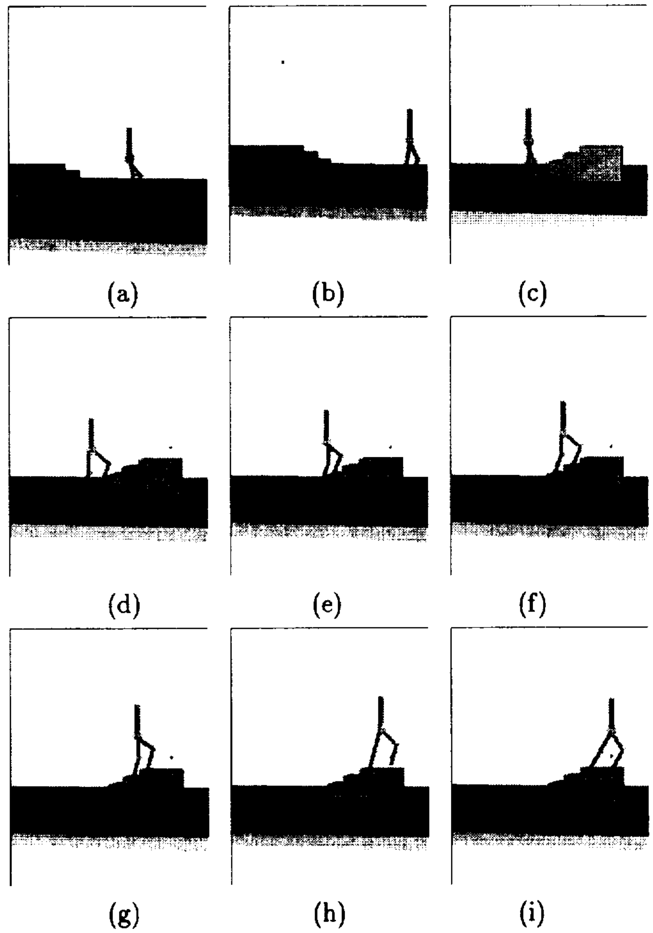

Figure 5: Continued bipedal movements for transitions *walking_on_flat_surface* $\Rightarrow$ *stepping_upwards* $\Rightarrow$ *walking_on_flat_surface*.

following expressions, respectively:

$$\begin{aligned} A'B' &= -[l_{11}\sin(a_{11}) + l_{12}\sin(a_{11}+a_{12})] \\ &\quad +[l_{21}\sin(a_{21}) + l_{22}\sin(a_{21}+a_{22})] \end{aligned} \tag{7}$$

and

$$\begin{aligned} A'C' &= -[l_{11}\sin(a_{11}) + l_{12}\sin(a_{11}+a_{12})] \\ &\quad +d\sin(a_{31}) \end{aligned} \tag{8}$$

Figures 4 and 5 provide a sequence of snapshots showing how the creature adapts to the local environment during its forward movements. As can be noted, several configuration changes are required at each of the following transitions: *walking_on_flat_surface* $\Rightarrow$ *stepping_downwards* $\Rightarrow$ *walking_on_flat_surface* $\Rightarrow$ *stepping_upwards* $\Rightarrow$ *walking_on_flat_surface*.

What is interesting to observe from Figures 5(b), (h), and (i) is that when the creature encounters the flat surface for the second and third times, it immediately executes its already conditioned reflexes, as searched from its behavioral conditioning network.

Six-Legged Movements

Figure 6 shows a six-legged creature being placed in a virtual environment that consists of a door frame and a sloped surface. The creature is commanded to *move straight ahead and turn right once the environment boundary is encountered.* Such a movement plan is specified by users through a window interface.

During its movements, the creature is equipped with several artificial proximity sensors to detect the obstacles and terrain conditions as encountered. These surrounding conditions will become part of the behavioral adaptation triggers. Generally speaking, this creature must (1) correctly orient its main body and (2) find the precise coordinated joint movements, in order to optimize the following concurrent motion objectives:

- movement toward a virtual goal straight ahead, and
- front collision-free movement.

The primitive behavioral pattern used by the creature in this case is *tripod gait.*

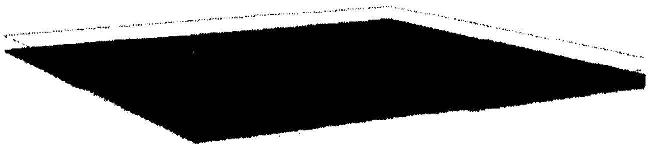

Figure 6: A six-legged creature commanded to move forward. In the course of its movements, the time duration and angular joint positions for all legs must be dynamically adjusted in reaction to the local environments.

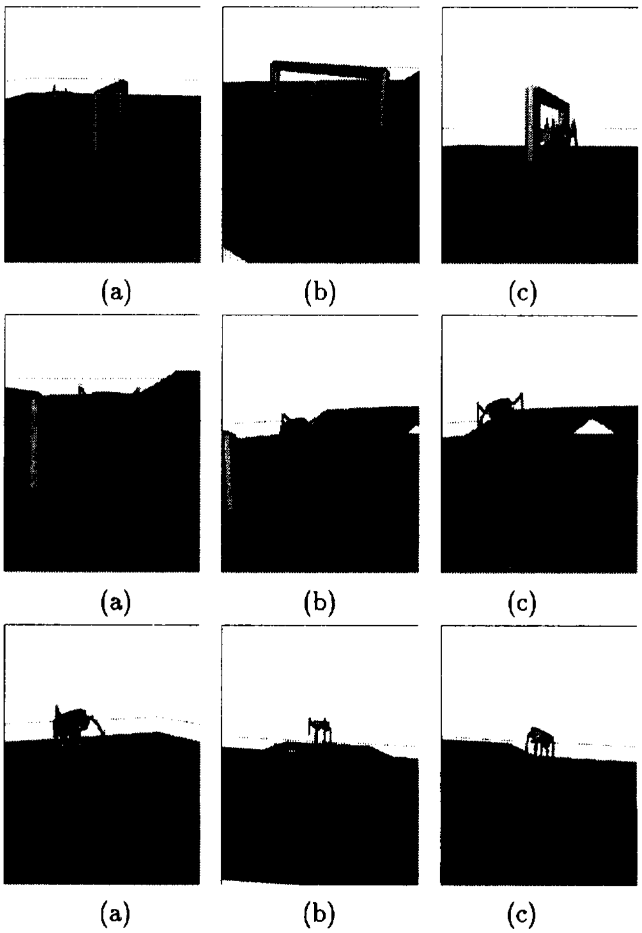

Figure 7: The snapshots of the creature in approaching the door frame, passing through the door frame, climbing up the slope, walking on the raised floor, and walking down the slope.

Figure 7 shows a series of snapshots on the creature's movements in which both the initial positions and the durations of the leg joints are changed in reaction to different surrounding conditions. These changes are dynamically generated (conditioned) as a result of the evolution strategy based adaptation that searches, within the selected qualitative behavioral template, for all the hip joints and the lower limb joints, the total of which can be as many as $6 \times 2(\textit{hip angles}) + 6 \times 1(\textit{lower limb angles})$.

System Details

As far as the definition of primitive behavioral patterns is concerned, links pertaining to a particular structural branch can be selected and grouped through a graphical user interface. At the same time, their motion range and the degree of motion overlap can be conjectured and specified. Further to the group definition, individual groups can also be sequenced to form a primitive behavioral pattern with or without triggering conditions.

In the present implementation, an artificial creature can be commanded through two means: (1) aim of movements by input devices such as a mouse, and (2) high-level movement commands. Some of the recognizable commands are concerned with simple movement behaviors, such as **`move_straight`**, **`move_over`**, **`move_under`**, **`turn_left`**, and **`turn_right`**, while others are more complex motion behaviors, such as **`avoid`**, **`follow`**, and **`block`**.

Concluding Remarks

In this paper, we have described our recent work on building synthetic autonomous agents (or animated creatures) that can dynamically select near-optimal behaviors and incrementally construct conditioned stimulus-response chains. In order to minimize the computation involved, the embedded mechanism for behavioral adaptation and learning utilizes a set of predefined primitive behavioral patterns (e.g., gaits) dictating qualitatively how various parts alternate in the course of movements. The actual adaptation of an agent's behavior is implemented using an evolution strategy based optimum seeking technique.

As has been demonstrated in our experiments, the evolution strategy based selection is well suited to the addressed problem of behavioral selection and learning under different external stimuli. This is largely due to the fact that such a technique directly accepts and evaluates numerical optimality constraints. Our experimental results have also shown that the developed adaptation mechanism is robust and efficient in finding the right motion parameters at run-time. The acquired stimulus-response pairs, or conditioned reflexes, can readily be retained by the creature and gracefully integrated into the existing pairs to develop complex behaviors.

One of the obvious limitations in our present implementation is the lack of dynamical constraints, such as the effects of surface friction and inertia. It is our present research goal to incorporate such constrains into the objective function of evolution strategy based parameterization.

References

Arai, K. 1993. Keyframe animation of articulated figures using partial dynamics. In Thalmann, N. M., and Thalmann, D., eds., *Models and Techniques in Computer Animation.* Tokyo: Springer-Verlag. 243–256.

Auslander, J.; Fukunaga, A.; Partovi, H.; Christensen, J.; Hsu, L.; Reiss, P.; Shuman, A.; Marks, J.; and Ngo, J. T. 1995. Further experience with controller-based automatic motion synthesis for articulated figures. *ACM Transactions on Graphics* 14(4):311–336.

Back, T.; Hoffmeister, F.; and Schwefel, H.-P. 1991. A survey of evolution strategies. In Belew, R., and Booker, L., eds., *Proceedings of the Fourth International Conference on Genetic Algorithms.* Los Altos, CA: Morgan Kaufmann Publishers.

Ching, W., and Badler, N. 1992. Fast motion planning for anthropometric figures with many degrees of freedom. In *Proceedings of the 1992 IEEE International Conference on Robotics and Automation, Nice, France*, 2340–2345.

Goldberg, D. E. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning.* Reading, MA: Adison-Wesley publishing company.

Green, M., and Halliday, S. 1996. A geometric modeling and animation system for virtual reality. *Communications of the ACM* 39(5):46–53.

Harvey, I. 1992. Species adaptation genetic algorithms: The basis for a continuing SAGA. In Varela, F. J., and Bourgine, P., eds., *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life.* Cambridge, MA: The MIT Press/Bradford Books. 346–354.

Holland, J. H. 1975. *Adaptation in Natural and Artificial Systems.* Ann Arbor: The University of Michigan Press.

Loizidou, S., and Clapworthy, G. J. 1993. Legged locomotion using HIDDS. In Thalmann, N. M., and Thalmann, D., eds., *Models and Techniques in Computer Animation.* Tokyo: Springer-Verlag. 257–269.

Maes, P. 1991. A bottom-up mechanism for behavior selection in an artificial creature. In Meyer, J. A., and Wilson, S., eds., *Proceedings of the first International Conference on Simulation of Adaptive Behavior.* Cambridge, MA: The MIT Press.

Maes, P. 1995. Artificial life meets entertainment: Lifelike autonomous agents. *Communications of the ACM* 38(11):108–114.

Mataric, M. J. 1992. Minimizing complexity in controlling a mobile robot population. In *Proceedings of the 1992 IEEE International Conference on Robotics and Automation.*

Michalewicz, Z. 1992. *Genetic Algorithms + Data Structures = Evolution Programs.* Berlin: Springer-Verlag.

Rechenberg, I. 1973. *Evolutionsstrategie – Optimierung technischer Systeme nach Prinzipien der biologischen Evolution.* Stuttgart: Frommann-Holzboog.

Schwefel, H.-P. 1995. *Evolution and Optimum Seeking.* New York: John Wiley & Sons, Inc.

Sims, K. 1994. Evolving 3D morphology and behavior by competition. In Brooks, R. A., and Maes, P., eds., *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems.* Cambridge, MA: The MIT Press. 28–39.

Vasilonikolidakis, N., and Clapworthy, G. J. 1991. Design of realistic gaits for the purpose of animation. In Thalmann, N. M., and Thalmann, D., eds., *Computer Animation '91.* Tokyo: Springer-Verlag. 101–114.

Ventrella, J. 1994. Explorations in the emergence of morphology and locomotion behavior in animated characters. In Brooks, R. A., and Maes, P., eds., *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems.* Cambridge, MA: The MIT Press. 436–441.