

A Simulation-Theory Inspired Social Learning System for Interactive Characters

Daphna Buchsbaum
Icosystem Corporation
Cambridge, MA 02138, USA
daphna@icosystem.com

Bruce Blumberg and Cynthia Breazeal
Media Lab
Massachusetts Institute of Technology
Cambridge, MA 01239, USA
bruce@media.mit.edu
cynthiab@media.mit.edu

Andrew N. Meltzoff
Institute for Learning and Brain Sciences
University of Washington
Seattle, WA 98195, USA
meltzoff@u.washington.edu

Abstract— Humans (and many other animals), display a remarkably flexible and rich array of social competencies, including the ability to interpret, predict and react appropriately to the behavior of others. We believe that developing systems that have these same sorts of social abilities is a critical step in designing animated characters and robots who appear intelligent and capable in their behavior, and who are intuitive and engaging for humans to interact with. Research in human development suggests that 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.

In this paper, we present Max T. Mouse, an anthropomorphic animated mouse character who uses his own motor and action representations to interpret the behaviors he sees his friend Morris Mouse performing (a process known as Simulation Theory in the cognitive literature). Max's imitation and social learning system allows him to identify simple goals and motivations for Morris's behavior, which we consider a front-end component needed for the larger project of developing characters who act in accord with a full theory of mind.

Index Terms— Intelligent Robots, Sociable Robotics, Human-Robot Interaction (HRI), Animation, Cognitive Science

I. INTRODUCTION

We want to build animated characters and robots capable of rich social interactions with humans and each other. 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.

While very few species exhibit the most complex forms of imitation [1] [2], and perhaps no non-human animal possesses a full theory of mind [3], 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 [4] [5] [6].

A. Related Work

In recent years, numerous imitative robotic and animated characters have been developed, and it is beyond the scope of this paper to review them all (for a partial review see [7] [8] and the discussion in [9]). Similarly, work has been done on allowing robots to learn new movements or tasks through observation and instruction (examples include [10] [11] [12] [13] [14] among others).

However, in our system's case, the end goal is not imitation or the learning of new tasks and actions. Instead, imitation is used as a way of bootstrapping the character's understanding of another's motivations and goals. In other words, the processes and representations used to generate our character's behavior are also used to recognize and make sense of the behavior it observes, and important step towards the ultimate goal of socially intelligent animated characters and robots.

II. 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 beliefs, intents, emotions, desires, etc.) from observable behavior (e.g., their movements, facial expressions, speech, actions, etc.). This competence is often referred to as a 'theory of mind' [3] [15].

A. Simulation Theory

Simulation Theory is one of the dominant hypotheses about the nature of the cognitive mechanisms that underlie theory of mind (e.g. [16] [17]). It can perhaps best be summarized by the clichè "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 a Simulation Theory-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.

1) *From Imitation to Understanding Other Minds:* Meltzoff proposes that infants’ grasp that others are “like me”, and their experience with imitating conspecifics lies at the root of their coming to simulate and understand the minds of others. For instance, Meltzoff [4] hypothesizes that the human infant’s basic ability to translate the perception of another’s act into the production of their own act provides a basis for elaborating self-other similarities, and for learning the connection between behaviors and the mental states producing them (see [4], for a fuller description of this “like me” theory).

There are a number of ways in which imitation could help bootstrap the development of a theory of mind [6]. 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 to 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.

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 (broadly construed) social learning system uses hierarchical action structures (described below), 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.

B. Hierarchical Action Structures

Hierarchical, motivationally-driven behavior selection mechanisms have frequently been suggested in the animal behavior literature (see for instance [18] and [19] for some classic examples). Timberlake [20] 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).

Using the idea of hierarchically organized action systems, such as those proposed by Timberlake,

Byrne and Russon [21] suggest that imitation occurs at multiple stages of the action hierarchy: from imitating individual movement primitives at the lowest level, to imitating the sequential arrangement of behaviors, 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 one level of imitation. Since our behavior architecture is based on a hierarchical action

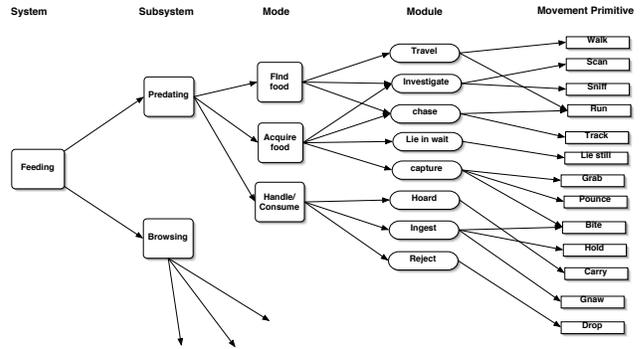


Fig. 1. An example motivational system for animal feeding (after Timberlake 1989 [20])

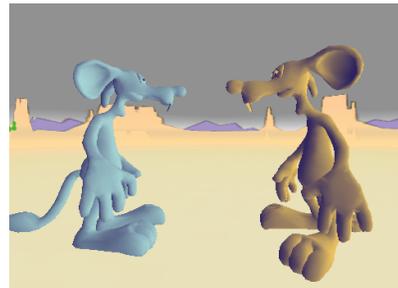


Fig. 2. Max (right) and Morris (left) in the virtual desert

system, we are in an excellent position to explore and take advantage of imitative learning at multiple levels of the action hierarchy.

III. MAX AND MORRIS MOUSE

Max and Morris (see figure 2) are the latest in long line of interactive animated characters developed by the Synthetic Characters Group at the MIT Media Lab (see for example [22] and [23]). 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. (For more information please see [22] [23] and [24].)

A. The Motor System

The motor representation used by our characters is a multi-resolution, directed, weighted graph, known as a *posegraph* (introduced in [25]). 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 movement 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 behaviors (e.g., sitting, jumping, waving), and therefore often match up closely with motor requests from the action system (described below). 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 movements, in order to identify and imitate them. Therefore, inspired by Simulation Theory, our characters recognize and imitate movements they observe by comparing them with the movements they are capable of performing themselves, a process we will discuss in greater detail later on.

B. Action System

A character's *action system* is responsible for behavior arbitration—choosing what behavior the character engages in and when it does so. While a movement is an individual motion primitive, **a behavior is a movement or series of movements placed in an environmental and motivational context**. Movements represent stand-alone physical motion, while behaviors are performed in response to environmental triggers, motivational drives and desired goal states. Individual behaviors are represented in our system as *action tuples* [23] and are organized into a hierarchical structure composed of motivational subsystems (which are described below). 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. Within each level of the action hierarchy, tuples compete probabilistically for expression, based on their action and trigger values.

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 [23]).

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 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.

The 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 4), are tuples such as *get food*, and *eat food*, and below *get food* are in turn *reach for food* and *jump for food*. Notice that the leaf tuples (e.g. *jump*, *reach*, *eat*) are generally the ones to request specific movements from the motor system (see figure 5).

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 Timberlake [20].

C. 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 body parts on Morris's body, such as his hands, nose and feet (see figure 3). Currently, Max is hard-wired to know the correspondence between his own body parts and Morris's (e.g. that his right hand is like Morris's right hand), but research with human infants [5] suggests that learning may underlie such mappings, and previous projects with artificial agents have also featured characters using learned correspondences [9], and a similar extension is planned for this research.

As Max watches Morris, he roughly parses Morris's visible behavior into individual movements. Max locates points in time when Morris was momentarily still, or where he passed through a transitional pose, such as standing, both of which could signal the beginning or end of a movement. Max then tries to identify the observed movement, by comparing it to all the movement representations contained within his own posegraph. To do this, Max compares the trajectories of Morris's body parts to the trajectories his own limbs would take while performing a given movement. This process allows Max to come up with the closest matching movement in his repertoire, using as few as seven visible body parts (as of writing, we have not tested the system using fewer than seven). By performing his best matching movement Max can imitate Morris.

D. Action Identification

By matching observed movements to his own, Max is able to imitate Morris. Max can also use this same ability to try and identify which set of action tuples (i.e. which path through his action hierarchy) he believes Morris is currently performing. Max keeps a record of *movement-action tuple correspondences*, that is, which behavior he is generally trying to carry out when he performs a particular movement (e.g. the *reaching* movement is most often performed during the *getting* action tuple). 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 tuples' trigger contexts to determine which of these action tuples was possible under the current circumstances. These trigger contexts are known as *can-I triggers* because they fire based on the presence of conditions

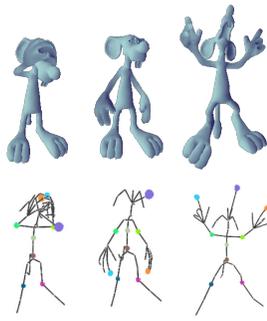


Fig. 3. This figure shows Morris in 3 poses. The top row is Morris as we see him, while the bottom row is Morris as seen through Max's synthetic vision. The colored spheres on Morris's body are key body parts whose location is tracked by the synthetic vision system.

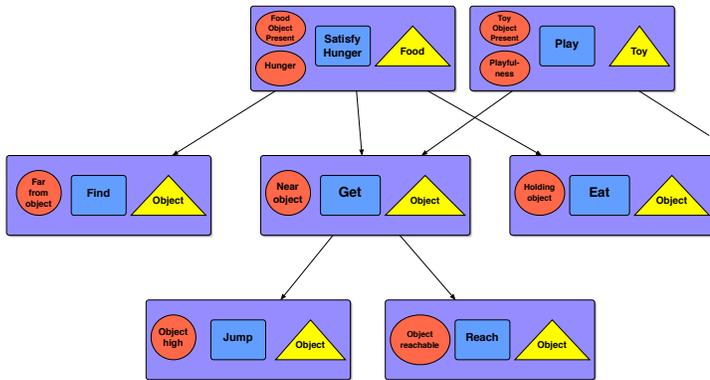


Fig. 4. A simplified diagram of Max's hunger motivational subsystem (the top level of his play motivational subsystem is shown as well). Large rectangles represent *action tuples*. Circles are *trigger contexts*, triangles are *objects*, and small rectangles are *actions* (do-until contexts not shown).

which make a behavior possible to carry out. The following sections describe this process in greater depth.

1) *An Identification Example:* Let's say that Max sees Morris reaching for a piece of cheese in order to eat it. How does Max identify that behavior as reaching for the cheese, and how does he know that it is part of the *hunger* motivational subsystem (shown in figure 4)?

When Max sees Morris reach, the first thing he does is identify the movement he sees Morris performing (as described in section III-C). In this case that movement is *reaching*. Next, Max searches his map of movement-action correspondences to find out which of the bottom-level (i.e. leaf) action tuples in his action hierarchy have requested this movement in the past. Recall that a movement is an individual motion primitive, such as *reaching*, *jumping* or *eating*, while an action tuple is a behavior that occurs in a motivational and environmental context, and requests that the motor system carry out particular movements. In this case, Max finds that he has only performed the *reaching* movement during his *reach* action tuple (figure 5), so it's likely that Morris was reaching for an object.

Next, Max traces back up his action hierarchy from the reaching action tuple. By tracing up his action hierarchy from the *reach* tuple,

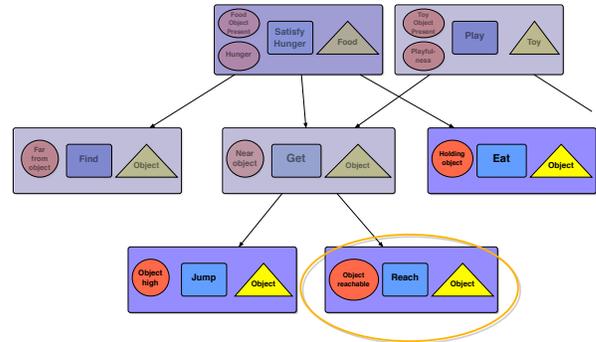


Fig. 5. Identifying the leaf action tuples where the *reach* movement is used

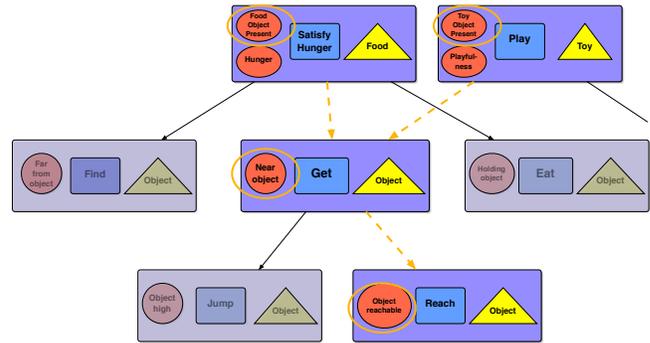


Fig. 6. Evaluating the *can-I* triggers along the paths to Max's reaching action tuple

Max finds that *reach* is part of the *get* tuple, which is used by a number of motivational subsystems. In the example shown in figure 6, *get* is used by both the *hunger* and *play* motivational subsystems. Since Max uses the *reach* action tuple in a number of motivational contexts, he must decide which one of these contexts best matches Morris's current situation, in order to decide *what* Morris is reaching for, and *why* he is reaching for it. To do this, Max once again uses a simulation theoretic approach—he checks which path through his action hierarchy it would have been possible for him to follow if he were in Morris's place.

Max evaluates the *can-I* triggers at each level of his action hierarchy, on the way to the *reach* action tuple (shown in figure 6). In this simple example, there are two *can-I* triggers—*food object present* and *toy object present*—at the top level of the hierarchy. These triggers are known as *object selection* triggers, since they check whether any appropriate objects are available and then, if they are, select one of them as the *object of attention* for all the action tuples below this one in the hierarchy. One subtle but important point here is that Max must check whether there were any food or toy objects available **at the time he saw Morris reaching**, not whether they are available currently, which is what he would check if he himself were hungry or playful.

Here, Max finds that there was a piece of cheese present when Morris began reaching, but no toy object present. Max can therefore eliminate the path from the *play* action tuple to the *reach* action tuple, leaving only the path from *hunger* to *reach* as a possible behavior Morris could have been performing. Now, Max must verify that it was possible for Morris to be reaching for the piece of cheese in order to satisfy his hunger. To do this, Max simply checks

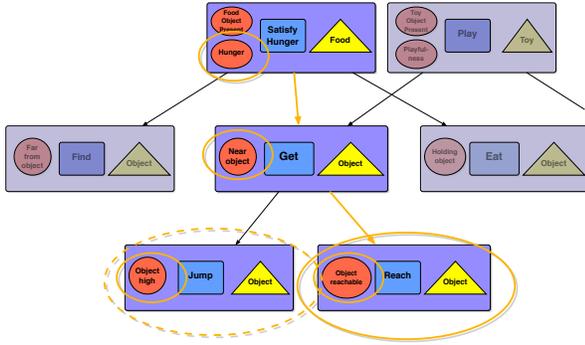


Fig. 7. If the observed movement is ambiguous, or part of multiple bottom level tuples, evaluating the *can-I* triggers can help identify the observed behavior. Here, Max determines that Morris was reaching rather than jumping because the food object was within reach rather than high in the air.

the remaining *can-I* triggers between the *hunger* and *reach* tuples (figure 6).

The next *can-I* trigger is a *proximity* trigger for the *get* action, which checks whether the *object of attention*—in this case the piece of cheese—is close enough to ‘get’ (but not so close that the character is already holding it). If the piece of cheese is close enough to Morris that he could have gotten it by reaching or jumping (without needing to walk anywhere), then this trigger will return true. Now, there is only one *can-I* trigger remaining—the *object reachable* trigger for the *reach* tuple, which is another proximity trigger that checks whether the object is within reach. Since the cheese was within Morris’s reach, this trigger will also return true, and Max concludes that Morris was reaching for the piece of cheese in order to satisfy his hunger (see also figure 7 for an example with multiple possible bottom level tuples).

E. 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-tuple is *hunger* (see figure [ref:fig:hungersys2]). 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 (figure 4). By searching his own action system for the action-tuple that Morris is most likely to be performing, Max can identify likely *should-I* triggers and *success do-untils* for Morris’s current behavior. For example, if Max sees Morris eat, he can match this with his own eating action-tuple, 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-tuple. Notice that in this second case, Max does not need to discern the purpose of jumping and reaching separately, since these are both sub-behaviors of *get* in his own hierarchy. (For related theory on how human infants use their own behavior to decode the goals and intentions of others, see [26] [4]). We are currently developing mechanisms that allow Max to use the trigger and *do-until* information from his best matching action-tuple 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 as working on cooperative tasks.

F. Learning About Objects

Finally, 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 (the hunger 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 (for further details on the implementation of this process, please see [24]).

IV. DISCUSSION AND FUTURE WORK

We have seen that Max the Mouse (and by extension, other animated characters and robots developed using the Synthetic Characters architecture), is capable of imitating other characters, identifying simple motivations and goals for their behavior, and learning from their actions. Furthermore, the implementation of these abilities relies strongly on mechanisms and approaches suggested by the cognitive literature, such as hierarchical action structures and Simulation Theory (for a discussion of implications for cognitive science, please see [24]). While this represents an important component of the larger project of creating characters who act in accord with a full theory of mind, there is a ways to go before we achieve this goal. The following sections describe just a few of the ways in which this work could be extended.

A. Understanding Emotions

This paper has primarily focused on goal-directed actions—actions which try to to satisfy a particular motivation or carry out an intention, generally towards an object in the world. However, humans often perform actions that are instead emotionally communicative—conveying a particular affective state. People smile, shrug, give a thumbs up, cross their arms and frown, wring their hands and so on. Correctly interpreting these sorts of affective movements is a critical part of human social interaction [27]—emotion recognition is even considered a significant predictor of social competence in children [28].

The characters Max and Morris already have a large repertoire of emotionally significant actions, and a number of autonomic variables devoted to their current emotional state. By applying the action identification techniques described in this paper to emotionally driven portions of their action systems, Max and Morris could use Simulation Theory to identify each other’s emotions. (For related work on agents who interpret and display emotion see for example Breazeal’s work with Kismet [29] and Picard’s Affective Computing research [30]).

B. Cooperative Behavior

One critical aspect of human social behavior is our ability to cooperate and work on joint tasks. The American Heritage Dictionary defines cooperating as “to work together towards a common end or purpose” [31]. Thus, in order for one character to engage in cooperative behavior with another it must be able to recognize and adopt the other’s goal.

The first step of cooperative behavior, recognizing the other’s goal, has already been accomplished (at least at a low level) in this work. Max is able to identify the goals of Morris’s actions by looking at the goals of his own actions. We are currently implementing the second step—acting on this goal in order to help accomplish it.

Finally, there are multiple levels of goals in the action hierarchy, and it will be an interesting question to explore which ones to help satisfy when. For instance, if Max sees Morris reaching for a piece of cheese, instead of trying to get Morris the out-of-reach cheese (satisfying a low-level goal), Max could recognize Morris’s hunger, and get him some ice cream instead (satisfying a higher level goal). for some related work on robot-human collaboration see [32] and [33] and on robot perspective-taking see [34]).

ACKNOWLEDGMENT

We would like to thank the members of the Synthetic Characters and Robotic Life Groups of the MIT Media Lab for their help with, and contributions to, this project.

REFERENCES

- [1] K. Dautenhahn and C. Nehaniv, Eds., *Imitation in animals and artifacts*. Cambridge, MA: MIT Press, 2002.
- [2] C. Heyes and B. Galef, Eds., *Social Learning in Animals: The Roots of Culture*. Boston: Academic Press, 1996.
- [3] D. Premack and G. Woodruff, “Does the chimpanzee have a theory of mind?” *Behavioral and Brain Sciences*, vol. 1, no. 4, pp. 515–526, 1978.
- [4] A. Meltzoff, “Imitation and other minds: The ‘like me’ hypothesis,” in *Perspectives on Imitation: From Cognitive Neuroscience to Social Science*, S. Hurley and N. Chater, Eds. Cambridge: MIT Press, March 2005, vol. 2, pp. 55–77.
- [5] A. Meltzoff and M. Moore, “Explaining facial imitation: A theoretical model,” *Early Development and Parenting*, vol. 6, pp. 179–192, 1997.
- [6] A. Meltzoff and J. Decety, “What imitation tells us about social cognition: a rapprochement between developmental psychology and cognitive neuroscience,” *Transactions of the Royal Society of London B*, vol. 358, pp. 491–500, 2003.
- [7] C. Breazeal and B. Scassellati, “Robots that imitate humans,” *Trends in Cognitive Sciences*, vol. 6, pp. 481–487, 2002.
- [8] S. Schaal, “Is imitation learning the route to humanoid robots?” *Trends in Cognitive Sciences*, vol. 3, pp. 233–242, 1999.
- [9] C. Breazeal, D. Buchsbaum, J. Gray, D. Gatenby, and B. Blumberg, “Learning from and about others: Towards using imitation to bootstrap the social understanding of others by robots,” *Artificial Life*, vol. 11, no. 1-2, January 2005.
- [10] C. Atkeson and S. Schaal, “Robot learning from demonstration,” in *International Conference on Machine Learning*, 1997, pp. 12–20.
- [11] S. Schaal, “Learning from demonstration,” in *Advances in Neural Information Processing Systems*, M. Mozer, M. Jordan, and T. Petsche, Eds. Cambridge, MA: MIT Press, 1997, vol. 9, pp. 1040–1046.
- [12] A. Fod, M. Mataric, and O. Jenkins, “Automated derivation of primitives for movement classification,” *Autonomous Robots*, vol. 12, no. 1, pp. 39–54, January 2002.
- [13] M. Mataric, V. Zordan, and M. Williamson, “Making complex articulated agents dance; an analysis of control methods drawn from robotics, animation, and biology,” *Autonomous Agents and Multi-Agent Systems*, vol. 2, no. 1, pp. 23–44, March 1999.
- [14] N. Nicolescu and M. Mataric, “Natural methods for robot task learning: Instructive demonstrations, generalization and practice,” in *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Australia, July 2003.
- [15] H. Wellman, D. Cross, and J. Watson, “Meta-analysis of theory-of-mind development: The truth about false belief,” *Child Development*, vol. 72, pp. 655–684, 2001.
- [16] A. Goldman, “Imitation, mind reading, and simulation,” in *Perspectives on Imitation: From Cognitive Neuroscience to Social Science*, S. Hurley and N. Chater, Eds. Cambridge: MIT Press, March 2005, vol. 2, pp. 79–93.
- [17] R. Gordon, “Intentional agents like myself,” in *Perspectives on Imitation: From Cognitive Neuroscience to Social Science*, S. Hurley and N. Chater, Eds. Cambridge: MIT Press, March 2005, vol. 2, pp. 95–106.
- [18] N. Tinbergen, *The Study of Instinct*. Oxford University Press, 1951.
- [19] R. Dawkins, “Hierarchical organization: A candidate principle for ethology,” in *Growing points in ethology*, P. Bateson and R. Hinde, Eds. Cambridge University Press, 1976.
- [20] W. Timberlake and G. Lucas, “Behavior systems and learning: from misbehavior to general principles,” in *Contemporary learning theories: Instrumental conditioning theory and the impact of biological constraints on learning*, R. Mowrer and S. Klein, Eds. Hillsdale, NJ: Erlbaum Associates, 1989, pp. 237–275.
- [21] R. Byrne and A. Russon, “Learning by imitation: A hierarchical approach,” *Behavioral and Brain Sciences*, vol. 21, no. 5, pp. 667–684, 1998.
- [22] R. Burke, D. Isla, M. Downie, Y. Ivanov, and B. Blumberg, “Creature smarts: The art and architecture of a virtual brain,” in *Proceedings of the 2001 Computer Game Developers Conference*, 2001.
- [23] B. Blumberg, M. Downie, Y. Ivanov, M. Berlin, M. Johnson, and B. Tomlinson, “Integrated learning for synthetic characters,” in *Proceedings of the 29th annual ACM conference on Computer Graphics and Interactive Techniques*, vol. 21. SIGGRAPH, 2002, pp. 417–426.
- [24] D. Buchsbaum, “Imitation and social learning for synthetic characters,” Master’s thesis, Massachusetts Institute of Technology, Media Laboratory, September 2004.
- [25] M. Downie, “behavior, animation and music: the music and movement of synthetic characters,” Master’s thesis, Massachusetts Institute of Technology, Media Lab, 2001.
- [26] A. Meltzoff, “Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children,” *Developmental Psychology*, vol. 31, pp. 838–850, 1995.
- [27] D. Keltner, “Social functions of emotions at four levels of analysis,” *Cognition and Emotion*, vol. 13, pp. 505–521, 1999.
- [28] J. Leppanen, “Does emotion recognition accuracy predict social competence in childhood?” *Psychologia*, vol. 36, pp. 429–438, 2001.
- [29] C. Breazeal, “Emotion and sociable humanoid robots,” *International Journal of Human-Computer Studies*, vol. 59, pp. 119–155, 2003.
- [30] R. Picard, *Affective Computing*. Cambridge: MIT Press, 1997.
- [31] *The American Heritage Dictionary*, 4th ed. Houghton Mifflin Company, 2000.
- [32] C. Breazeal, G. Hoffman, and A. Lockerd, “Teaching and working with robots as a collaboration,” in *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi Agent Systems*. AAMAS, July 2004.
- [33] A. Lockerd and C. Breazeal, “Tutelage in socially guided robot learning,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. IROS 2004, forthcoming 2004.
- [34] J. Trafton, N. Cassimatis, M. Bugajska, D. Brock, F. Mintz, and A. Schultz, “Enabling effective human-robot interaction using perceptive-taking in robots,” in *Proceedings of the International Conference on Systems, Man and Cybernetics*. IEEE, to appear 2005.