Toward Teaching a Robot "Infant" using Emotive Communication Acts

Cynthia Breazeal (Ferrell) and Juan Velásquez
Massachusetts Institute of Technology
Artificial Intelligence Laboratory
545 Technology Square, Room 938, 935
Cambridge, MA 02139 USA
email: cynthia@ai.mit.edu, jvelas@ai.mit.edu

Abstract

This paper presents ongoing work towards building an autonomous robot that learns in a social context. The mode of social interaction is that of a caretaker-infant pair where a human acts as the caretaker for the robot. By placing our robot, Kismet, in an environment with a human caretaker who actively assists and guides Kismet's learning, this work explores robot learning in a similar environment to that of a developing infant. In doing so, this approach attempts to take advantage of this special sort of environment and the social interactions it affords in facilitating and constraining learning. This paper proposes an approach where emotive channels of communication are employed during social robot-human interactions to shape and guide what the robot learns.

keywords: human-robot interaction, scaffolding, emotions

Introduction

Our work focuses not on robot-robot interactions, but rather on the construction of robots that engage in meaningful social exchanges with humans. By doing so, it is possible to have a socially sophisticated human assist the robot in acquiring more sophisticated communication skills. Specifically, the mode of social interaction is that of a caretaker-infant pair where a human acts as the caretaker for the robot. By treating the robot, Kismet, as an *altricial system*[1] whose learning is assisted and guided by the human caretaker, this work explores robot learning in a similar environment to that of a developing infant.

It is known that an infant's emotions and drives play an important role in generating meaningful interactions with the caretaker (Bullowa 1979). These interactions constitute learning episodes for new communication behaviors. In particular, the infant is strongly biased to learn communication skills that result in having the caretaker satisfy the infant's drives (Halliday 1975). The infant's emotional responses provide important cues which the caretaker uses to assess how to satiate the infant's drives, and how to carefully regulate the complexity of the interaction. The former is critical for the infant to learn how its actions influence the caretaker, and the latter is critical for establishing and maintaining a suitable learning environment for the infant. Similarly, the caretaker's emotive responses to the infant shape the continuing interaction and can guide the learning process.

The behavior engine of Kismet (our robot) is designed to generate analogous interactions for a robot-human pair as for an infant-caretaker couple. As such, it integrates perception, attention, behavior, motivations, motor skills and expressive acts. Since an infant's emotions and drives guide and shape much of his behavior, the robot's motivational system also plays a prominent role in influencing the robot's focus of attention, behavior, expressive acts, and learning.

Previous work has addressed how the behavior engine uses motivations and facial expressions to maintain an appropriate level of stimulation during social interaction with humans (Breazeal(Ferrell) 1998a,b). This is a critical skill for the kinds of social learning that mothers and infants engage in, for it helps the mother tune her actions so that they are appropriate for the infant. This paper presents work in progress to extend Kismet's behavior engine to incorporate a learning mechanism for acquiring emotional memories as described in (Velásquez 1998), which is inspired by the experimental findings of (Damasio 1994) and (LeDoux 1996). These emotional memories provide a bridge by which the caretaker can provide the robot with rich and ongoing forms of reinforcement

[1] We want to emphasize that Kismet relies upon nurturing acts of the caretaker to shape and guide what it learns. This is in contrast to simply being a neonatal-like system which may or may not have access to a benevolent caretaker in its environment.

during social interactions through emotive channels of communication, which in turn can be used to shape and guide what the robot learns.

Kismet is shown in figure 1 displaying a range of emotive expressions analogous to anger, fatigue, fear, disgust, excitement, happiness, interest, sorrow, and surprise. It consists of two active stereo systems, vision and audio, embellished with facial features for emotive expression.

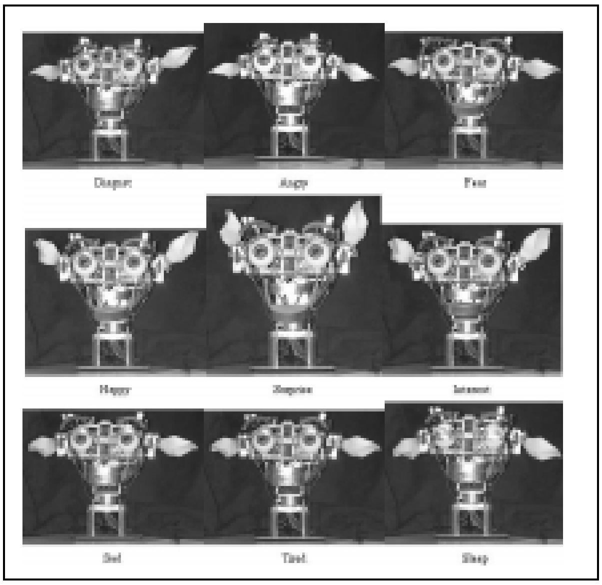

Figure 1 Kismet displays a variety of facial expressions corresponding to emotive and expressive states.

This paper is organized as follows: first we discuss relevant work in developmental psychology and emotion theory that has inspired our use of emotion-based mechanisms as a basis for learning in a social infant-caretaker context. These ideas have shaped the mechanisms proposed in our approach. We then argue for several special implications such a mechanism would have, with respect to learning new behavior, when situated in a socio-cultural environment. Finally, we present the state of the work in progress, summarizing some early results and discuss planned extensions.

The Role of Motivations in Learning Behavior

Motivations encompass drives, emotions, and pain which all play several important roles for both arbitrating and learning behavior. In Ethology, much of the work in motivation theory tries to explain how animals engage in appropriate behaviors at the appropriate time to promote survival (Tinbergen 1951, Lorenz 1973). For animals, internal drives influence which behavior the animal pursues, for example, feeding, foraging, or sleeping. Furthermore, depending on the intensity of the drives, the same sensory stimulus may result in very different behavior. For example, a dog will respond differently to a bone when it is hungry than when it is fleeing from danger.

It is also well accepted that animals learn things that facilitate the achievement of biologically significant goals. Work in Ethology has argued that motivations provide an impetus for this. In particular, the motivational system provides a reinforcement signal that guides what the animal learns and in what context. For instance, operant learning can be viewed as an animal's discovery of new appetitive behaviors (applying an existing skill in a new stimulus situation) that brings it closer to attaining some goal. When an animal has a strong drive that it is trying to satisfy, it is primed to learn new appetitive behaviors. Appetitive behaviors bring the animal into an appropriate relation with the world such that the relevant consummatory behavior can become active and satiate that drive. For this reason, it is much easier to train a hungry animal with a food reward than a satiated one (Lorenz 1973).

As another example, Mowrer (1960) postulates that learning is best thought of in two stages. In the first stage the animal learns to associate a particular emotion with a specific stimulus, and in the second stage the animal learns what behavior serves to alleviate that emotion. For instance, in Mowrer's experiments a rat learns to leave a box upon hearing a tone (previously that tone is paired with a painful shock). According to Mowrer's theory, first the rat learns to fear the tone and afterwards learns that leaving the box reduces that fear. It is argued that the emotional state allows for more flexible learning while providing a strong source of motivation for exploring other alternatives if early attempts fail. In Mowrer's case, he argues that the emotional state explains why the rat will explore alternate routes to reduce its fear if a barrier blocks the exit.

For a robot, an important function of the motivation system (emotions, drives, and pain) is to regulate behavior selection so that the observable behavior appears coherent, appropriately persistent, and relevant given the internal state of the robot and the external state of the environment. In this way, drives and emotions bias and influence behavior (and the manner in which they are expressed) but do not completely determine it. Emotions and drives are also necessary for establishing the context and impetus for learning as well as providing a reinforcement signal. Previous work in autonomous agent research has used drives as a mechanism for behavior arbitration (Maes 1990, McFarland & Bosser 1993), and Blumberg (1996) used motivations (called *internal variables*) to implement operant conditioning so that human user could train an animated dog new tricks.

Learning with a Human Caretaker

Parents take an active role in shaping how and what infants learn by means of *scaffolding*. Traditionally, scaffolding is thought of in social terms where a more able adult (the caretaker) manipulates the infant's interactions with the environment to foster novel abilities. Commonly, it involves reducing distractions, marking the task's critical attributes, reducing the number of degrees of freedom in the target task, providing ongoing reinforcement through expressive displays of face and voice, and enabling the subject to experience the end or outcome of a sequence of activity before the infant is cognitively or physically able of seeking and attaining it for himself (Wood, Bruner & Ross 1976). During this social process, the use of emotional expressions and gestures facilitates and biases learning throughout these exchanges. The emotive cues the parent receives during these social exchanges serve as feedback so the parent can adjust the nature and intensity of the structured learning episode to maintain a suitable learning environment where the infant is neither bored nor over-whelmed.

This view of scaffolding emphasizes the intentional contribution of the adult in providing conscious support and guidance to enable the infant to learn new skills. The scaffolding is a pedagogical device where the adult exploits temporarily engineered emergence of function by the infant to push him a little beyond his current abilities and in the direction the adult wishes him to go. For instance, by exploiting the infant's instinct to perform a walking motion when supported vertically in the standing position, the adult encourages him to learn how to walk before he is physically able. The infant's repeated experience of unplanned outcomes helps him discover new possibilities for effective action and to foster permanent adaptive change. This kind of learning is *serendipitous*, i.e. learning via accidental, fortunate discovery.

In addition, during early interactions with the caretaker, an infant's motivations and emotional displays are critical in establishing the foundational context for learning episodes, from which the infant can learn shared meanings of communicative acts (Bullowa 1979). During early face-to-face exchanges with his mother, an infant displays a wide assortment of emotive cues such as coos, smiles, waves, and kicks. At such an early age, the infant's basic needs, emotions, and emotive expressions are among the few things his mother thinks they share in common. Consequently, she imparts a consistent meaning to the infant's expressive gestures and expressions, interpreting them as meaningful responses to her mothering and as indications of the infant's internal state. Curiously, experiments by Kaye as reported in Bullowa (1979) argue that the mother actually supplies most if not *all* the meaning to the exchange when the infant is so young. The infant does not know the significance his expressive acts have for his mother, nor how to use them to evoke specific responses from her. However, because the mother *assumes* her infant shares the same meanings for emotive acts, her consistency *allows* the infant to discover what sorts of activities on his part will get specific responses from her. Routine sequences of a predictable nature can be built up which serve as the basis of learning episodes (Bullowa 1979). Furthermore, it provides a context of mutual expectations.

For example, early cries of an infant elicit various care-giving responses from his mother depending upon how she initially interprets these cries and how the infant responds to her mothering acts. Over time, the infant and mother converge on specific meanings for different kinds of cries. Gradually the infant uses subtly different cries (i.e., cries of distress, cries for attention, cries of pain, cries of fear) to elicit different responses from his mother. The mother reinforces the shared meaning of the cries by responding in consistent ways to the subtle variations. Evidence of this phenomena exists where mother-infant pairs develop communication protocols different from those of other mother-infant pairs (Bullowa 1979).

This form of scaffolding is referred to as *emergent scaffolding* by (Hendriks-Jansen 1996). It relies on the mother-infant dyad being seen as two tightly coupled dynamic systems. In contrast to the previous case where the adult guides the infant's behavior to a desired goal state, here the ongoing activity arises from the continuous mutual adjustments between the two participants. For instance, the interaction between a suckling infant and the mother who jiggles him whenever he pauses in feeding creates a recognizable interactive pattern that emerges from low level actions. This pattern of behavior encourages the habit of turn taking upon which face to face exchanges will later be built. Hence, some activity patterns exhibited by newborns have no place in adult behavior. They may simply serve as a bootstrapping role to launch the infant into an environment of adults who think in intentional terms, communicate through language, use tools, and manipulate artifacts. Within this cultural context, these same skills are transferred from adult to infant.

For our purposes, we now focus our attention on a single scaffolding act for the remainder of this paper: providing the robot with rich, ongoing reinforcement. For humans, the caretaker supplies the infant with reinforcement though emotive acts of communication, so we propose to do the same for Kismet. This is a rich channel of communication since it can influence the robot at an emotional level, and as argued above, the robot is designed to have an innate bias to behave in a way that maintains a positive emotional state.

This form of scaffolding is also interesting because it can be used in both a traditional as well as emergent manner. Used in the traditional manner, if the caretaker has a particular behavior she wants the robot to learn, she can tune her assistive acts to reinforce that particular behavior. For instance, the caretaker can positively reinforce forward progress by showing signs of happiness, negatively reinforce backward progress by showing signs of distress, and explicitly notify the robot of success by exhibiting happy excitement with a touch of surprise for emphasis. If necessary, she can take an active role in helping the robot experience the desired behavior before it is capable of achieving success on its own. However, this is an inherently social process, and the caretaker must continually monitor the robot and tune her reinforcing acts so they are appropriate. Hence, even though the caretaker assumes the intentional stance of teaching the robot a particular behavior, she cannot ignore that she and the robot form a mutually regulatory pair where each member's behavior affects that of the other. By consistently and repeatedly engaging in this process, the robot could eventually learn to associate a positive emotional state with the desired behavior. This effectively "tags" that behavior as being worthy of pursuit in its own right.

Toward a Mechanism: Insights from Emotion Theory

Given that we propose to provide the robot with rich, ongoing reinforcement though emotive channels of communication, how might one design a mechanism(s) to accomplish this? Specifically, how could one transfer an emotion-based signal displayed on the caretaker's face (or heard in her voice) to the robot's emotional system? Furthermore, once the robot has access to that emotive information, how can it learn to associate that emotive content with the immediate behavior (stimulus condition and action)? Fortunately, there are several insightful works from emotion theory that can help us.

The first insight comes from the study of facial expressions in infants. Trevarthen (1979) among others have argued that infants are born with a rich repertoire of facial expressions. Specifically they note that newborns exhibit facial expressions which are similar to their adult counterparts for the emotions of pleasure, displeasure, fear, surprise, confusion and interest. A more controversial finding is that infants are born with the ability to mimic the facial expressions of adults (Meltzoff & Moore 1977), as well as adult vocalizations (Bullowa 1979).

Along a different line of research, the role of facial expression in activating and regulating emotional experience (i.e. emotion feeling state) remains controversial. However several studies offer evidence that naturally occurring emotion expression can determine or influence emotion experience (Tomkins 1962, Izard 1990). For instance, these studies suggest that a person can make himself feel happier simply by smiling. Izard (1990) argues that the neural substrates and interconnections of involuntary facial expressions suggest that they are a relatively efficient and direct mechanism for regulating emotion experience. He further argues that these circuits are likely to be present in infants since the involuntary expressions are encoded in the genome and probably have innate connections to the neural substrates of emotion experience. Given Izard's arguments, it is conceivable that spontaneous facial expressions of an infant can influence its emotional state as well.

By combining these two theories, one can envision a mechanism for socially transferring positive or negative reinforcement from the caretaker to the robot during social interactions. For instance, when the robot does something the caretaker deems to be good, she smiles and acts enthused. The robot responds by mimicing the caretaker's happy expression, thereby inducing itself to be in a good emotional state. This mimicing behavior has the added benefit of giving the caretaker a visual feedback cue as to whether or not the intended emotive signal was transferred to the robot. Note that one could also design a mechansism where the perceived emotion on the caretaker's face (or heard in her voice) influences the robot's emotive system directly[2], bypassing the mimicing phase. This may be necessary in the case where visible expression of emotion is not observable but is present in the caretaker's voice.

The next step is to design a mechanism by which this positive emotional state becomes associated with the act and circumstance the robot just experienced. This would allow the robot's emotive state to serve as a reinforcement signal during learning episodes. In this case, the expression on the caretaker's face enables the robot to "color" or "tag" the encountered events or objects with an emotional marker that is consistent with the emotion being aroused by the caretaker's expression.

Drawing once again on ideas from emotion research, we envision this "tagging" process as an instance of emotional conditioning. LeDoux's work in the mechanisms of fear conditioning (LeDoux 1996) exemplifies these ideas and provides direct evidence on the ability of emotional systems to associate different contextual stimuli significant for the type of learning described in this paper. Following this view, we think of emotional systems as consisting a set of inputs, an appraisal mechanism, and a set of outputs. The appraisal mechanism, specific to each emotional system, is programmed by evolution to detect certain input or trigger stimuli that are relevant to the proper functioning of the robot. We will refer to these triggers as *natural releasers* or *primary emotions*. But these appraisal mechanisms also have the ability to learn about stimuli that tend to be associated with and are predictive of these natural releasers. These *learned releasers* or *secondary emotions* are in effect "marked" or "tagged" with the emotional information associated with the natural releasers. Thus, learned releasers can also unleash certain patterns of response that have been useful in dealing with situations that have been previously activated by the emotional system only through natural releasers.

This "tagging" process fits well within Damasio's *somatic marker hypothesis* (Damasio 1994), in which he argues that the activation of a covert, nonconscious biasing mechanism plays a significant role in human learning and reasoning. The main idea behind this hypothesis is that decisions that are made in circumstances similar to previous experience (whose outcome could be potentially harmful, or potentially advantageous) induce a somatic response that are used to label future outcomes that are important to us, and to signal their danger or advantage. Thus, when a negative somatic marker is linked to a particular future outcome, it serves as an alarm signal that tells us to avoid that particular course of action. However, if a positive somatic marker is linked instead, it becomes an incentive to make that particular choice.

For our purposes, the combination of these two theories has inspired a mechanism for emotional memories that generates emotional markers for significant stimuli **--** i.e., those that are present at the time the emotional systems become active due to the occurrence of a natural releaser (Velásquez 1998). This mechanism gives us the functionality required for "tagging" events and objects with the emotional information that is provided by the caretaker's reinforcement signals during social interaction with the robot. This type of emotional learning can then be used in conjunction with other types of learning mechanisms to learn new behaviors, as suggested below.

---

[2] Either way, having a mechanism by which the robot's emotional state can mirror that of the caretaker's emotive display(s) has intriguing implications for building a robot capable of empathy and compassion **--** such as the robot being in a positive emotional state when the caretaker is happy, and in a negative emotional state when the caretaker shows displeasure. As a consequence, the robot could then be motivated to learn and behave in ways that keeps the caretaker happy, or even cheers the caretaker up when unhappy.

Implications for Altricial Learning in Robotics

By combining the ideas from the previous three sections, we are working toward giving Kismet the ability to learn in an environment with social interactions similar to those afforded to a human infant. Most importantly, the environment contains a benevolent caretaker who takes an active role in helping the robot learn how it can better satisfy its drives and maintain a positive emotional state. Thus far, we have focused on one important scaffolding act which the caretaker may provide -- the ability to socially transfer positive and negative reinforcement to the robot through emotive channels of communication. The ability for a human to manipulate the robot at an "emotional" level has a couple of important implications for learning new behaviors.

Learning New Appetitive Behaviors

The first implication is that the caretaker can help the robot discover *new appetitive behaviors* -- namely, new ways to achieve existing goals, such as maintaining its drives within homeostatic bounds. For instance, one of Kismet's drives is to be stimulated by things in the environment, such as a toy. Kismet's toys tend to be brightly colored and contrast strongly with the rest of the environment. Hence the color of the toy should attract Kismet's attention because of its saliency, and Kismet should be biased to look at it simply due to low level attentional mechanisms. A detailed description of Kismet's motivation and attention systems can be found in (Breazeal(Ferrell) 1998a,b).

Here's a hypothetical situation: let us say that while Kismet's **`stimulation`** drive is outside of the homeostatic regime (causing the robot to be "bored"), Kismet happens to look at the toy simply because of its saliency. The caretaker observes Kismet's "sad" expression, which is indicative of Kismet's "bored" state, and sees Kismet looking at the toy. The caretaker might assume that Kismet is looking at the toy because it "wants" the toy and is "sad" because no one is playing with Kismet using the toy.

Although this intentional interpretation of Kismet's behavior is not true, it nonetheless gives the caretaker reason to start playing with Kismet using the toy. As a result of being stimulated by the toy, Kismet's **`stimulation`** drive is satiated, which serves as a reinforcing signal. Furthermore, while playing with Kismet, the caretaker may be using pleasing sounding vocalizations as caretakers often do when playing with infants. This can be another reinforcing signal, for it brings Kismet's emotive system into a positive state. If this scenario is repeated often enough, Kismet may learn that when the caretaker is present, Kismet's **`stimulation`** drive will be satiated simply by looking at the toy. Hence, this **`Look-At-Toy`** behavior is a new appetitive behavior that Kismet can employ to reduce its **`stimulation`** drive and place itself in a positive emotive state.

Once the robot has measures of how its motivational state (level of drives and emotions) changes with experience, a learning mechanism is required to incorporate this knowledge into behavior. Blumberg (1996) presents a reinforcement-based algorithm that implements operant conditioning which could be suitable for this type of learning, and has demonstrated its effectiveness by teaching an animated dog new behaviors.

Learning New Consummatory Behaviors

In addition to helping the robot learn new appetitive behaviors, the emotive system gives the robot a means for learning *new consummatory behaviors* with the help of the caretaker. Recall the discussion on scaffolding. For our purposes this discussion is interesting because one could postulate that when an infant is learning a *novel* behavior, there is *no a priori representation of the goal for that behavior in the infant's mind*. Hence, the infant's actions are not goal-directed initially. However, the goal may very well be represented in the *mind of the adult* who is guiding the infant to experience the desired goal state through performing that novel behavior.

The emotion based learning mechanisms discussed previously provides a way by which these social forms of scaffolding can be used by a human to *transfer* a goal representation from the caretaker's mind to the robot in the form of a new consummatory behavior. To do so, the caretaker can exploit the learning mechanics of Kismet's motivation system to place the robot in a positive emotional state. From the robot's perspective, it receives rich and ongoing reinforcement about its actions within specific contexts by witnessing the caretaker's expressive acts such as facial expressions or emotive vocalizations. By having the caretaker communicate this information socially, the robot has access to measures of progress toward the goal as well as achievement of the goal *before* that goal is explicitly represented within the robot. In addition, through her scaffolding acts, the caretaker can actively guide the robot toward the goal and allow the robot to experience the goal *before* the robot is capable of achieving that goal on its own.

By consistently and repeatedly engaging in this process, the robot's emotional memories could eventually learn to associate a positive emotional state with the surrounding context and the robot's actions. As a result, that behavior (surrounding context and robot action) is "tagged" as being worthy of pursuit in its own right, and the robot is motivated to perform that new behavior simply because it places the robot in a positive emotional state. For this reason, we consider the new behavior to be of the *consummatory* type because its activation serves to directly satisfy a basic need of the robot -- i.e. maintain a positive emotional state.

Once the robot has integrated a new consummatory behavior into its behavior system, it can use a learning mechanism similar to the reinforcement-based contingency learning of Blumberg (1996) to learn *new appetitive* behaviors for that consummatory behavior. Granted, the new consummatory behavior ultimately sub-serves one of the basic goals of the overall system -- i.e. to be in a positive emotive state. However, to an observer, the robot's behavior may appear to be goal-directed toward a *new goal* -- the robot may either perform the consummatory behavior spontaneously when circumstances permit, or it may pursue alternate appetitive behaviors in order to activate the new consummatory behavior, simply because it places the robot in a positive emotional state. After all, people are often motivated to do things because it makes them feel good, not only because they are tired, hungry, thirsty, et cetera.

The communication of positive and negative reinforcement though emotive channels gives the caretaker a way of guiding the learning process, and shaping what the robot learns by teaching it new ways to achieve existing goals, or perhaps even new goals themselves. Note that the caretaker is critical for learning novel behavior because she helps the robot *identify* and *pursue* new goals by communicating this information through emotive channels of communication. The caretaker does so by actually *training* the robot to be in a good emotional state when performing that behavior. Having an insightful, intentional caretaker guiding the learning process is a tremendous advantage for human infants because the caretaker helps the infant learn what behaviors are worth pursuing in the first place. Otherwise, how might a new-born decide this for itself without the benefit of being raised in a culture? By placing our robot "infant" in a similar environment with access to similar information channels, we wish to explore issues regarding how the robot and what it learns can benefit in similar ways.

Progress Towards Building an Altricial Robot that Learns

We want to design a robot that is biased to learn how its emotive and behavioral acts influence the caretaker in order to satisfy its own drives and to maintain a positive emotional state. Toward this end, we have endowed the robot with a motivational system that works to maintain its drives within homeostatic bounds and could motivate the robot to learn behaviors that satiate them. The motivation system should also work to maintain a positive emotional state and play an important role in helping the robot learn which behaviors to pursue or avoid to promote this state.

Further, we have provided the robot with a set of emotive expressions that are easily interpreted by a naive observer as analogues of the types of emotive expressions that human infants display. This allows the caretaker to observe the robot's emotive expressions and interpret them as communicative acts. She assumes the robot is trying to tell her which of its needs must be tended to, and she acts accordingly. She also assumes the robot's facial expressions are indicative of its emotional state and acts to educe positive expressions. This establishes the requisite routine interactions for the robot to learn how its emotive acts influence the behavior of the caretaker, which ultimately serves to satiate the robot's own drives and maintain a positive emotive state.

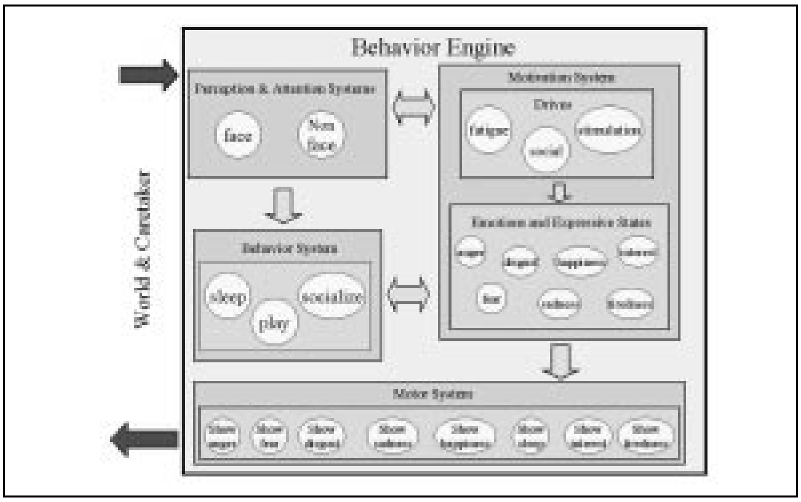

Figure 2 This figure illustrates the behavior engine implementation to date. Please see (Breazeal(Ferrell) 1998a,b) for a detailed description.

For the remainder of this paper, we describe our progress to date in this endeavor. However, due to space constraints, we only give a very brief description of the behavior engine. We refer the interested reader to (Breazeal(Ferrell) 1998a,b) for an in-depth description.

A framework and current implementation for Kismet's behavior engine is shown in figure 2. The overall system is implemented as an agent-based architecture similar to that of (Blumberg 1996, Maes 1990, Brooks 1986, Minsky 1986). The system architecture consists of five subsystems: the *perception and attention systems*, the *motivation system*, the *behavior system*, and the *motor system*. The perception system extracts salient features from the world, the motivation system maintains internal state in the form of drives and emotions, the attention system determines saliency based upon perception and motivation, the behavior system implements various types of behaviors as conceptualized by Tinbergen (1951) and Lorenz (1973), and the motor system realizes these behaviors as facial expressions and other motor skills.

The Perceptual Subsystem

From its visual input, the robot extracts two percepts, *face* and *non-face*. The *face* percept affects the *social* drive and is computed using a ratio template technique first proposed by Sinha (1994) and later adapted for this system by Scassellati (1998). The method looks for a characteristic shading pattern of human assuming a frontal viewpoint. The intensity of the *face* percept is given by the amount of visual motion of a detected face. Any other motion is attributed to a *non-face* stimulus which affects the *stimulation* drive.

The Motivational System: Drives

The robot's **drives** serve three purposes. First, they influence behavior selection by preferentially passing activation to some behaviors over others. In a similar manner, they also influence the **emotive** state of the robot. By doing so, the robot's facial expressions reflect how well the robot's drives are being maintained. This provides expressive cues that the caretaker can read to gauge the interaction, and tune it to appropriate level of intensity. In this way, the robot can work with the caretaker to establish and maintain a suitable learning environment where it is neither under stimulated nor overwhelmed. Last, the **drives** provide a learning context -- the robot learns skills that serve to satiate them.

The design of the robot's drive subsystem draws heavily from the ethological views (Lorenzm 1973), (Tinbergen 1951). Briefly, the intensity of each drive reflects the ongoing needs of the robot and the urgency for tending to them. The robot must act to maintain its drives within homeostatic bounds to promote its well being. However, unless satiated, each drive increases in intensity away from the homeostatic regieme.

Currently, Kismet has three **drives**. The activation energy of each ranges between [-*max*, +*max*], where the magnitude of the **drive** represents its intensity. In general, each **drive** is partitioned into three regimes: an *under-whelmed regime* (large positive values) where the robot is being under stimulated, an *over-whelmed regime* where the robot is being over stimulted (large negative values), and the *homeostatic regime* (desired mid-range values).

- **`Social drive`**: One **`drive`** is to be social, i.e. to be in the presence of people and to be stimulated by people. On the under-whelmed extreme the robot is **`lonely`**, i.e., it is predisposed to act in ways to get into face to face contact with people. On the over-whelmed extreme, the robot is **`asocial`**, i.e. it is predisposed to act in ways to disengage people from face to face contact. The robot tends toward the **`asocial`** end of the spectrum when a person is over-stimulating the robot. This may occur when a person is moving to much, is too close to the camera, an so on.

- **`Stimulation drive`**: Another **`drive`** is to be stimulated, where the stimulus can either be generated externally by the environment or internally through spontaneous self-play. On the under-whelmed end of this spectrum, the robot is **`bored`**. This occurs if the robot has been inactive or unstimulated over a period of time. On the over-whelmed part of the spectrum, the creature is **`distressed`**. This occurs when the robot receives more stimulation than it can effectively handle, and predisposes the robot to reduce its interaction with the enviroment, perhaps by closing its eyes, turning its head away from the stimulus, and so forth.

- **`Fatigue drive`**. This **`drive`** is unlike the others in that its purpose is to allow the robot to shut out the external world instead of trying to regulate its interaction with it. This is the time for the robot to do "internal housekeeping" without having to worry about the external world. Currently while the robot "sleeps", *all* **`drives`** return to their homeostatic regimes so that the robot is in a good motivational state when it awakens.

The Behavior Subsystem

As mentioned previously, **`drives`** cannot satiate themselves. They become satiated whenever the robot is able to evoke the corresponding *consummatory behavior*. Hence, at any point in time, the robot is motivated to engage in behaviors that maintain its **`drives`** within their homeostatic regime. For instance, if the robot lacks social stimulation, it is biased to interact with people.

Some of Kismet's consummatory behaviors require a perceptual contribution (typically provided by the person interacting with the robot) to become active. These consummatory behaviors can become active through environmental stimulation alone if it is strong enough. This has an important consequence for regulating social interaction with a person. Namely, if the nature of the interaction is too intense, the robot becomes motivated to reduce the amount of stimulation. For instance, if the caretaker is interacting with the robot too intensely, the **`social drive`** may move into the overwhelmed regime. When this occurs, the robot exhibits displeasure, which is a cue for the caretaker to back off a bit.

Currently, there are three consummatory behaviors, each responsible for satiating its affiliated drive. The activation level of each behavior can range between [0, *max*]. In general, both internal and environmental factors are used to determine whether or not they should be activated. Ideally, the behavior becomes active when the **`drive`** enters the under stimulated regime and remains active until it returns to the homeostatic regime.

- **`Socialize`** acts to move the **`social drive`** back toward the **`asocial`** end of the spectrum. It is potentiated more strongly as the **`social`** drive approaches the **`lonely`** end of the spectrum. Its activation level increases above threshold when the robot can engage in face to face interaction with a person, and it remains active for as long as this interaction is maintained. Only when active does it reduce the intensity of the drive.

- **`Play`** acts to move the **`stimulation drive`** back toward the **`confused`** end of the spectrum. It is potentiated more strongly as the **`stimulation drive`** approaches the **`bored`** end of the spectrum. The activation level increases above threshold when the robot can engage in some sort of stimulating interaction, either with the environment such as visually tracking an object, or with itself such as playing with its voice. It remains active for as long as the robot maintains the interaction, and while active it continues to move the **`drive`** toward the over-whelmed end of the spectrum.

- **`Sleep`** acts to satiate the **`fatigue drive`**. When the **`fatigue drive`** reaches a specified level, the **`sleep`** consummatory behavior turns on and remains active until the **`fatigue drive`** is restored to the homeostatic regime. When this occurs, it is released and the robot "wakes up".

This behavior also serves a special "motivation reboot" function for the robot. If the caretaker fails to act appropriately and any drive reaches an extreme, the robot is able to terminate bad interactions by going to **`sleep`**. This gives the robot a last ditch method to restore all its **`drives`** by itself.

The Motivational System -- Emotions:

For Kismet, **`emotions`** serve two functions. First, the establish the **`emotive`** expression of the robot. Second, because the **`drives`** contribute to the **`emotional`** state of the robot, which is reflected by its facial expression, the **`emotions`** play an important role in communicating the state of the robot's "needs" to the caretaker and the urgency for tending to them.

The organization and operation of the emotion subsystem is strongly inspired by various theories of emotions in humans (Ekman & Davidson 1994), (Izard 1990), and most closely resembles the framework presented in (Velásquez 1996). As such, the robot has several **`emotion`** processes. Although they are quite different from emotions in humans, they are designed to be rough analogs --- especially with respect to the accompanying facial expressions. Kismet's emotional processes represent different families of related affective responses. Each member of an emotion family shares certain mechanisms and characteristics, including similarities in antecedent events, expression, likely behavioral response, and physiological patterns. These characteristics differ between emotion families, distinguishing one from another.

In the literature on human emotions (Ekman & Davidson 1994), there are four factors that can elicit an emotion: neurochemical, sensorimotor, motivational, and cognitive. These releasers can be either hard-wired (called a *natural releaser*), or acquired through experience (called a *learned releaser*). Currently, all releasers from the first three categories are natural releasers and all cognitive ones are learned releasers.

- Neural: Includes the effects neuroactive agents (e.g., neurotransmitters, and brain temperature) that can lead to emotion and which can be mediated by hormones, sleep, diet, and environmental conditions.
- Sensorimotor: Includes sensorimotor processes, such as facial expressions, body posture, and muscle action potentials, that not only regulate ongoing emotion experiences but can also elicit emotion.
- Motivational: Includes all motivations that lead to emotion (e.g., innate responses to foul odors or tastes producing disgust, pain or aversive stimulation causing anger, and emotions like sadness eliciting others such as anger).
- Cognitive: Includes all type of cognitions that activate emotion, such as appraisal of events, comparisons, attributions, beliefs and desires, and memories.

Currently, all releasers from the first three categories are *Natural Releasers* and all cognitive ones are *Learned Releasers*.

The activation of an **`emotion`** is a nonlinear function of its input, which includes both the input from all associated releasers and the excitatory (positive) and inhibitory (negative) input from other **`emotion`**. This is summarized in Equation (1):

$$A_{it} = f\left(\Psi(A_{it-1}) + \sum_k R_{ki} \cdot W_{ki} + \sum_l \mu_{li} \cdot A_{lt}\right) \quad (1)$$

Where A_{it} is the activation of `emotion` i at time t; A_{it-1} is its activation at the previous time step; $\Psi()$ is the function that controls the temporal decay of the activation of `emotion` i; R_{ki} is the value of Releaser k, and W_{ki} is its associated weight, where k ranges over the set of releasers for `emotion` i; μ_{le} is the strength of the excitatory (positive) or inhibitory (negative) input from `emotion` l, where A_{lt} is its activation value at time t; and *f* is a limiting function such as the standard ramp and logistic (sigmoid) functions.

Numerically, the activation level of each **`emotion`** can range between [0, *max*] where *max* is an integer value determined empirically. Although the **`emotions`** are always active, their intensity must exceed a threshold level before they are expressed externally. When this occurs, the corresponding facial expression reflects the level of activation of the **`emotion`**. Once an **`emotion`** rises above its activation threshold, it decays over time back toward the base line level (unless it continues to receive inputs from other processes or events). Hence, unlike **drives**, **emotions** have an intense expression followed by a fleeing nature.

Fast Primary Emotions:

This model of an emotion process allows for the distinction between different affective phenomena. For instance, *primary emotions* are modeled as the activation (via natural releasers) of one particular emotion process such as **disgust** or **fear**. These primary emotions play an essential role in the preparation of appropriate emotional responses that are important for social interaction.

So far, there are eight primary emotions implemented on Kismet: **anger**, **disgust**, **fear**, **happiness**, and **sadness** are analogs of the primary emotions in humans. The last three **emotions** are somewhat controversal in classification, but they play in an important role in learning and social interaction between caretaker and infant so they are included in the system: **suprise**, **interest**, **excitement**.

Emotional Memories and Secondary Emotions:

Emotion processes also have the capacity of acquiring Learned Releasers, which, as we previously mentioned, correspond to stimuli that tend to be associated with and predictive of Natural Releasers. Secondary Emotions correspond to the activation of **emotions** via Learned Releasers. These secondary emotions occur only after we begin experiencing feelings and start making orderly associations between objects and situations, and primary emotions. Secondary emotions also support social exchange, but more important, they provide the foundation for different kinds of learning, including the type of social learning discussed above.

These secondary emotions have been modeled with an associative network comparable to Minsky's K-lines (Minsky 1986), in which primary emotions are connected to the specific stimuli (e.g., executed behavior, objects or agents) that have elicited them during the robot's interaction with the world. During emotional learning, connections within this network are changed according to a modified Hebbian rule. The sum of the weights of all incoming learned connections to an **emotion** process are kept constant, through multiplicative normalization based on the existing excitatory and inhibitory connections between **emotions**.

The Motor Subsystem

For each **emotion** there is a recognizable accompanying facial expression. These are implemented in the motor system among various motor transducer processes. The low level face motor primitives control the position and velocity of each degree of freedom. At the next level, the motor skill processes implement coordinated control of the facial features such as wiggling the ears or eyebrows independently (i.e. those motions typically coordinated when performing a facial expression). Next are the face expression processes which direct all facial features to show a particular expression -- their intensity (speed and displacement of facial features) can vary depending on the intensity of the emotion evoking the expression. Blended expressions are computed by taking a weighted average of the facial configurations corresponding to each evoked emotion.

Experiments and Results

In this section we report on a variety of that experiments we have carried out to test the ideas presented in this paper. A variety of social interaction experiments between Kismet and a human caretaker were performed. These experiments focus on generating and maintaining a social interaction of suitable intensity for the robot grounded in emotive cues.

Not all experiments have been carried out on Kismet. In particular, those involving the formation of emotional memories have used another robot in our lab, called Yuppy, as a testbed. However, the architecture of both robots and their underlying mechanisms are very similar despite differences in implementation details. Porting these
mechansisms to the Kismet platform is discussed in the Future Work section of this paper.

Kismet: Regulating Social Interaction with a Human Using Emotive Displays

A number of experiments have been performed with Kismet using the behavior engine shown in figure 2 and discussed previously. Each experiment involved a human interacting with the robot either through direct face-to-face interaction, by waving a hand at the robot, or using a toy to play with the robot. Due to space constraints, we report on only one set of these experiments, previously unpublished. The interested reader can refer to Breazeal(Ferrell) (1998a,b) for a report of different experiments.

Data was recorded on-line in real-time during interactions between a human and the robot. Figures 4 and 5 plot the activation levels of the appropriate **`emotions`**, **`drives`**, behaviors, and percepts. **`Emotions`** are always plotted together with activation levels ranging from 0 to 2000. Percepts, behaviors, and **`drives`** are often plotted together. Percepts and behaviors have activation levels that also range from 0 to 2000, with higher values indicating stronger stimuli or higher potentiation respectively. **`Drives`** have activations ranging from -2000 (the over-whelmed extreme) to 2000 (the under-whelmed extreme).

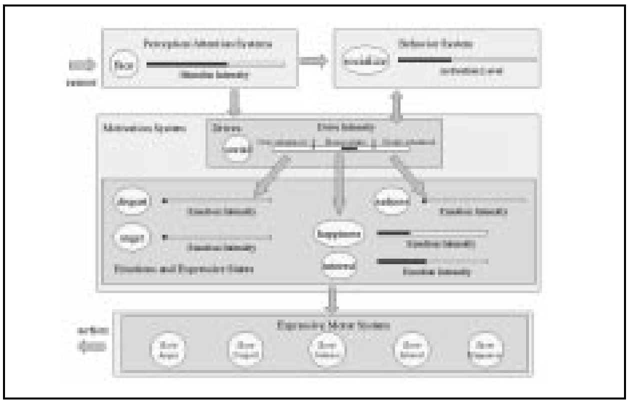

Figure 3. Diagram of the social circuit

In this set of experiments, the caretaker engages the robot in playful exchange with a small, black and white plush cow. Because the perceptual system classifies the cow's face as a *face* percept, these interactions are regulated by the **`social circuit`**, as shown in figure 3. The motion generated by the cow gives a rating of the stimulus intensity. During the exchange, the Kismet's facial expressions reflect its ongoing motivational state (i.e. it's mood) and provides the human with visual cues as to

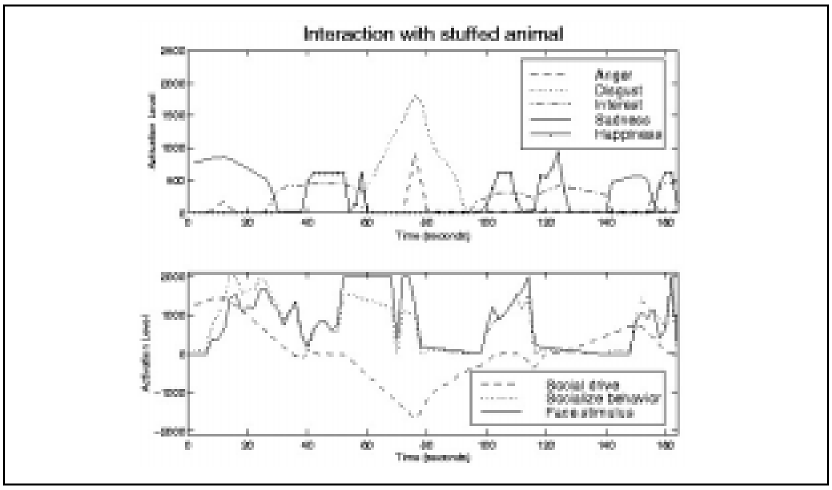

how to modify the interaction to keep the robot's drives within homeostatic ranges.

Figure 4. Playing with the robot using a plush cow toy

Figure 4 shows a plot of how Kismet's emotions, social drive, and behavior respond to varying the intensity of the *face* stimulus of a plush cow toy. The run begins with the **social drive** within the **lonely** regime and the robot looking "sad". At *t=5* the experimenter shows the robot the cow's face and moves the cow in small gentle motions. This corresponds to a stimulus of acceptable intensity level which restores the **drive** to the homeostatic regime. As a result the robot appears "interested" and "happy". From *50 >= t >= 78* the experimenter begins swinging the cow quickly in front of the robot's face. Because the stimulus is too intense, the **drive** moves into the **asocial** regime and the robot expression of "disgust" intensifies until eventually blended with "anger" as well. At *t=78* the experimenter removes the cow from the robot's visual field and allows the **drive** to return to the homeostatic regime. From *98 >= t >= 118* the cow's face is shown to the robot again which maintains the **drive** within the homeostatic regime and the robot displays "interest" and "happiness". From *118 >= t >= 145* the cow's backside is shown to the robot. The lack of a face stimulus causes the **social drive** to return to the **lonely** regime, but at *t >= 145* the cow is turned to face the robot and the **drive** is restored to the homeostatic regime until the conclusion of the run. The run ends with the robot in a "happy" and "interested" state.

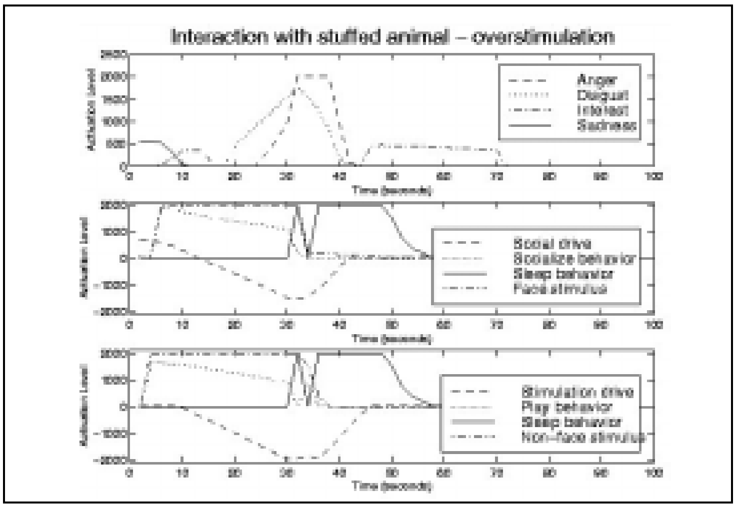

Figure 5. Overstimulating the robot with the plush cow toy

Infants fall into a disturbed sleep when put into an extremely anxious state for a prolonged time. Analogously for Kismet, if the interaction is over-whelming for long periods of time, Kismet will first show increasing signs of "disgust", eventually blending with increasingly intense signs of "anger", as the **social drive** continues to move toward the over-whelmed end of the spectrum. Figure 5 shows one example of this effect. When no relief is encountered and the **drive** hits its outer limit (t=30), the robot goes into an emergency sleep mode. As discussed previously, sleeping serves as a sort of "motivational reboot" for the robot by restoring all **drives** to their homeostatic ranges. Hence, upon "awakening", the robot is in a balanced, "interested" state.

These experiments demonstrate how the robot's emotive expressions are used to regulate the nature and intensity of playful social exchange, and how the nature of the interaction influences the robot's behavior. The result is an ongoing "dance" between robot and human aimed at maintaining the robot's drives within homeostatic bounds. If the robot and human are good partners, the robot remains "interested" and/or "happy" most of the time. These expressions indicate that the interaction is of appropriate intensity for learning.

Yuppy: Emotional Behavior and Emotional Conditioning

An extension to the Cathexis framework (Velásquez 1996, Velásquez 1997), which includes the mechanisms for emotional learning described above, has been used to develop and control an emotional pet robot that exhibits emotional behavior and is able to acquire secondary emotions as it operates in the real world. The robot, called Yuppy, is shown in Figure 6.

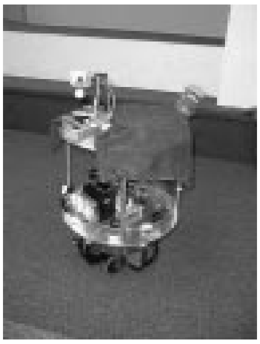

Figure 6 Yuppy, an Emotional Pet Robot

The Sensory System. Yuppy has different sensors, including two color CCD cameras (currently using only one as part of its vision system), an active stereo audio system composed of two microphones mounted on Yuppy's ears, IR sensors for obstacle avoidance, an air pressure sensor used to model simple touch perception in the form of painful and pleasurable stimuli, a pyro sensor aligned with the top camera used to detect changes in temperature due to the presence of people, and a simple proprioception system.

The Drive System. Yuppy's Drive System is composed of four different drives: **`Recharging-Regulation`**, **`Temperature-Regulation`**, **`Fatigue`**, and **`Curiosity`**, each of which controls internal variables representing the agent's **`battery`**, **`temperature`**, **`energy`**, and **`interest`** levels, respectively.

The Emotion Generation System. Yuppy's Emotion System includes **`emotions`** with *natural releasers* for the set of basic emotion families including **`Anger`**, **`Fear`**, **`Disgust`**, **`Happiness`**, **`Sadness`**, and **`Surprise`**. Some of the most representative releasers and their associated **`emotions`** include:

- *Interactions with Drive Systems*: Unsatisfied drives produce **`Distress`** and **`Anger`**; **`Distress`** is also produced, to a lesser extent, when drives are oversatisfied; and satiation of drives releases `Happiness`.
- *Interactions with the environment*: In general, all pink-reddish objects release **`Happiness`** to some extent, although pink Styrofoam "bones" elicit the most; Yellow objects found in the environment release **`Disgust`**; Darkness releases activity in the **`Fear emotion`**; and loud noises release **`Surprise`**.
- *Interactions with People*: When people interact with the robot, two of the several possible interactions include petting and disciplining the robot. These actions will generate representations of pleasure and pain, respectively. Pleasure releases **`Happiness`**, and pain releases **`Distress`** and **`Anger`**.

The Behavior System. The robot's Behavior System is composed of a distributed network of approximately nineteen different self-interested behaviors, directed in most part toward satisfying its needs and interacting with people. Examples of such behaviors include **`Search-For-Bone`**, **`Recharge-Battery`**, **`Wander`**, **`Avoid-Obstacle`**, **`Approach-Person`**, and **`Express-Emotion`**.

The user interacts with the robot providing it with different stimuli (e.g., approaching the robot, petting the robot, and showing the bone).

Yuppy exhibits emotional behaviors under different circumstances. For instance, when the robot's **`Curiosity drive`** is high, Yuppy wanders around, looking for the pink bone. When it encounters one, the activity of the **`Happiness emotion`** increases and specific behaviors, such as **`Wag-Tail`** and **`Approach-Bone`** become active. On the other hand, as time passes by without finding any bone, the activity of its **`Distress emotion`** rises and appropriate responses, such as **`Droop-Tail`**, get executed. Similarly, while wandering around, it may encounter dark places which will elicit fearful responses in which it backs up and changes direction.

Besides regulating action-selection and generating emotional behaviors through primary emotions, the robot learns secondary emotions which are stored as new or modified cognitive elicitors based on the associative network model described before. In one instance, the **`Fear emotion`** acquires a new releaser for loud sounds. In this classical scenario of *fear conditioning*, a natural releaser (Pain generated as a person disciplines Yuppy) generates a fearful response (**`Cower`** behavior). The sound stimulus by itself, however, does not produce any activation in the **`Fear emotion`**, thus the **`Cower`** behavior does not become active either. If both stimuli are presented simultaneously, the **`Fear`**

emotion creates a new cognitive releaser for the sound stimulus. After only one trial, the newly formed releaser for the loud sound is capable of producing some activation of the **Fear emotion**. After several more trials, the connection between the sound releaser and the **Fear emotion** is strong enough to produce activation of the cowering behavior, and thus an emotional memory is formed.

In another scenario, after locating a pink bone (natural releaser) and approaching it, the robot interacts with the person carrying the bone. Depending on these interactions (e.g., the person pet or hit the robot), the robot is able to create positive or negative emotional memories with respect to people, and the selection of future behaviors such as approaching or avoiding them are influenced.

These results showed that emotional conditioning was possible under the proposed model, and that it could further be used to associate and predict stimuli that accompanied natural releasers. These results are interesting and promising, since, as we have described before, they provide the foundations for several different kinds learning systems, including that of social learning.

Future Work

Future work includes integrating the mechanisms of learning emotional memories into Kismet's existing behavior engine, and establishing the appropriate connections so that the caretaker's emotive signals are mimicked by the robot and influence its emotive state. Work in perceiving the caretaker's emotive cues both visually and through auditory channels is underway. Once this is done, experiments will be performed to see if Kismet can acquire emotional memories. Once successful, these emotionally labeled states can be combined with reinforcement based learning methods (such as that of Blumberg) to explore socially teaching the robot new appetitive behaviors. Ultimately we would like to explore teaching the robot new goals in the form of new consummatory behaviors.

Acknowledgments

Support for this research was provided by a MURI grant under the Office of Naval Research contract N00014--95--1--0600. The visual processing routines described where graciously ported to Kismet by Brian Scassellati. Support for research on Yuppy was provided by Yamaha Motor Corporation.

References

Blumberg, B. (1996), Old Tricks, New Dogs: Ethology and Interactive Creatures, PhD thesis, MIT.

Breazeal (Ferrell), C. (1998a), A Motivational System for Regulating Human-Robot Interaction, *in* "Proceedings of AAAI-98".

Breazeal (Ferrell), C. (1998b), Early Experiments using Motivations to Regulate Human-Robot Interaction, *in* "Proceedings of AAAI Fall Symposium 98".

Brooks, R. (1986), "A robust layered control system for a mobile robot", IEEE Journal of Robotics and Automation RA-2, 253--262.

Brooks, R., et al. (1998), Alternative Essences of Intelligence, in "Proceedings of AAAI-98".

Bullowa, M. (1979), Before Speech: The Beginning of Interpersonal Communicaion, Cambridge University Press, Cambridge, London.

Damasio, A. (1994), Descartes Error: Emotion, Reason, and the Human Brain, G.P. Putnam's Sons, New York, NY.

Ekman, P. & Davidson, R. (1994), *The Nature of Emotion: Fundamental Questions*, Oxford University Press, New York.

Halliday, M. (1975), Learning How to Mean: Explorations in the Development of Language, Elsevier, New York, NY.

Hendriks-Jansen, H. (1996), Catching Ourselves in the Act, MIT Press, Cambridge, MA.

Izard, C. (1990), "Facial Expressions and the Regulation of Emotions", Journal of Personality and Social Psychology.

LeDoux, J. (1996), The Emotional Brain: The Mysterious Underpinnings of Emotional Life, Simon and Schuster, New York, NY.

Lorenz, K. (1973), Foundations of Ethology, Springer-Verlag, New York, NY.

Maes, P. (1990), "Learning Behavior Networks from Experience", ECAL90.

McFarland, D. & Bosser, T. (1993), Intelligent Behavior in Animals and Robots, MIT Press, Cambridge, MA.

Meltzoff, A. & Moore, M. (1977), "Imitation of facial and manual gestures by human neonates", Science.

Minsky, M. (1986), The Society of Mind, Simon and Schuster.

Mowrer, O. (1960), Learning Theory and Behavior, John Wiley and Sons, New York, NY.

Scassellati, B. (1998), Finding Eyes and faces with a Foveated Vision System, *in* 'Proceedings of AAAI-98', AAAI Press.

Sinha, P. (1994), 'Object recognition via Image Invariants: A Case Study', *Investigative Ophthamology and Visual Science* **35**, 1735-1740.

Tinbergen, N. (1951), *The Study of Instinct*, Oxford University Press, New York.

Tomkins, S. (1962), Affect, imagery, consciousness, Springer.

Trevarthen, C. (1979), Communication and cooperation in early infancy: a description of primary intersubjectivity, *in* M.Bullowa, ed., "Before Speech", Cambridge University Press, pp. 321--348.

Velasquez, J. (1996), Cathexis: A Computational Model for the Generation of Emotions and their Influence in the Behavior of Autonomous Agents, Master's thesis, MIT.

Velásquez, J. (1997). Modeling Emotions and Other Motivations in Synthetic Agents, *in* "Proceedings of AAAI-97".

Velásquez, J. (1998), When Robots Weep: Emotional Memories and Decision-Making, *in* "Proceedings of AAAI-98".

Wood, D., Bruner, J.S. & Ross, G. (1976), The role of tutoring in problem-solving, Journal of Child Psychology and Psychiatry 17, 89--100.