

Interactive Training for Synthetic Characters

Song-Yee Yoon^{†,‡}, Robert C. Burke[†], Bruce M. Blumberg[†] and Gerald E. Schneider[‡]

sy@media.mit.edu

rob@media.mit.edu

bruce@media.mit.edu

jerry@mit.edu

[†] Synthetic Characters Group
Media Laboratory

Massachusetts Institute of Technology
Cambridge, MA 02139

[‡] Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, MA 02139

Abstract

Compelling synthetic characters must behave in ways that reflect their past experience and thus allow for individual personalization. We therefore need a method that allows characters to learn. But simply adding traditional machine learning algorithms without considering the characters' own motivations and desires will break the illusion of life. Intentional characters require interactive learning. In this paper, we present the results of Sydney K9.0, a project based on the Synthetic Characters creature kernel framework. Inspired by pet training, we have implemented a character that can be trained using the "clicker training" technique. Clicker training utilizes the natural desires of an animal and employs operant conditioning procedures for shaping their behavior. The necessary plasticity of system interconnections shaped by associations and rewards that is required by clicker training was integrated into the creature kernel framework. The implemented system also includes a module named DogEar that is designed for collecting real-world acoustic data, such as human voice commands, integrated into the creature kernel's perception system. This provides a seamless interface between the simulated and real worlds. Detailed implementation and interaction results are presented.

Introduction

Synthetic characters are artificial creatures that have their own motivations and desires and can interact with human beings in real time (Kline and Blumberg 1999). The capacity to adapt and the ability to convey intentions (Dennett 1987) have been pointed out as key factors that maintain the illusion of life during human interactions with characters designed to have their own mind and will.

In particular, a character's ability to learn makes the interactions between humans and characters a sustainable relationship. For example, toys that learn through interactions are appealing as creatures that can reflect their past experiences. They are personalizable, and are seen to exist as unique beings.

Introducing learning abilities into synthetic characters brings with it several challenges if we hope to maintain a lifelike impression. For example, we want characters to learn what they are being taught, but remain in character. Yet, through this training, we would like to be able to shape characters, and personalize them as our friends and companions. Even if they are learning, they should have their own drives and affects, and show these states through their actions and emotional expressions. This raises the issue of how to train a being that has its own intentions. More fundamentally, learning is a hard problem for any system but one which can be greatly facilitated by the presence of a skilled trainer who can guide the learning process and thereby reduce the complexity of the learner's task.



Figure 1. This figure shows two of the characters that inhabit the Sydney K9.0 virtual world - Sydney and Fridge. Human participants who interact with this system can train the virtual dog, Sydney. Fridge is the graphical instantiation of the user, who holds and controls the training stick to guide Sydney's attention. Fridge can reward Sydney with milkbones when the human participant shakes the milkbone box.

Pet trainers face a similar problem, since real animals also have their own drives and interests. Among the various training techniques that have been tried with pets, the "clicker training" method, derived from operant conditioning, has proven to be particularly successful with animals ranging from dogs to dolphins (Wilkes 1995;

Pryor 1999). Inspired by this idea from pet training, our project implements a clicker training session.

Related Work

Building an artifact that can learn is an extensively studied area (Ballard 1997). In the case of situated characters, however, additional constraints – which naturally originate due to the complexity of the surrounding world – need to be dealt with. Terzopoulos and Tu (1994) integrated learning into graphical creatures where the learning focused on locomotion for surviving in the simulated physical world. Relevant cues were already given and creatures were assumed to know what to learn and pay attention to. In real world situations, however, deciding what to attend to and learning relevant cues and the right associative links are significant challenges (Scherl and Levesque 1993). These are areas where extensive state-based search methods (Allen, Hendler and Tate 1990) have been used for dealing with the complexity of the environment and as a strategy for learning, and this has been applied to the domain of self-animated graphical creatures (Funge, Tu and Terzopoulos 1999).

This is exactly the problem that must be addressed by animals in the natural world. Studies of animal behavior (Lorenz 1981; Skinner 1938) have defined operant conditioning as a powerful learning scheme for a situated creature (Sutton and Barto 1998). Operant conditioning has been implemented in a number of studies with promising results (Drescher 1993; Touretzky and Saksida 1996; Blumberg, Todd and Maes 1996). However, the question addressed by our work is different. Namely, assuming you have a system that can learn using operant and classical associations, what is the best way to teach the system? To answer this question, we take our inspiration from animal training.

Clicker Training

One of the most popular methods that pet trainers use is clicker training (Wilkes 1995), which is based on operant conditioning and consisting of two associative learning processes. By repeatedly giving a reward (i.e. something like a food treat that is motivationally important to the dog) right after sounding a hand-held “clicker” device, the dog forms an associative link between the clicker sound and the positive reinforcer. After multiple repetitions, the clicker becomes a secondary reinforcer that reliably signals that a primary reinforcer will be forthcoming. In addition, the clicker acts as an “event marker” indicating the exact behavior and/or configuration for which the dog is to be rewarded. A behavior that is consistently rewarded will tend to increase in frequency per Thorndike’s Law (Wilkes 1995). A clicker trainer uses this phenomenon together with techniques such as luring or shaping to create the desired behavior. In luring, the dog is lured into a configuration that can then be rewarded. In shaping,

the trainer guides the dog’s search by rewarding successive approximations to the desired behavior. Once the behavior is in place, a clicker trainer then teaches the dog the context in which this behavior is a reliable strategy for being rewarded. Typically, the context is a gesture or an utterance. In a sense, since the dog already “knows” the behavior, it can focus on learning the context in which the behavior is maximally reliable. Backward chaining is typically used to teach sequences of actions.

There are two properties of the clicker sound that make it effective for training. Firstly, the clicker emits a sound of a very distinct frequency that is easily discernable from other auditory inputs. Secondly, it provides a precise “end of behavior” marker because it has very short temporal characteristics. Thus, while verbal praises such as *good dog* could be used for the same purpose, the duration of such an utterance is sufficiently long that a number of behaviors could have been performed while the utterance was being made (Wilkes 1995). The clicker sound provides the dog with a more accurate sense of which behavior is being rewarded.

Clicker training works effectively for pets because it solves several problems faced by a learner in a complex and unpredictable environment. The interactive nature of the training and the distinctive clicker marking help the learner to solve the temporal and behavioral credit assignment problems (Sutton and Barto 1998), which are crucial for successful learning in a dynamic environment. The training proceeds by forming a chain of links that ultimately leads to a motivationally important state for the animal. The training starts with forming the simple associative link between the clicker sound and reward, then the rest of the procedure results in backwardly expanding the link to a more complex behavioral repertoire, including only elements that ultimately lead to reward. Feedback from the trainer, such as the clicker sound, helps the pet to determine exactly why it is being rewarded and what it is supposed to learn and focus on. At the same time, this simplifies the job of the trainer.

Sydney K9.0

In this section, we describe a project called **Sydney K9.0**, which implements three main characters: “Sydney,” an interactive dog; “Fridge,” his friend the walking refrigerator; and “Nabee,” the butterfly. We extend the creature kernel framework (Yoon 2000) to include an interactive learning ability in building the characters for this project. In addition, Sydney is endowed with an auditory system described below.

An example interaction is shown in Figure 2. The participant has all four elements of the interface available: a wearable microphone, a training stick, a milkbone box and the clicker. The user can move the training stick in any direction and the movement and resulting location of the stick is reflected by a corresponding change in the training stick that Fridge holds. Sydney can sense the

location of the tip of the training stick using a virtual sensor located on the virtual training stick in Fridge's hand. The clicker sound is treated as a special sound signal that is perceived as distinct from the speech signal, even though both are auditory signals. When Sydney's behavior deserves to be rewarded, the human participant shakes the milkbone box, which results in a bone appearing on screen, which Sydney can eat to satisfy his hunger drive. A sensor is installed inside the milkbone box that is activated whenever the box is shaken.



Figure 2. Sydney K9.0 implements a virtual dog training session where Fridge, a refrigerator in the virtual world, is both the trainer and the proxy for the human participant. The participant can provide input to the system with the aid of various interfaces: a wearable microphone, a training stick, a milkbone box and a clicker.

The participant can train Sydney using the clicker training method described above. Since this training proceeds while respecting the synthetic character's motivations and intentions – just as when we train real dogs – we believe it proposes a method for training an intentional character without compromising the character's illusion of life (Thomas and Johnston 1981). A graphical display of the interaction between Sydney and Fridge is shown in Figure 1.

Now we discuss how the characters in **Sydney K9.0** were built, and then describe the external components, such as speech processing for DogEar; and internal components, such as Sydney's behavior, learning and memory implementation.

Characters

All three characters in this project – Sydney the dog, Fridge the refrigerator, and Nabee the butterfly – are implemented within the creature kernel framework as suggested by Blumberg (Blumberg 1996; Blumberg and Galyean 1995) and later expanded by Kline (Kline and Blumberg 1999) and Yoon (Yoon 2000).

Creature kernel. Creature kernel is a framework for building interactive synthetic characters. Within this

framework, a creature is built internally as a sum of four subcomponents: the motivation, behavior, motor and perception systems. These four components coordinate to make the character operate as a functioning whole.

The motivation, behavior and perception systems are organized as a network of basis units. For example, the behavior system is modeled as a hierarchically connected network of behavior units as shown in Tinbergen's animal behavior model (Tinbergen 1951). The perception system is a system of sensors that extract information about what is going on inside and outside of the character. It includes both the virtual vision sensors for seeing another character in the same virtual world, and physical sensors for receiving input (such as the training stick) coming from the physical world. The motivation system is composed of affect and drives. Affect refers to the creature's emotional state, and is associated with the character's feelings and moods. Drives are motivators, and include endogenous drives and externally induced desires. The behavior system performs action selection i.e., it chooses the most relevant behavior to perform given the perception and motivation input. The behavior system in turn sends a signal to the motor system, which is a set of action primitives and available motor skills, indicating which action should be performed and how it should be performed. In our system, where characters are 3D animated creatures, the motor system is composed of a set of animations, and it renders the skill chosen by the behavior system on the screen.

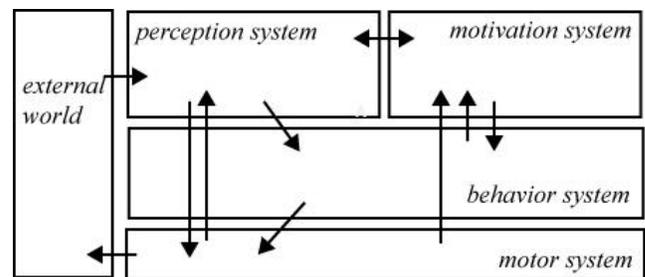


Figure 3. A schematic diagram of a synthetic character. Arrows represent the information flow between systems. Well-coordinated communication among these four systems—perception, motivation, behavior and motor—is required for a character to successfully function in a dynamic world.

Figure 3 shows the schematic diagram of the creature kernel. Arrows represent information flow between components. Any loop made by a succession of arrows represents a possible sequence, e.g., from the creature's perception of the world, to internal information processing and then to exertion on the outside world. All three characters in this project are implemented using this framework. Detailed explanation of this kernel operation is found in Yoon (2000).

Creature Kernel Implementation

This section further explains the creature kernel framework as it pertains to the implementation of Sydney. Sydney is based on a version of the creature kernel that has been augmented to allow for the processing of real-time voice commands from the human participant, and to allow formation of associations between verbal commands and corresponding tricks during the clicker training session.

The Motivation System. Sydney's motivation system consists of two parts – the drive system and the affect system. Sydney has three main drives: hunger, curiosity and fatigue. He feels hungry if he does not consume an adequate amount of food for a long time, and thus Sydney perceives as an incentive or positive reinforcement something that will satisfy his hunger drive (i.e., a snack). This drive provides Sydney with the motivation to learn through interaction with the human participant because learning and performance of a desired behavior provide Sydney with milkbones as rewards.

Curiosity drives Sydney to explore the world. Drives rise at different rates while satisfying stimuli are not provided, or while the key stimuli that instigate them are perceived. The Curiosity drive rises when Sydney stays at one place for a long period of time, and exploring the world lowers this curiosity. Sydney is also interested in interacting with other creatures such as Nabee, the butterfly. So, for example, when Nabee is close to Sydney, his curiosity rises until it often surpasses the other drives. Fatigue brings Sydney back home for a rest, and as he rests his Fatigue drive is reduced. After the fatigue level diminishes below the level of another drive, Sydney gets up to initiate behaviors that will satisfy his currently most pressing drive.

The Behavior System. Sydney's behavior system is composed of two main parts. The first is the *primary* behavior system, which is a set of behavior nodes and sequences that can be used to satisfy goals. In essence, it is composed of behavior units such as walk, run, lie down, sit down, roll over, beg, etc., which form sequences within the behavior network to satisfy goals and drives.

The other part of Sydney's behavior system is the *autonomous* behavior system, which contains more reflex-like behavior such as ear movement – triggered when Sydney detects sound – and autonomous stand, which is a low-priority default behavior that runs to avoid situations where no animation would be running because no particular behavior is chosen by the action selection mechanism.

The Motor System. As for Sydney's motor system, animation files made by animators are provided as the means for enabling Sydney to perform certain skills and display them on the screen. Leaf node behavior units, at

the termination of the behavior system's branches, are where the behavior system blends into the motor system. The motor system uses a verb-adverb model (Rose, Cohen and Bodenheimer 1998), whereby the behavior system specifies a verb/adverb and the motor system does real time motion interpolation among animation examples to produce motion with the appropriate emotional content. The animation blending capabilities of the motor system are described in more detail by Johnson (1999).

The Perception System. Sydney's perception system is expanded to cope with real-time voice inputs from human participants. His perception system is divided into two parts: virtual sensors, and sensors responsible for receiving inputs from the physical world. Virtual sensors enable Sydney to access all of the returnable features of the targeted objects (Kline and Blumberg 1999). The perception network is extended by additional sensors that return levels of beliefs about whether the objects are far, close, etc.

The sensors responsible for processing real world inputs, include *DogEar*, an interface to the human voice commands. The system is built to use two PCs. One PC is devoted to receiving data through the microphone and running *DogEar*, and the other PC updates the creature kernel and renders graphics. Implementation detail on *DogEar* is found in the next section.

DogEar

The *DogEar* system has been implemented to mediate verbal communication between a human participant and Sydney. It receives the human participant's utterance data, and converts the raw data to a Cepstral coefficient format that is received and stored by Sydney's auditory memory.

Data Acquisition. The sounds in the system are sampled at 11025 Hz. Sound bites are obtained by using a thresholding algorithm that averages the signal over windows of 512 samples. Recording starts when the signal is above the threshold, and ends when it has been below the threshold for three successive windows. The sample is then trimmed at each end to the nearest zero-crossing.

Representation. We have chosen to use a vector of cepstral coefficients as a representation. Like a real dog, Sydney does not comprehend language, nor does he have a concept of language. What Sydney interprets is the acoustic pattern of the speech signal; thus, a Cepstral coefficient representation is a sufficient and appropriate encoding of the necessary information. Cepstral analysis is a technique that removes the pitch ripple from high-resolution speech spectra, as examined by Rabiner and Juang (1993). The goal of cepstral analysis is to obtain the vocal tract response after removing the pitch ripple.

The *DogEar* system performs this task by filtering the log-magnitude of the signal with an inverse FFT. This is followed by truncation of the coefficients beyond the pitch

frequency, and then a forward FFT (Intel 1998). In the Fourier domain, we are using 10 filters placed linearly on a scale from 100 Hz up to 2 kHz, followed by 10 additional filters laid out on a Mel scale up to 6400 Hz. Analysis is performed using a window size of 512 samples and an overlap of 256 samples per window. At present, in order to facilitate classification of feature vectors of equal length, cepstral coefficients are obtained for exactly the first 20 windows. If the utterance is sufficiently short that it does not fill 20 windows, the remaining windows are filled with silence. We thus end up with 400-dimensional cepstral coefficient feature vectors.

Integration

The protocol that allows DogEar to communicate with the creature kernel system includes two binary bits indicating the existence of new speech and clicker inputs, and the Cepstral coefficient representation of new utterance data.

Classifier and Short-Term Memory. The main system includes a short-term memory module with a fixed number of memory cells. These cells, referred to as memory groups below, each represent uniquely classified utterances. This memory module is very simple in implementation; its main role is in storing patterns. When a new utterance is heard, it is compared to the utterances in all the groups to see if the distance of the newly arrived data in any of the groups is closer than a preset threshold. Dynamic Time Warping (Rabiner 1993, Intel 1998) is used as a distance metric.

Interaction. From Sydney's standpoint, the information he receives which is relevant to training comes from auditory inputs, utterances of the human participant and the clicker sound, and virtual visual inputs. It is assumed that Sydney has the ability to distinguish the clicker sound from utterances, due to its distinctive frequency.

Without the human participant's presence, Sydney freely wanders around, motivated by his internal drives and affects aroused endogenously as well as perceptually. At realization of the human's presence, through a speech signal that is perceived as an attentional bid (from Fridge), Sydney runs toward Fridge and awaits a command because he has learned that the interaction and performing the right behaviors lead to reward. As the user's speech command is issued, Sydney performs a trick that, he believes, matches the spoken command. If he was right, the human participant shakes the milkbone box in his hand that produces a milkbone in the virtual world. Sydney runs toward the milkbone, consumes it and thereby reduces his level of hunger. A reward solidifies his belief in the link between the command and the behavior he just performed. When the behavior was not right, the lack of reward weakens his belief in the associative link between the perceived acoustic pattern and the behavior. This makes the behavior less likely to be performed the next time he hears a similar utterance. In the absence of a

command Sydney will probabilistically choose a behavior based on the likelihood of reward.

Learning

The "eating" behavior is rewarding because it reduces the level of hunger. Thus, perceiving a milkbone is rewarding for Sydney. The unconditioned stimulus, food, and the unconditioned response, eating that leads to hunger drive satisfaction, is an innate behavior context chain. This pathway is represented as the chain of links marked **A** in Figure 4. The first thing that a dog trainer does is to form a linkage between the clicker sound and something rewarding. It could be any type of food that the dog enjoys or petting or praise that causes positive affective arousal in the dog. A strong link between the clicker sound and the unconditioned reward stimulus makes the dog regard the clicker sound itself as rewarding, and thus the trainer can use this clicker sound as a positive reinforcer for further training. This associative link between two perceptual stimuli is marked **B** in Figure 4.

Finally, operant conditioning builds a chain of links labeled **C** in Figure 4. A verbal command is issued and, when Sydney performs the appropriate behavior, he is rewarded with the clicker sound. It forms an associative link between the command and the behavior that leads to the clicker sound. The sound has a positive reinforcing value due to the expectation of a reward to follow. Update rules and learning algorithms for integrating learning into the creature kernel are described by Yoon, Blumberg and Schneider (2000).

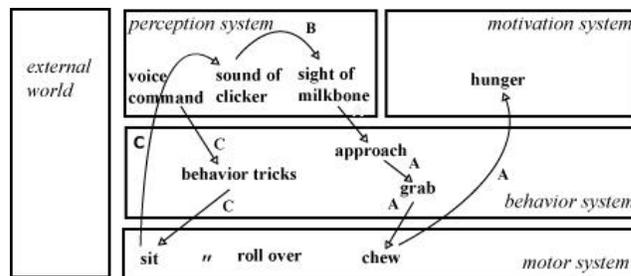


Figure 4. Clicker training, effectively, builds up a behavioral representation within Sydney's creature kernel. The chain labeled **A** is an unconditioned stimuli-unconditioned response (US-UR) link that preexists in the creature kernel. During the first phase of the clicker training, an associative link between two stimuli, clicker sound and food, a piece of milkbone in this project's case, is formed and this is labeled **B** in this figure. The third associative link chain is formed through an operant conditioning process and this is labeled **C**. The link encompasses the sequence of spoken command → performance of a corresponding trick → clicker sound. The link between a spoken command and a trick is strengthened as the outperforming of the link is rewarded by the clicker sound, which has a positive reinforcing value as it is followed by an unconditioned rewarding stimulus.

Observations and Discussion

Preliminary interactive experiences have been provided to experienced programmers, as well as novice users who have never interacted with software creatures before.

Participants who had pet-training experience felt comfortable extending their experience to the virtual pet training session. One response we received to a recalcitrant Sydney while adjusting the noise controls in a crowded environment was, "he's acting that way because you aren't treating him gently enough. I know how dogs behave since I have my own pet dog at home." The user went on to (gently) teach Sydney how to roll over on command.

Since the acoustic data processing does not assume *a priori* knowledge of language or grammar, Sydney's acoustic system has proven robust. At one point during a demonstration, Sydney had learned to respond appropriately to commands in English, French, Japanese and Norwegian.

Perhaps most importantly, users were consistently impressed by Sydney's lifelike behavior. These observations demonstrate the plausibility of using the clicker training method to allow synthetic characters to learn while maintaining the illusion of life that comes from their actions and reactions.

Limitations of this implementation include the lack of ability to process affect information in the verbal commands. As our participants indicated, real dogs discern contextual information from cues such as prosody and pitch in utterances. This extension will be implemented by augmenting DogEar to extract these acoustic cues and transfer them to the behavior system.

More fundamentally, we need to implement luring and shaping. Clicker training requires variation on the part of the animal to produce behavior that can be rewarded. The trainer can wait for the creature to produce the behavior, but it is often more efficient to lure the animal into performing it. For example, to teach the dog to lie down, a trainer can lure the dog to lie down and then click so that it learns it is a behavior that can potentially bring a reward.

In Nature an adaptable mammal must have a variety of learning types available to it for optimizing its chances of survival (Lorenz 1981). We believe our system is general enough to enable a number of these learning types in addition to operant conditioning, including habituation, concept formation, and formation of affective responses to objects or places (see Yoon et al. 2000). A full evaluation of the learning abilities of a synthetic character like Sydney will require detailed analysis of all these learning types, a task we are now at the stage of undertaking.

Acknowledgments. Many thanks to all the members of the Synthetic Characters Group, and Gary Wilkes for introducing us to the world of clicker training.

References

Allen, J., Hendler, H. and Tete, A, eds. 1990. *Readings in Planning*, San Mateo, CA.: Morgan Kaufmann.
Ballard, D. H. 1997. *An Introduction to Natural Computation*. Cambridge, MA.: MIT Press.

Blumberg, B. and Galyean, T. 1995. Multi-Level Direction of Autonomous Creatures for Real Time Virtual Environment. In *Proceedings of SIGGRAPH 95*: New York, NY. ACM SIGGRAPH

Blumberg, B. M. 1996. Old Tricks, New Dogs: Ethology and Interactive Creatures. Ph.D. diss., The Media Laboratory, Massachusetts Institute of Technology.

Blumberg, B. M., Todd, P. M. and Maes, P. 1996. No Bad Dogs: Ethological Lessons for Learning in Hamsterdam. From Animals to Animats 4, *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. Cambridge, MA.: MIT Press.

Dennett, D. 1987, *The Intentional Stance*, Cambridge, MA.: MIT Press.

Drescher, G. 1991. *Made-up minds: a constructive approach to artificial intelligence*. Cambridge, MA.: MIT Press.

Funge, J., Tu, X. and Terzopoulos, D. 1999. Cognitive Modeling: Knowledge, Reasoning and Planning for Intelligent Characters. In *Proceedings of SIGGRAPH 99*, 29-38.: New York, NY. ACM SIGGRAPH

Intel, Inc. 1998. Intel® Recognition Primitives Library V4.0 Documentation. Santa Clara, CA.: Intel Corporation. Available at <http://developer.intel.com/vtune/perflibt/RPL>

Johnson, M. P. 1999. Multi-Dimensional Quaternion Interpolation, In *ACM SIGGRAPH99 Conference Abstracts and Applications*, page 258. New York, NY. ACM SIGGRAPH.

Kline, C. and Blumberg, B. 1999. The Art and Science of Synthetic Character Design. In Proceedings of the AISB 1999 Symposium on AI and Creativity in Entertainment and Visual Art, Edinburgh, Scotland.

Lorenz, K. 1981. *The Foundations of Ethology*. New York, NY.: Springer-Verlag.

Pryor, K. 1999. *Clicker Training for Dogs*. Waltham, MA.: Sunshine Books, Inc.

Rabiner, L. and Juang, B. 1993. *Fundamentals of Speech Recognition*. New York, NY.: Prentice-Hall.

Rose, C., Cohen, M. F., and Bodenheimer, B. 1998. Verbs and Adverbs: Multidimensional Motion Interpolation. *IEEE CG&A*, 18(5): 32-40. Piscataway, NJ.: IEEE Press.

Scherl, R. and Levesque, H. 1993. The frame problem and knowledge-producing actions. In *Proceedings of AAAI-93*. Menlo Park, Calif.: AAAI Press.

Skinner, B. F. 1938. *The Behavior of Organisms*. New York, NY.: Appleton Century Crofs.

Sutton, R. and Barto, A. 1998. *Reinforcement Learning*. Cambridge, MA.: MIT Press.

Terzopoulos, D. and Tu, X. 1994. Artificial fishes: Autonomous locomotion, perception, behavior and learning in a simulated physical world. *Artificial Life* 1:327-351.

Thomas, F. and Johnston, O. 1981. *The Illusion of Life: Disney Animations*. New York, NY.: Abbeville Press.

Tinbergen, N. 1951. *The Study of Instinct*. New York, NY.: Oxford University Press.

Touretzky, D. S. and Saksida, L. M. 1996. Skinnerbots. From Animals to Animats 4, *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. Cambridge, MA.: MIT Press.

Wilkes, G. 1995. Click and Treat Training Kit Version 1.2. Mesa, AZ.: www.clickandtreat.com.

Yoon, S. -Y. 2000. Affective Synthetic Characters. Ph D. diss. Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology. *Forthcoming*.

Yoon, S.-Y., Blumberg, B. M. and Schneider, G. E. 2000. Motivation Driven Learning for Interactive Synthetic Characters. In *Proceedings of the Fourth International Conference on Autonomous Agents*. New York, NY.: ACM Press. *Forthcoming*.