Social Learning in Humans, Animals and Agents

Daphna Buchsbaum[1], Bruce Blumberg[2] and Cynthia Breazeal[2]

[1]Icosystem Corporation
10 Fawcett St. Cambridge MA 02138
[2]MIT Media Lab
20 Ames St. Cambridge MA 02139
daphna@icosystem.com
bruce@media.mit.edu
cynthiab@media.mit.edu

Abstract

We want to build animated characters and robots capable of rich social interactions with humans and each other, and who are able to learn by observing those around them. An increasing amount of evidence suggests that, in animals and humans, the ability to learn by watching others, and in particular, the ability to imitate, could be crucial precursors to the development of appropriate social behavior, and ultimately the ability to reason about the thoughts, intents, beliefs, and desires of others.

We have created a number of imitative characters and robots, the latest of which is Max T. Mouse, an anthropomorphic animated mouse character who is able to observe the actions he sees his friend Morris Mouse performing, and compare them to the actions he knows how to perform himself. Max's imitation and social learning system allows him to identify simple goals and motivations for Morris's behavior, an important step towards developing characters with a full theory of mind. In this paper, we describe the cognitive basis for Max's social learning capabilities, and explore the implications of our work for future research in both artificial and natural systems.

Introduction

Some of the most exciting new applications being developed for artificial creatures require them to cooperate with humans and each other as socially capable partners. As character designers, it is possible to gain valuable insights into how social intelligence might operate and be acquired by looking to the fields of developmental psychology and animal behavior. It appears that, among animals, learning from the behavior of others (known as social learning) is by no means a single monolithic process. Rather, species sample widely from a spectrum of overlapping social competencies (Heyes and Galef 1996, Whiten 2004), ranging from using information about others to help focus their attention, to emulating other's actions and goal states. While very few species exhibit the most complex forms of imitation, and perhaps no non-human animal possesses a full theory of mind (Premack and Woodruff 1978), the abilities animals do possess allow them to consistently exploit their social environment in ways that far outstrip our current technologies.

Furthermore, many of the simpler behavior-reading abilities present in animals may represent prerequisites for the more complex mind-reading abilities humans possess. An increasing amount of evidence suggests that, in human infants, the ability to learn by watching others, and in particular, the ability to imitate, could be crucial precursors to the development of appropriate social behavior, and ultimately the ability to reason about the behaviors, emotions, beliefs, and intents of others (Meltzoff 1996, Meltzoff and Moore 1997, Meltzoff and Decety 2003).

Approach

In previous work, we began to explore the role of imitation and social learning in artificial intelligence, by implementing a facial imitation architecture for an interactive humanoid robot (Breazeal et al. 2004). In this paper, we present a novel system that provides artificial creatures with a cognitive architecture inspired by the literature on animal social learning, including a robust mechanism for observing and imitating whole gestures and movements. Critically, the characters presented in this paper are able to use their imitative abilities to bootstrap simple mechanisms for identifying each other's low-level goals and motivations and learning from each other's actions, bringing us several steps closer to the goal of creating socially intelligent artificial creatures.

Social Learning in Humans and Animals

There is a rich research literature available investigating social learning in a variety of animal species, particularly non-human primates. Much of this literature has been devoted to partitioning socially mediated learning into various subtypes (for a review see Whiten 2004, Call and Carpenter 2002, Byrne and Russon 1996 and also Heyes and Galef 1996). The primary contribution of this research to the design of socially adept artificial systems may lie not in the divisions between types of social learning that have

occupied much of the research agenda, but rather in the spectrum of potential social learning situations and mechanisms these divisions highlight. Here, we draw attention to some of the most commonly cited ways in which one organism can potentially learn by observing another (the categories we use have been roughly adapted from Whiten 2004).

Attention Shifting. The animal's attentional focus is affected by others actions. This includes *stimulus enhancement*, where the observer becomes more likely to attend to and interact with stimuli it has noticed a model attending to.

Significance Learning. Using other's behavior and reactions as cues about the significance of objects in the environment. This includes *social referencing*, where the observer alters his reaction to a stimulus based on the observed behavior of a model, and *affordance learning*, where the animal learns certain properties of the environment, or of objects in the environment, through observation.

Impersonation. Copying the form of another's action. This category encompasses relatively simple behaviors such as *response facilitation*, where the observer becomes more likely to perform an action already in its repertoire, as a result of seeing a model perform that action, as well as behaviors such as *mimicry*,

Emulation and true imitation. In mimicry, the observer replicates the physical movements of the model, while in *emulation* it is the end-state generated by the model's actions that is replicated. In *true imitation* the observer attempts to replicate not only the model's actions, but also their perceived goals.

Learning About Others. Information about conspecifics is gathered over the course of interactions. This includes learning the positions of different group members in a dominance hierarchy, and *perspective-taking*, where one animal takes actions that take another's visual point-of-view into account. It also includes more advanced theory of mind, where one animal must model some aspect on another's internal mental state.

Imitation and Hierarchical Action Structures

Hierarchical, motivationally-driven behavior selection mechanisms have frequently been suggested in the animal behavior literature (see for instance Tinbergen 1951 and Dawkins 1976 for some classic examples). Timberlake (Timberlake 1989) has proposed a particularly detailed theory of hierarchical behavioral structures in animals, known as the behavior systems approach. According to Timberlake, an animal's action hierarchy is composed of behavioral systems, each of which is associated with an innate motivation or drive, such as feeding, self-defense, or socializing. Within a motivational system, each level of the hierarchy contains increasingly specific, sequentially organized actions for satisfying the associated drive (an example motivational system is shown in Figure 1).

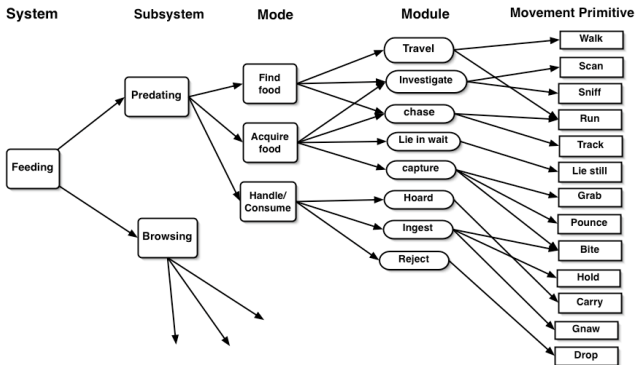

Figure 1: An example motivational system for feeding (after Timberlake 1989).

Using the idea of hierarchically organized action systems, such as those proposed by Timberlake, Bryne and Russon (Byrne and Russon 1996) have proposed another way in which to broaden the definition of imitation. They suggest that much animal imitation occurs at the "program" level, where an animal with a hierarchical action system learns a program for organizing its actions by observing the hierarchical structure of another animal's behavior. Subsequently, it is this hierarchical organization that is imitated, rather than the surface form of the other animal's movements. *Program-level imitation* is contrasted with what they define as *action-level imitation*, in which it is the specific physical movements of the model that are replicated. Byrne and Russon suggest that most task-oriented imitation is program-level imitation, whereas action-level imitation is more rare, and may serve a primarily social purpose. In addition to Byrne and Russon's observational studies, some support for this theory comes from Whiten's experimental demonstrations of imitation of sequential (and potentially hierarchical) action structures in chimpanzees, and imitation of hierarchical behaviors by young children (Whiten 2002).

Byrne and Russon's theory emphasizes the idea that imitation may operate at a number of levels, and outlines a possible mechanism by which this could occur—the perception and production of hierarchical action structures. They suggest that imitation occurs at multiple stages of the action hierarchy: from imitating individual movement primitives at the lowest level, to imitating the arrangement of behavioral modes and modules (to borrow Timberlake's terminology), to adopting the high-level goal or motivation at the top of the hierarchy.

Most previous work in robotic imitation has focused on teaching robots or animated characters individual actions meant to solve a particular task, taking advantage of only the lowest level of imitation. Since our behavior architecture is based on a hierarchical action system, we are in an excellent position to explore and take advantage of imitative learning at other levels of the action hierarchy.

Understanding Other's Minds

For artificial creatures to possess human-like social intelligence, they must be able to infer the mental states of others (e.g., their thoughts, intents, beliefs, desires, etc.) from observable behavior (e.g., their gestures, facial expressions, speech, actions, etc.). This competence is often referred to as a theory of mind (Premack and Woodruff 1978).

Simulation Theory. Simulation Theory (ST) is one of the dominant hypotheses about the nature of the cognitive mechanisms that underlie theory of mind (Gordon 1993). It can perhaps best be summarized by the cliché "to know a man is to walk a mile in his shoes." Simulation Theory posits that by simulating another person's actions and the stimuli they are experiencing using our own behavioral and stimulus processing mechanisms, humans can make predictions about the behaviors and mental states of others, based on the mental states and behaviors that we would possess in their situation. In short, by thinking "as if" we were the other person, we can use our own cognitive, behavioral, and motivational systems to understand what is going on in the heads of others.

From a design perspective, Simulation Theory is appealing because it suggests that instead of requiring a separate set of mechanisms for simulating other persons, we can make predictions about others by using our own cognitive mechanisms to recreate how we would think, feel, and act in their situation—thereby providing us some insight into their emotions, beliefs, desires, intentions etc. We argue that an ST-based mechanism could also be used by robots and animated characters to understand humans, and each other, in a similar way. Importantly, it is a strategy that naturally lends itself to representing the internal state of others and of the character itself in comparable terms. This would facilitate an artificial creature's ability to compare its own internal state to that of a person or character it is interacting with, in order to infer their mental states or to learn from observing their behavior. Such theories could provide a foothold for ultimately endowing machines with human-style social skills, learning abilities, and social understanding.

Imitation and Simulation Theory. Meltzoff proposes that the way in which infants learn to simulate others is through imitative interactions. For instance, Meltzoff (Meltzoff 1996) hypothesizes that the human infant's ability to translate the perception of another's action into the production of their own action provides a basis for learning about self-other similarities, and for learning the connection between behaviors and the mental states producing them.

Simulation Theory rests on the assumption that the other is enough "like me" that he can be simulated using one's own machinery. Thus, in order to successfully imitate and be imitated, the infant must be able to recognize structural congruence between himself and the adult model (i.e., notice when his body is "like" that of the caregiver, or when the caregiver's body is "like" his own). The initial "like me" experiences provided by imitative exchanges could lay the foundation for learning about additional behavioral and mental similarities between self and other.

There are a number of ways in which imitation could help bootstrap a Simulation Theory-type ToM (Meltzoff and Decety 2003). To begin with, imitating another's expression or movement is a literal simulation of their behavior. By physically copying what the adult is doing, the infant must, in a primitive sense, generate many of the same mental phenomena the adult is experiencing, such as the motor plans for the movement. Meltzoff notes that the extent to which a motor plan can be considered a low-level intention, imitation provides the opportunity to begin learning connections between perceived behaviors and the intentions that produce them.

From Social Animals to Social Characters

The cognitive literature provides compelling evidence for the presence and usefulness of social learning in human and non-human animals. Furthermore, a number of themes can be seen in the cognitive theories described here, which can be used to guide our design of socially capable artificial creatures.

Multiple Levels of Social Learning. Social learning and imitation may happen at many levels of behavioral granularity.

Multiple Sources of Information. There are multiple sources of information contained in an action, and each provides opportunities for different kinds of social learning.

Motivationally-driven Action Hierarchies. One possible way in which to represent multiple sources of information and multiple levels of behavioral granularity is by using a motivationally-driven hierarchical action structure.

Simulation Theory. Simulation Theory, where the creature uses itself to help interpret another's behavior, may be an especially useful approach to developing social abilities.

Bootstrapping from Imitation. Being able to identify and imitate another's behavior may be the first step towards more complex interpretation of that behavior.

With these points in mind, we would like to explore the ways in which Simulation Theory, and other cognitively inspired mechanisms, can be used by one character to learn from another's behavior. Our implementation of a simulation-theoretic social learning system uses hierarchical action structures, and attempts to exploit multiple levels of social learning and multiple sources of observational information. Finally, we use our characters' ability to recognize and reproduce observed movements as the starting point for developing more complex social skills, such as identifying simple motivations and goals, and learning about objects in the environment.

Max and Morris Mouse

Max and Morris are the latest in long line of interactive animated characters developed by the Synthetic Characters Group at the MIT Media Lab (see for example Burke et al. 2001, Blumberg et al. 2002, and Tomlinson et al. 2002). They were built using the Synthetic Characters C5m toolkit, a specialized set of libraries for building autonomous, adaptive characters and robots. The toolkit contains a complete cognitive architecture for synthetic characters, including perception, action, belief, motor and navigation systems, as well as a new, high performance graphics layer for doing Java-based OpenGL 3D Graphics. A brief introduction to a few of these systems will be given here, but it is beyond the scope of this paper to discuss them all in detail (for more information please see Burke et al. 2001, Blumberg et al. 2002). For a complete description of the social learning architecture discussed in this paper, and its implementation, please see Buchsbaum 2004.

The Motor System

The motor representation used by our characters is a multi-resolution, directed, weighted graph, known as a *posegraph* (introduced in Downie 2001). To create a character's *posegraph*, source animation material is broken up into poses corresponding to key-frames from the animation, and into collections of connected poses known as movements. Animations can be generated and played out on the character in real-time by interpolating down a path of connected pose nodes, with edges between nodes representing allowable transitions between poses. The graph represents the possible motion space of a character, and any motor action the character executes can be represented as a path through its posegraph.

Within the posegraph representation, *movements* are of particular importance to us here. Movements generally correspond to things we might intuitively think of as complete actions (e.g., sitting, jumping, waving), and therefore often match up closely with requests from the behavior system. While the pose representation provides us with greater motor knowledge and flexibility, the movement representation is often a more natural unit to work with. More critically, because movements correspond closely to motor primitives, or to simple behaviors, they also represent the level at which we would like to parse observed actions, in order to identify and imitate them. Therefore, inspired by Simulation Theory, our characters recognize and imitate actions they observe by comparing them with the movements they are capable of performing themselves, a process we will discuss in greater detail later on.

Action System

A character's *action system* is responsible for behavior arbitration—choosing what behavior the character engages in and when it does so. Individual behaviors are

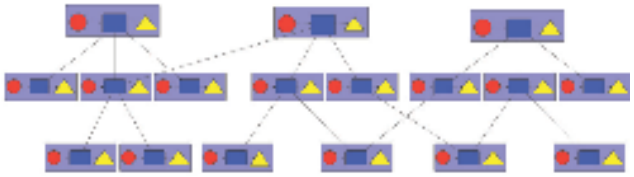

Figure 2: An example action system. Rectangles represent tuples. Circles are trigger contexts, triangles are objects, and rectangles are actions (do-until contexts not shown). There are three motivational subsystems in this example action

represented in our system as *action tuples* (Blumberg et al. 2002) and are organized into a hierarchical structure composed of motivational subsystems (which are described below). An example action system is shown in Figure 2. Each action tuple contains one or more *actions* to perform, *trigger contexts* in which to perform the action, an optional *object* to perform the action on, and *do-until contexts* indicating when the action has been completed.

The *action* is a piece of code primarily responsible for sending high-level requests for movements or movement sequences to the motor system. The requests can range from something relatively simple such as to "look at" an object, to more complex actions like "reach for the cheese". Actions in tuples towards the top of the hierarchy are more general (e.g. "satisfy hunger"), and become more specific farther down, with leaves in the action tree corresponding to individual requests to the motor system (e.g. "perform the eating movement"). Actions have associated values, which can be inherent (i.e. pre-programmed) or learned, and represent the utility of performing that action to the creature (for further discussion of action values see Blumberg et al. 2002).

Trigger contexts are responsible for deciding when the actions should be activated. In general, there are a variety of internal (e.g., motivations) and external (e.g. perceptions) states that might trigger a particular action. For instance, both the presence of food and the level of a character's hunger might be triggers for an *eat* action. Similarly, a tuple's *do-until contexts* decide when the action has completed.

Many behaviors, such as eating and reaching, must be carried out in reference to a particular object in the world. In our system, this object is known as the character's *object-of-attention*. In this work, all action tuples not at the top-level of the action hierarchy defer their choice of object to the tuple at the top of their motivational subsystem. Action tuples at the top of motivational hierarchies choose objects of attention most likely to satisfy the particular drive they serve (e.g. a *satisfy hunger* tuple might choose a nearby food object).

Action tuples are grouped into *action groups* that are responsible for deciding at each moment which tuple will be executed. Each action group can have a unique action selection scheme, and there can be only one tuple per

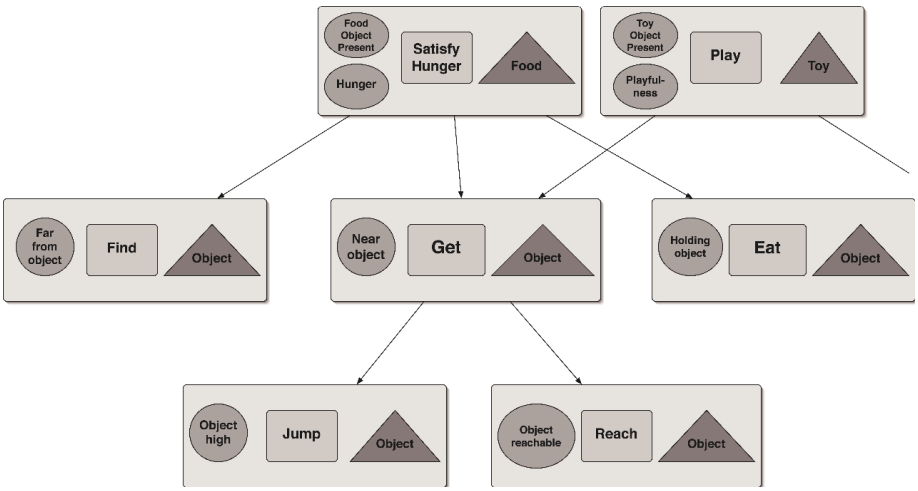

Figure 3: A simplified diagram of Max's hunger motivational subsystem (top-level of his play motivational system shown as well).

action group active at a time. All the action groups in this work use a probabilistic action selection mechanism, that chooses among all the tuples they contain based on their respective trigger and action values.

As mentioned earlier, the characters presented in this paper use an action system that is hierarchically organized and motivationally driven. This hierarchical organization means that each level of the action system has its own action group, containing increasingly specific, mutually exclusive, action tuples. At the top-level are tuples whose purpose is simply to satisfy a particular motivation or drive, such as a *play* or *hunger* drive. Since these tuples are in the same action group, only one of them may be active at a time, which keeps the character from dithering between competing drives. Below each of these motivational tuples, are tuples representing increasingly specific mechanisms for satisfying drives. For instance, below the *satisfy hunger* action tuple (whose sub-hierarchy is shown in Figure 3), are tuples such as *get food*, and *eat food*, and below *get food* are in turn *reach for food* and *jump for food*.

Figure 4: Left: Morris, as seen through Max's synthetic vision. Middle: Morris (left) and Max (right) in the virtual desert. Right: Morris covering his eyes, as seen by Max. Notice that some of the spheres marking his body parts are not visible.

Again, at each of these levels of the hierarchy, only one action tuple at a time may be active. For example, *satisfy hunger*, *get food* and *reach for food* could all be simultaneously active, but *reach for food* and *jump for food* cannot be active at the same time (which makes intuitive sense, since they would require the character to perform conflicting motions simultaneously). Finally, one important point about the hierarchical action structure used by the characters in this work is its striking similarity to the motivationally-driven hierarchical systems hypothesized by ethologists and cognitive scientists such as Whiten (Whiten 2002), Bryne and Russon (Byrne and Russon 1996), and Timberlake Timberlake 1989).

Imitation and Movement Recognition

Max the Mouse is able to observe and imitate his friend Morris's movements, by comparing them to the movements he knows how to perform himself. Max watches Morris through a color-coded *synthetic vision system*, which uses a graphical camera mounted in Max's head to render the world from Max's perspective. The color-coding allows Max to visually locate and recognize a number of key end-effectors on Morris's body, such as his hands, nose and feet (see Figure 4). Currently, Max is hard-wired to know the correspondence between his own effectors and Morris's (e.g. that his right hand is like Morris's right hand), but previous projects have featured characters using learned correspondences (Breazeal et al. 2004), and a similar extension is planned for this research.

As Max watches Morris, he roughly parses Morris's visible behavior into individual movements and gestures. Max locates points in time when Morris was momentarily still, or where he passed through a transitionary pose, such as standing, both of which could signal the beginning or end of an action. Max then tries to identify the observed

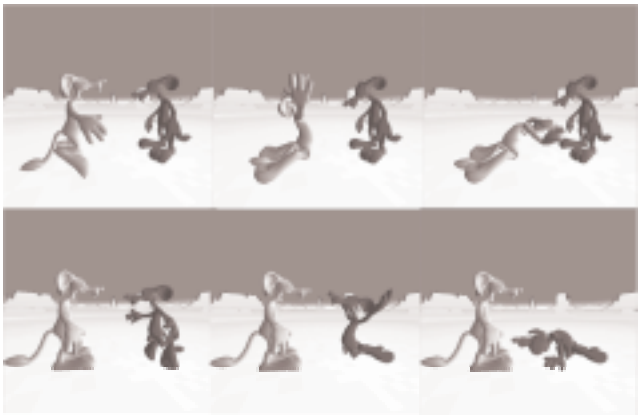

Figure 3: Top row: Morris demonstrates pounding the ground. Bottom row: Max imitates pounding the ground.

movement, by comparing it to all the movement representations contained within his own posegraph. To do this, Max compares the trajectories of Morris's effectors to the trajectories his own limbs would take while performing a given movement. This process allows Max to come up with the closest matching motion in his repertoire, using as few as seven visible effectors (as of writing, we have not tested the system using fewer than seven). By performing his best matching movement or gesture, Max can imitate Morris.

Action Identification

By matching observed gestures and movements to his own, Max is able to imitate Morris. Max can also use this same ability to try and identify which actions he believes Morris is currently performing. Max keeps a record of *movement-action correspondences*, that is, which action he is generally trying to carry out when he performs a particular movement (e.g. the 'reaching' gesture is most often performed during the 'getting' action). When he sees Morris perform a given movement, he identifies the action tuples it is most likely to be a part of. He then evaluates a subset of the trigger contexts, known as *can-I triggers*, to determine which of these actions was possible under the current circumstances. In this way, Max uses his own action selection and movement generation mechanisms to identify the action that Morris is currently performing.

Motivations and Goals

Another subset of trigger contexts, known as *should-I triggers*, can be viewed as simple motivations. For example, a should-I trigger for Max's eating action is hunger. Similarly, some do-until contexts, known as *success contexts*, can represent low-level goals. Max's success context for reaching for an object is holding the object in his hands. By searching his own action system for the action that Morris is most likely to be performing, Max can identify likely should-I triggers and success do-untils for Morris's current actions. For example, if Max sees Morris eat, he can match this with his own eating action, which is triggered by hunger, and know that Morris is probably hungry.

Similarly, Max can see Morris reaching for, or jumping to get, an object, and know that Morris's goal is to hold the object in his hands, since that is the success context for Max's own 'get' action. Notice that in this second case, Max does not need to discern the purpose of jumping and reaching separately, since these are both sub-actions of 'get' in his own hierarchy. We are currently developing mechanisms that allow Max to use the trigger and do-until information from his best matching action in order to interact with Morris in a more socially intelligent way. For instance, Max might see Morris reaching and help him get the object he is reaching for, bringing him closer to more advanced social behavior such working on cooperative tasks.

Learning About Objects

One important way in which Max can already learn by observing Morris is through a process similar to that of social referencing(which was described earlier in this paper). By watching Morris interact with unknown objects, Max can learn some of the affordances of these objects.

For example, let's say Max starts out knowing that cheese is edible, but not knowing anything about ice cream. Meanwhile, Morris knows that ice cream is an edible (and tasty) treat. If Max watches Morris reach for the ice cream and is asked to identify what Morris is doing he will shrug, indicating that he doesn't know why Morris is reaching. This is because none of the possible paths to the 'reach' tuple in Max's action system seem valid.

If however, Max sees Morris eat the ice cream cone, the story is different. At this point, Max notices something important—he only uses the 'eat' action tuple (and in turn the 'eating' movement) to satisfy one drive, because it is only part of one motivational subsystem. Since eating to satisfy hunger is the only purpose Max knows of for the eating movement, he checks to see if Morris could have been eating an unknown object. To do this, he re-evaluates his can-I triggers with a slight modification—he assumes that Morris's object-of-attention is a food object. In this example, Max would choose the ice cream as Morris's likely object of attention, and would find that Morris was in fact holding the ice cream, making it possible for him to be eating it. Max would conclude that Morris was eating the ice cream, and from this point on, Max would recognize ice cream as a potential food source (for further details on this process, please see Buchsbaum 2004).

Discussion: Simulation Theory as the Road to Social Characters

Simulation Theory is in many ways the unifying factor among the various social learning tasks and mechanisms tackled in our research. We chose to consistently use a

simulation-theoretic approach while addressing a wide variety of social learning problems, in part because of the strong supporting evidence from the cognitive research, but also in order to see just how far an artificial creature could get using itself as a model for other's behavior.

As it turns out, the answer appears to be: pretty far. From recognizing observed movements, to finding another character's object of attention, to identifying another's motivations and goals, Simulation Theory seems to be an effective approach to a range of social learning problems. Perhaps more importantly, as discussed in the following sections, it is often an approach that simplifies the problem at hand.

Movement Recognition and Motor Representation.

Motion parsing and movement recognition from visual data are extremely challenging problems, and currently represent very active research areas. In particular, on-line movement classification systems typically require a large set of training data, and rely on statistical models to extract motion features from the data that correlate with particular gestures (but see Bindiganavale 2000). However, by using their own movement repertoires as the example set, our characters are able to perform on-line movement classification without any training period, and using only a limited set of body part coordinates (However, unlike Bindiganavale's system, this system does not currently address characters of different morphology. This is an important drawback to our current system, and a critical area for future research to explore).

The Role of Motivationally-Driven Hierarchical Action Structures

In more cognitively complex tasks, the advantage provided by perception-production coupling can be even more striking, particularly when applied within a hierarchical action structure. For instance, when a character such as Max maps one of his action tuples onto the observed behavior of another character, he is not only provided with the information contained within that tuple—the action to be performed, the immediate goals of that action, and the environmental context that triggered it—but with all the information contained within the hierarchy that tuple belongs to.

When Max see Morris reaching for an object, such as a piece of cheese, and identifies this with his own reaching tuple, he knows not only that Morris is reaching for the cheese, but that he wants to get the cheese—if the cheese were higher up Morris might jump instead of reach, but both have 'getting' as their goal. Looking still higher up the hierarchy tells Max that Morris wants to get the cheese because he is hungry, and that he might want other food items as well.

Motivationally-driven hierarchical action structures seem to provide an ideal cognitive substrate for Simulation Theory, because they so neatly package together the key characteristics of an action—the movements involved, the motivations for the action, and the action's goal. Further, hierarchical structures facilitate not only the recognition of immediate goals and motivations, but also recognition of the hierarchy of goal-directed behavior that characterizes intentional action. In other words, using a hierarchical action structure produces a hierarchical intention structure, and allows such a structure to be recognized in others.

Building Blocks of Social Learning

Earlier in this paper, we introduced a number of different categories of social learning described in the cognitive and ethological literature. We also discussed a number of theories that suggested that these apparently different types of social learning might result from responding to different aspects of the stimuli, or might represent different uses of the same underlying behavior mechanisms and structures.

Our work on characters capable of social learning seems to strengthen the case for shared social learning mechanisms. At least in the artificial system presented here, it appears that, given a number of key mechanisms—namely movement recognition, a motivationally-driven hierarchical action structure and the ability to simulate another's point-of-view on that structure, a large number of seemingly disparate social learning abilities can be demonstrated. To give a few traditional examples:

Stimulus Enhancement. While this is not a skill our research has focused on particularly, stimulus enhancement can be accomplished within this architecture. In the system described here, the other character's object of attention is always identified. By subsequently adopting this object as its own object of attention, our character could demonstrate stimulus enhancement.

Mimicry and Movement or Action Level Imitation. Many different terms have been used for the simple reproduction of the movements produced by another. In this work we have used the terms imitation and mimicry. In any case, by reproducing the movement he sees Morris performing Max can mimic Morris.

'True' Imitation. Whenever Max identifies and demonstrates the action he believes Morris is performing, he copies not only the form of the movements, but the object they are aimed at, and what he believes to be their goal.

Emulation. Again, while it was not explicitly addressed in our research, the characters presented here are capable of emulating the results of an action. They would do this by focusing only on the matching do-until context of that action, rather than on the action as a whole.

Identifying Goals and Motivations. Perhaps most importantly, hierarchical action systems help a character to identify another's motivations and goals at multiple levels of granularity.

The ability of this system to potentially reproduce all of these forms of social learning seems to change the critical

research question from "which kind of social learning is occurring?" to "which kind of social learning is most appropriate here?" This is not a trivial question. What the previous list shows us is that attending to different aspects of the identified action, or different levels of the hierarchy, leads to different responses. How to tell whether the movement, or the object, or the result of the action is the critical part, whether the immediate goal or its parent is most relevant, are important and unresolved questions.

Unfortunately, success in an artificial system cannot definitively prove anything about natural systems—it can only suggest. Nevertheless, the success of this approach, coupled with previous ethological research pointing to the existence of hierarchical action systems in animals, lends support to the idea that differences in social learning abilities may represent differences in which levels of the hierarchy different animals imitate (in the broadest sense of the word), and which sources of information they attend to.

Conclusion

This paper presents an approach to creating imitative, interactive characters, inspired by the literature on human and animal social learning, and by the Simulation Theory view of social cognition. Additionally, it introduces our ongoing work towards creating animated characters and robots that are able to understand simple motivations, goals and intentions, a critical step in creating artificial creatures who are able to interact with humans and each other as socially capable partners. Finally, it discusses the role played by certain key cognitive mechanisms, such as hierarchical action structures and perception-production coupling, in facilitating social learning, and suggests that the success of these mechanisms in an artificial system underscores their potential importance in natural ones.

References

Bindiganavale, R.N. 2000. Building Parameterized Action Representations from Observations. PhD diss., Dept. of Computer and Information Science, University of Pennsylvania.

Blumberg, B.; Downie, M.; Ivanov, Y.; Berlin, M.; Johnson, M.P.; and Tomlinson, B. 2002. Integrated learning for synthetic characters, *ACM Transactions on Computer Graphics*, 21: 417–426.

Breazeal, C.; Buchsbaum, D.; Gray, J.; Gatenby, D.; and Blumberg, B. Learning from and about others: Towards using imitation to bootstrap social understanding in robots. *Artificial Life*, Forthcoming,

Buchsbaum, D. 2004. Imitation and Social Learning for Synthetic Characters. Master's thesis, Media Lab, Massachusetts Institute of Technology.

Burke, R.; Isla, D.; Downie, M.; Ivanov, Y.; and Blumberg, B. 2001. Creature smarts: The art and architecture of a virtual brain. In Proceedings of the 2001 Computer Game Developers Conference.

Byrne, R.W. and Russon, A.E. 1996. Learning by imitation: A hierarchical approach. *Behavioral and Brain Sciences* 21(5):667–684.

Call, J. and Carpenter, M. Three sources of information in social learning. 2002. In Dautenhahn, K. and Nehaniv, C.L. eds. *Imitation in Animals and Artifacts.* Cambridge, MA: MIT Press. 211–228.

Dawkins, R. 1976. Hierarchical organization: A candidate principle for ethology. In Bateson, P.P.G. and Hinde, R.A. eds. *Growing points in ethology.* Cambridge University Press, 1976.

Downie, M. behavior, animation and music: the music and movement of synthetic characters. 2001. Master's thesis, Media Lab, Massachusetts Institute of Technology.

Gordon, R. 1993. Folk psychology as simulation. *Mind and Language,* 1:158–171.

Heyes, C.M. and Galef, B.G. eds. 1996. *Social Learning in Animals: The Roots of Culture.* Boston, MA: Academic Press.

Meltzoff, A.N. The human infant as imitative generalist: A 20-year progress report on infant imitation with implications for comparative psychology. 1996. Heyes, C.M. and Galef, B.G. eds. *Social Learning in Animals: The Roots of Culture*. Boston, MA: Academic Press. 347–370.

Meltzoff, A.N. and Moore, M.K. 1997. Explaining facial imitation: A theoretical model. *Early Development and Parenting* 6:179–192.

Meltzoff, A.N. and Decety, J. 2003. What imitation tells us about social cognition: a rapprochement between developmental psychology and cognitive neuroscience. *Transactions of the Royal Society of London B,* 358:491–500.

Premack, D. and Woodruff, G. 1978. Does the chimpanzee have a theory of mind? Behavioral and Brain *Sciences* 1(4):515–526.

Timberlake, W. and Lucas, G.A. 1989. Behavior systems and learning: from misbehavior to general principles. In Mowrer, R.R. and Klein, S.B. eds. *Contemporary learning theories: Instrumental conditioning theory and the impact of biological constraints on learning*. Hillsdale, NJ: Erlbaum Associates. 237–275.

N. Tinbergen. 1951. *The Study of Instinct.* London: Oxford University Press.

Tomlinson, B.; Downie, M.; Berlin, M.; Gray, J.; Lyons, D.; Cochran, J.; and Blumberg, B. 2002. Leashing the alphawolves: Mixing user direction with autonomous emotion in a pack of semi-autonomous virtual characters. In *Proceedings of the Symposium on Computer Animation*.: ACM Press.

Whiten, A. Imitation of sequential and hierarchical structure in action: experimental studies with children and chimpanzees. In Dautenhahn, K. and Nehaniv, C.L. eds. *Imitation in Animals and Artifacts*. Cambridge, MA: MIT Press. 191–209.

A. Whiten. *The dissection of socially mediated learning.* Forthcoming.