# A Habit System for an Interactive Robot

**Kai-yuh Hsiao and Deb Roy**
Cognitive Machines
MIT Media Lab
Email: eepness@mit.edu and dkroy@media.mit.edu

### Abstract

Human behavior is a mix of automatic perceptually-triggered habits and intentional goal-directed actions. We have implemented a habit system that automatically executes actions based on internal and external context. The result of the habit system is a homeostatic mechanism with implicit goals that balances the physical needs of a physical robot with other factors such as interaction with human partners. By acting automatically on salient context elements, the robot intersperses the use of its conversational interface with self-preservation and curiosity drives. Ultimately, having both a goal and a habit system governing robot behavior will allow a very rich platform for embodied, grounded semantics. We present details of the model and implementation, and discuss insights into the embodied nature of semantics that arises from this work.

## Habitual and Goal-Directed Actions

Human behavior is mostly a matter of habit. On most levels, actions taken by humans are carried out automatically, as soon as the proper perceptual context arises. Intentional, goal-directed actions comprise only a part of human behavior, and even very complex intentional actions can be rehearsed to the point of being ingrained into habit. The dividing line between habit and intentional action is furthermore very slippery and difficult to introspect.

However, the difference between habit and intentional action is relatively clear. Intentional actions are taken in a goal-directed manner, so that actions can be re-evaluated rapidly amid a changing environment. However, if an intentional action is repeated in the same perceptual context, it begins the gradual slip into habitual action. Underlying this habituation is the assumption that the environment is static enough for repeated actions to have merit, but dynamic enough that intentional actions are necessary for some aspects of behavior.

Because of the split between habitual and intentional action, there is also a split in the anticipatory nature of human action. Intentional acts are selected and performed in explicit anticipation of a result. However, habitual acts are selected solely because of the current perceptual context, and are thus essentially reactive, but it is still possible to view them as being performed in implicit, previously-trained anticipation of a result.



### Building a Habit System

We are in the process of building a robot system with both habitual and intentional action. Habits are performed based on context, and intentional actions are performed in explicit service of a goal. As with the interactions between the prefrontal cortex and the basal ganglia of the primate brain (Miller & Wallis 2004), our system's behavioral control stems from the interaction between a goal system and a habit system. The habit system monitors internal and external percepts for contexts that trigger habits, and the goal system adds coherence to the habit system. It does this by using explicit representations of the habit system's implicit goals to inhibit irrelevant habits, thus adding coherence to habit execution.

Towards this end, we have begun by constructing a simple habit system, which acts according to internal and external context. The *internal context* of our system is a set of factors such as maintaining the mental model. The *external context* of our system includes factors coming from the environment, such as motor heat levels, proximity to surfaces, and utterances from humans. These contextual elements compete for the attention of the system, and the system automatically executes the habit associated with the winning element.

Not coincidentally, the contexts and the habits are designed to act homeostatically, to maintain the "health" of the system; one could imagine initially acquiring such habits based on repeated rewards of intentional actions. Thus, the habits include taking actions to reduce motor heat, avoid collisions with surfaces, and interact coherently with humans. As a habit system, though, these actions are triggered by the motivating context element, in pursuit of implicit goals, as opposed to being goal-directed.

The purpose of building this system is ultimately to create a conversational assistive robot capable of learning about its environment and interspersing its own physical and mental needs with the desires of the interacting human. An example of a physical need is keeping motor heat low; an example of a mental need is looking around to keep the mental model up to date. All habit-triggering context elements are viewed as *motivations* of the system as a whole (we use the term *drive* interchangeably), because of the way they compete for

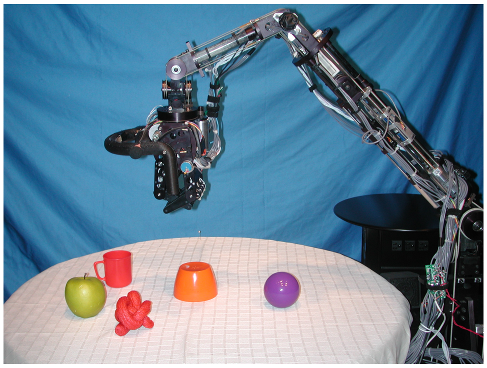

Figure 1: Ripley is a seven-degree-of-freedom robot with gripper and cameras integrated into the robot's "head," shown here looking at objects in its tabletop domain.

attention to motivate the system to execute the related habits.

The basis for our habit system is essentially similar to action selection systems previously studied in the animats and artificial agents community. In works such as (Blumberg 1994), (Minsky 1986), (Tyrrell 1993), and numerous others, it is suggested that no single behavior should consistently dominate over other behaviors. Instead, a number of behaviors compete for control of the system depending on the urgency of their underlying motivations.

The primary contributions of our habit system are 1) introducing human interaction elements with grounded language directly into an action selection system, and 2) to show how goal-directed inhibition could conceivably fit atop the system to introduce explicit goal pursuit in the future.

The remainder of this paper discusses the various components that have gone into our system so far, and concludes with a discussion of interesting properties and future directions.

## Ripley: An Interactive Manipulator Robot

Ripley is a robot arm designed for human interaction (see Figure 1) (Hsiao, Mavridis, & Roy 2003). It has seven degrees of freedom, starting with five series-elastic actuators (Pratt, Krupp, & Morse 2002), and ending with two standard motors for rotating and opening the gripper claw.

Along with force and position sensing in each DOF, Ripley is also instrumented with cameras mounted on the gripper for visual input. Motor heat is modeled by software processes that continuously integrate forces exerted at each joint, with appropriate decay during periods of motor shutdown. A microphone is worn by the robot's human communication partner so that Ripley can hear the human's speech during interactions. The audio stream drives the interaction between Ripley and the human partner (see "Verbal Interaction," below).

Ripley operates in a tabletop domain. Typical actions in its domain include viewing, grasping, and lifting objects that are on the table. By picking up objects, Ripley can then place them elsewhere or offer them to the human partner (such as in Figure 2). In addition to seeing the color, position, and size of objects, Ripley can also gauge the weight of an object while lifting it, by using its force sensitivity. Other actions include looking upwards in order to locate the human partner using a face detector (Viola & Jones 2001) to facilitate its interaction with the human partner.

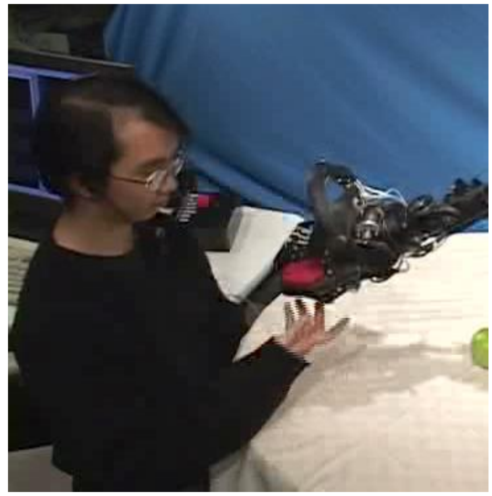

Figure 2: In response to the verbal request, "Hand me the red one on my left," Ripley hands the appropriate object to its human partner.

## Ripley's Mental Model

Ripley's ability to interact with its environment is greatly facilitated by its object-tracking *mental model* (Roy, Hsiao, & Mavridis 2004). By maintaining a persistent representation of each encountered object, as well as the human partner, in the mental model, Ripley has a sense of object permanence, retaining knowledge of objects even when not directly perceiving them. This is especially important because Ripley's cameras are mounted on its end effector, so its view of the workspace is limited and constantly shifting as it performs actions. Further interactions with an object, such as grasping, weighing, or observing from a different viewpoint, can also supplement Ripley's knowledge of the object.

The mental model is implemented as an internal three-dimensional simulation of the robot's environment. Object locations and properties are estimated from the visual input, and the human's position is found by running a face detector (Viola & Jones 2001) on the visual input. See Figure 3 for a graphical visualization of the mental model.

The end result is a representation of objects in the robot's environment, along with a representation of the positions of the human partner and the robot itself. This sense of object permanence allows the robot to respond meaningfully to requests sent down from the parser system, such as "Pick up the blue one," regardless of the current orientation of the cameras. The robot can also view the scene from the human's perspective, to process requests such as "Hand me the one on my left."

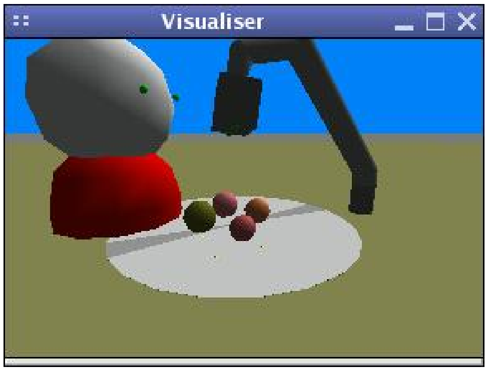


Figure 3: Ripley's mental model tracks the position of the robot, the human interaction partner (represented by the snowman-like stack on the left), and all the objects that have come into Ripley's limited visual perspective.

### Verbal Interaction

Ripley is capable of carrying out actions specified by its partner. It can also respond to simple questions about the objects. Ripley deals with verbal interaction by parsing the speech, finding word referents within its current mental model, and either executing the action, responding, or asking for more information (Roy, Hsiao, & Mavridis 2004).

We use Sun's Sphinx-4 speech recognizer (Carnegie Mellon University, Sun Microsystems Laboratories, Mitsubishi Electric Research Laboratories 2004), which passes its output to our group's Bishop system (Gorniak & Roy 2004). Bishop takes the sequence of words in an utterance, uses a chart parser to determine the parts of speech, and then uses compositionality rules to establish potential referents for each phrase within the current mental model.

Ripley's vocabulary includes names for properties ("blue," "heavy"), spatial relations ("on the left," "close to"), verbs ("pick up," "put down," "hand me"), as well as a set of sentence templates for requesting clarifications ("Do you mean this one?"). The denotative meaning (i.e., what the words refer to) of these words comes from their associations with low level sensory-motor schemas (Roy in press). The connotative meaning (i.e., how Ripley "feels" about the concepts underlying the words) emerges from the interaction between Ripley's verbal response drive and other non-verbal drives. We return to this topic later.

After extracting potential referents and either an action or a question, the system determines whether there is a single course of action to take. If there is, the robot then carries out the action, or responds with an appropriate statement, such as "The blue object is on the left." If no unique referent is provided for an action requiring a unique referent, the robot responds with a question to resolve the ambiguity, such as "Do you mean this one, or this one?" spoken while pointing appropriately. The human can then give a more specific request, such as "Pick up the blue one on the left."

This small amount of dialogue complexity is simple but effective for responding to basic verbal requests. Furthermore, because the interaction system leverages the information in the mental model, the robot is capable of fluently responding to the partner's utterances based on accumulated information about the scene, not just its current perspective.

## Drives / Motivations

We view the context elements that trigger habit execution as motivations or drives of the system, because a human whose perceptual context includes being tired is "motivated" to rest, just as a human whose perceptual context includes curiosity about a new object is "motivated" to examine the object. This sense of motivation has both a habitual and an intentional component to it; a sketch of the intentional component will be given in the discussion section.

Thus, each motivation of the system is modeled as an element of the context of the system. Each of these context elements are constantly in competition based on salience to capture attention, take control, and execute their associated habits. This has the effect of interspersing the various habits of the system according to current priority (salience metrics are designed to constitute a sort of priority), which essentially leads to a homeostatic mechanism for the robot.

Motivations compete based on their current salience values, which vary according to a hand-coded set of parameters. The salience value for each motivation ranges from 0.0 to 1.0. We designed all the drive values to approach 1.0 asymptotically, so even if a value is close to 1.0 (and is thus very high-salience), it is still possible for it to be exceeded. Using asymptotic values instead of unbounded linear values simplifies the graphical representation of the saliences, better reflects the biological analogue of neural activity, and makes it easier to program salience values by constraining the numerical scale on which they meaningfully compete.

The major motivations and their effects, which constitute the implicit goals of the associated habits, are depicted in Figure 4. We also explain how these motivations stem from two basic categories: "self-preserving" and "human-assisting." A diagram showing the relations between motivations and actions is given in Figure 5. We will use the terms "drive" and "motivation" interchangeably.

### Avoiding Fatigue (Keeping Motor Heat Low)

We consider motor heat to be analogous to muscle fatigue in animals, and high motor heat is thus viewed as a sort of discomfort or pain to the robot. The actual heat value is monitored by summing and decaying on the exerted force at each joint. Specifically,

$$h_n(t) = (h_n(t-1) + f_n(t))e^{\frac{-\Delta_t}{\tau}}$$

where $h_n(t)$ is the computed heat level for motor $n$ at timestep $t$, $f_n(t)$ is the force being exerted at motor $n$, $\Delta_t$ is the actual time elapsed (in seconds) since the last timestep, and $\tau$ is a time constant selected to approximate the cooling characteristics of our motors. Under normal conditions, we use a time constant of 5 minutes (300 seconds) in modeling heat dissipation.

The seven degrees of freedom in Ripley are grouped into three groups – two degrees at the base, one at the elbow, and four controlling the orientation and force of the gripper, with

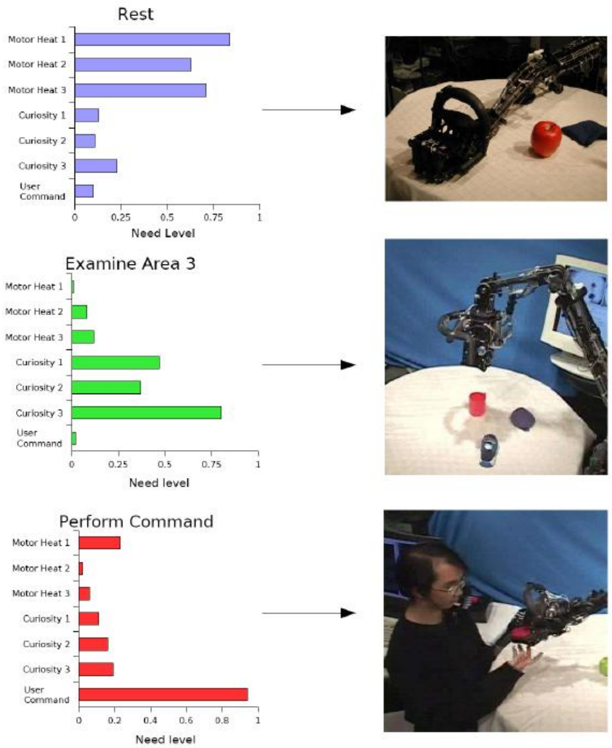


Figure 4: Ripley has three basic top-level motivations. High motor heat levels motivate Ripley to rest, high curiosity levels motivate Ripley to examine specific parts of its environment, and when interaction with the human indicates an action to be taken, Ripley is compelled to perform the action.

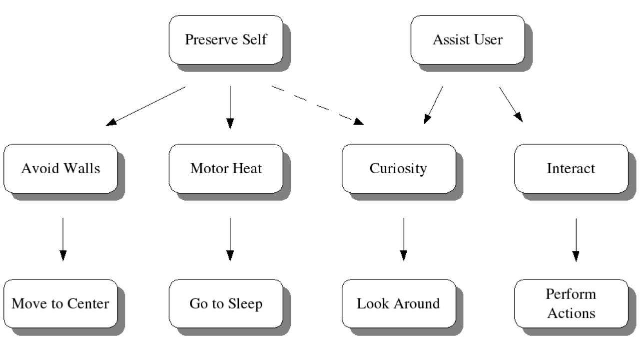


Figure 5: All motivations and resulting actions stem from the two basic goals that the robot should preserve its self and assist its partner. The connection from "Curiosity" to "Preserve Self" is shown with a dotted line because our system only uses curiosity to aid actions stemming from human interaction, but in principle it could also be used to track objects to avoid collisions, which would be self-preserving.

each group containing identical types of motors – so motor heat is summarized by tracking the maximum heat level in each motor group over time. Thus, the system watches three heat levels, $H_1(t) = \max(h_1(t), h_2(t))$, $H_2(t) = h_3(t)$, and $H_3(t) = \max(h_4(t), h_5(t), h_6(t), h_7(t))$. Each of these three monitored heat levels is treated as a single motivating drive that vies for control of the system. The salience with which each heat level competes for attention/control is a value in $(0, 1)$ scaled according to a sigmoid curve,

$$P_H = 1/[1 + e^{\frac{-4(H_n(t)-(a+b)/2)}{a-b}}]$$

where $a$ is the lowest value of $H_n(t)$ at which that motor type has been judged (by personally monitoring $H_n$ and physically touching the motors) to be dangerously hot, and $b$ is the highest value of $H_n(t)$ at which that motor type is comfortably cool.

The sigmoid curve compels each heat drive salience to a value between 0 and 1, asymptotically approaching 1.0 as heat increases (in line with the asymptotic behavior mentioned above), but with most of the range between 0 and 1 used for the motor's typical operating heat levels. When one of the heat motivations is granted control, it compels the robot to move to a safe position and turn off its motors – a preprogrammed action that we refer to as "going to sleep".

While a heat reduction drive has control and the robot is sleeping, the other drives continue to compete for control based on their current saliences. With the robot asleep, the motor heat levels will slowly drop according to the decay rates of the motor heat model, which decreases the salience of the heat levels. This in turn increases the chances that another drive will seize attention. There is also the possibility that another drive will increase in salience to exceed the salience of the heat level.

Motivations such as motor heat level fall into the more general category of "self-preserving" drives. Keeping motor heat levels generally low does very little to assist the human in any immediate sense, but rather it is focused more on keeping the robot's motors from taking damage, thus prolonging the usable life span of the robot (we previously had occasional motor burnouts, a trend which has subsided since implementing the heat monitoring). Naturally, it is also in the long-term interest of the human for the robot to require less maintenance, so in a sense it is also serving the person, but in an immediate way it is more clearly a self-preserving tendency.

## Curiosity (Keeping the Mental Model Up-to-date)

In our system, curiosity refers to a drive that causes the robot to look around at various areas of the table and its surrounding environment. The drive to look around is part of the controller's design because of its importance to the mental model – looking in various directions allows the mental model to keep its knowledge of nearby objects, including the location of the human, up to date. Having an up-to-date mental model is an anticipatory action that enables the robot to respond more quickly to requests from the human partner, if and when they are issued, by eliminating the need to take a second look around the table to find referents for an utterance.

As currently implemented, each direction of observation has been assigned its own drive level, and the salience of each of these levels is programmed to rise asymptotically over time towards an assigned maximum. When the robot has not looked in a particular direction in some amount of time, that drive will tend to grab the attention of the robot and move it to look in the desired direction, at which point the drive is reset to a low level, allowing other motivations to grab attention. Also, if in the course of performing other actions, such as carrying out verbal instructions or moving to its sleeping position, the robot happens to be able to simultaneously look and update that portion of its mental model, the corresponding curiosity motivation will also reset itself. This allows the curiosity drives to ensure the robot is only driven to look at areas it has actually not seen recently.

Our current implementation has six zones for the robot to observe. Three of them are areas of the table surface, which are (from the robot's perspective) the top, the bottom-left, and the bottom-right of the table. The other three zones involve looking up to where the human's face might be. These involve the robot looking straight forward, to its left, and to its right. The specific drive levels are modelled according to

$$C_n(t) = 1 - e^{-\frac{t-t_0}{\tau}}$$

where $C_n(t)$ is the salience level of the $n$th curiosity drive, $t$ is the current time in seconds, $t_0$ is the last time drive $n$ was satisfied (i.e., the robot was looking in the appropriate direction), and $\tau$ is a time constant that causes the drive to increase at a particular rate. We typically assign $\tau$ to be 120 seconds for the table-observing zones, and 240 seconds for the face-observing zones. Once again, the saliences approach 1.0 asymptotically.

We opted to add the motivation to look around because it improves the response of the system to user interactions. In this sense, our system's curiosity is specifically an improvement designed to assist the human partner. However, curious behavior in general can also contribute to self-preservation, by keeping track of elements in the environment that can cause harm to the robot.

### Speech Interpretation

Ripley's third basic behavior is to use its above-mentioned systems to carry on spoken interactions. When the human partner produces an utterance, the drive to respond jumps from zero to a fixed value (we have arbitrarily chosen 0.8 – adjusting this value makes for a more or less "obedient" personality), and in doing so has a high salience for getting the attention of the system. If it is successful at taking control, it then increases its salience the rest of the way to the maximum (1.0) and subsequently carries out the responses, question-asking, and action execution as described above. We use a 1.0 salience level because our interaction system is not necessarily robust to interruption, and this allows all interaction-related actions to complete before returning control. If the interaction drive is unsuccessful at taking this control, it simply responds with a verbal "Ask again later," and drops its value back to zero so whichever urgent behavior is already in control can continue.

Naturally, this behavior falls squarely under the "human-assisting" category. Its fixed salience values are chosen so that interactions are likely to complete, except when the value of another drive, such as motor heat, is already extremely high, in which case pursuing self-preservation makes more sense from a long-term perspective.

### Other Motivations

Two other motivations have been added to the system, but their habits come up much less frequently in the normal usage of the system. First, a wall-avoiding drive has been added, in which the robot keeps track of its head position relative to the nearby wall, which is about ten inches behind the robot's base. As it nears the wall, the salience of this drive increases, according to the same sigmoid curve as in the motor heat motivation, with $a = 5$ and $b = -5$, measured in inches from the robot base along the axis towards the wall. When high enough this drive grabs attention and causes the robot to move itself away from the wall, back towards the center of the table. Like the motor heat motivation, this reactive habit falls under the "self-preserving" category, on the grounds that striking the wall is damaging to robots. In future work, this motivation would be most useful if elaborated into a general "avoid collisions" behavior.

Finally, a "boredom" motivation has been implemented, with a constant low salience level of .1. When all other motivations have dropped below this level, the robot drops to a "boredom" state and stops addressing the other drives, on the grounds that they are all low enough not to need attention. Currently, this state relaxes the robot, allowing the human to freely move the robot, and performs no particular action; in future work this would be an ideal placeholder within which to place exploratory learning behaviors. If our curiosity motivation is analogous to an animal's drive to stay aware of its immediate environment, these periods of exploratory learning would be analogous to an animal playfully trying previously untested actions, or to learn more about relatively unfamiliar objects. Because this motivation sits at a constant level, it represents a threshold of inaction more than a separate motivation, and we omit it from our subsequent discussion.

## Selecting Actions

Given the basic types of drives – verbal commands, curiosity, heat levels, and wall avoidance – the system is designed to address the most salient context element at each moment, which ultimately has an automatic homeostatic effect. The premise is that any of the drives, when given attention, will execute habits that appropriately decrease its salience over time.

The system uses a biased competitive model of attention to select which drive to address at each moment (Desimone & Duncan 1995). Each possible subject of attention competes against the others according to its salience level, and the system directs its attention towards the highest one. Once attention is thus directed, the system adds a gradually diminishing positive bias to whatever it is currently paying attention to, so for some length of time it will be incrementally harder, but still possible, for other drives to preempt attention. The principle of biased competition also allows for other factors to bias the saliences, which will come into play when the goal-directed system tries to further inhibit actions not related to the current persistent goal.

The attentional bias is added to a value by initially decreasing its distance from 1.0 by one-third, and then gradually decaying this incremental value with a time constant of 1 minute. That is, suppose drive $n$ has value $D_n(t)$, and at time $t_0$, $D_n(t)$ is higher than all other $D_i(t)$, $i \neq n$. Then drive $n$ will receive the attention of the system, meaning that it takes control and carries out whatever habits are associated with it (heat drive causes sleep, curiosity causes looking, and interaction drive causes action execution).

Then, starting at time $t_0$, drive $n$ receives an attentional bias equal to $B_n(t) = \frac{1}{3}e^{-\frac{t-t_0}{\tau}}$, where $\tau$ is 60 seconds. The adjusted value of drive $n$ then becomes $(1 - D_n(t))B_n(t) + D_n(t)$. In words, the drive $n$ is treated by the system as if it moved closer to 1 by some proportion, while keeping with the asymptotic design. At $t_0$, this means that if drive $n$ just received control, its value will immediately inflate and ensure that it most likely keeps control. However, over the next minute, that inflation gradually decreases to allow other drives to compete more effectively. If another drive successfully takes control, the bias will be instantly removed from drive $n$, reset, and applied to the new drive instead.

Using this adjusted value prevents the system from *dithering*, where two motivations with similar priorities alternately grab control of the system and the time spent transitioning from one behavior to the other prevents either motivation from being addressed effectively.

### Changes in Drive Saliences = Personality Shifts

An interface (Figure 6) is used to monitor system salience levels and also to adjust the salience of each drive by a constant factor. Specifically, if $k_n$ is the factor for drive $n$ set by the interface, the adjusted salience is $D'_n = k_n(1-D_n(t)) + D_n(t)$, essentially a permanent attentional bias similar to the temporary attentional bias described in the previous section.

Increasing the sensitivity to motor heat causes the system to spend more time sleeping, and increasing the sensitivity to verbal commands causes the system to service all requests immediately, no matter what the level of other drives may be. In effect, shifts of the relative saliences of the top level drives lets us shift Ripley's "personality" from paranoid (high curiosity drive) to obedient (high verbal response drive) to simply lazy (high self-preserving drive). As a robot's behavior network grows in complexity, high-level behavioral design may be most easily accomplished by this style of salience adjustment.

## System Behavior

In the absence of human interaction, the robot constantly interleaves curiosity and sleep behaviors. With curiosity in control, Ripley looks around its environment, allowing the mental model to process objects and faces in each direction.


Figure 6: The graphical interface to the behavior system. On the left are bars indicating the strengths (saliences) of each motivation. From top to bottom, there are motivations to avoid wall collisions, a boredom threshold of activity, three heat indicators, six curiosity directions, and a motivation to carry out actions based on interaction with the human. The third heat indicator is currently in control. On the right side are sliders to manually adjust the saliences of the motivation types.

Over time, motor heat rises and Ripley moves towards a centered position and lies down on the table in order to allow its motors to cool off. As this happens, the curiosity drives continue to increase, but because the system biases towards its current activity, the heat levels have to drop for some period of time before the curiosity elements can regain control and begin looking around again.

If the robot's motivations are all fairly satisfied, it just sits still. In this state, if the person drags the robot towards the wall, it will attempt to take back control and move itself back towards the center of the table.

When the partner talks to Ripley, the drive to interact increases its salience to the preset value of 0.8. If another drive is currently higher than that value, the robot makes the preprogrammed verbal request for the person to "Ask again later." However, in most cases, Ripley will respond quickly to requests. With the exception of extreme situations (e.g., where Ripley has been forced not to move so that its mental model is completely out of date and thus Ripley is extremely curious), the robot stops resting or looking around and reaches for the appropriate object or makes an appropriate response, and then returns to its other activities.

### An Open System

An interesting aspect of the system stems from its *openness*. That is to say, the robot is not performing its behaviors in a vacuum of internal states, looping repetitively, or acting based solely on a prior plan, but rather it acts in response to changes in its environment and to actions by its partner. Essentially, the robot's internal representations and context awareness are open to a variety of externally derived influences, via its online processing. We believe that allowing immediate perceptual context to govern behavior is critical to embodied systems, and that any goal-directed processing should be a secondary influence on behavior.

For instance, the system's pre-programmed directions of observation, along with its mental model, react responsively to environmental changes and human actions. If the human moves or removes objects on the table, the robot updates its mental model while looking around and responds appropriately during subsequent interactions. Also, the updates of the mental model can occur while looking around due to the curiosity drive, or it can happen because the robot happened to look in that direction while performing user-assigned tasks or going to sleep. Furthermore, if Ripley looks to satisfy one of its curiosity zones, then that curiosity element considers itself satisfied and drops its value back to zero, regardless of what caused the robot to look in that direction. In fact, due to the compliant nature of the series elastic motor controllers, the human can even grab Ripley by the head and force the robot into arbitrary poses, and in the process cause Ripley's curiosity drive to be satisfied. Thus, far from acting out a fixed plan, the system's pre-programmed directions of observation interact responsively with other actions of the system.

Since Ripley's motor heat saliences model actual time-varying motor heat, anything that alters motor load affects the behavior of the system. Typical actions such as looking around will tend to increase motor heat to a moderate level. If the robot is instructed to pick up a relatively heavy object, motor heat will increase further than usual, and the robot will likely rest for a while afterwards. Most strikingly, if a person (or a wall) completely obstructs the motion of the robot from its intended actions, motor heat will increase very rapidly, and the robot will quickly stop its efforts and abruptly head for its sleeping position.

## Denotative and Connotative Semantics: A Robot's Perspective

From a theoretical perspective, one of our goals in building conversational robots is to develop new computational models of communicative meaning. We believe that, in contrast to symbolic models of computational semantics, which essentially define word meanings in terms of other word-like symbols (this is typically how semantic networks and first order predicate logic are used to encode lexical semantics in natural language processing systems), a theoretically significant alternative is to use an embodied platform like Ripley to ground linguistic meaning in terms of sensory-motor representations, using its basic perceptions to define words denotatively, and using the motivations of the habit system as a novel sort of connotative meaning.

In our approach to language grounding, the denotation of a word, when used in certain linguistic contexts, can be bound to its reference in terms of sensory expectations with respect to motor actions (Roy in press). For example, the word "red" denotes an expectation in visual color space, and the expectation holds with respect to an object which is visually observed (the motor action in this case is looking at the object and taking a visual reading of its optical form). For a second example, consider the denotation of "heavy." For Ripley, the word denotes an expectation of accumulated forces on its joints as it executes a lift procedure.

Linguists often distinguish what a linguistic structure denotes versus what it connotes. Connotation refers to the implied, subjective meaning of a word or phrase as opposed to its conventional, literal meaning. Computational semantics rarely addresses this aspect of meaning even though natural human language usage is steeped in connotations. By building Ripley, we see an emerging computational theory of connotation which can be summarized as follows:

> The connotative meaning of a word is the accumulated effect of the word's underlying sensory-motor structure with respect to the motivational system of the language user.

Continuing with the example of the meaning of "heavy", consider the connotation of the term for Ripley with respect to its three primary drives. In terms of doing what it is told to do, the meaning of "heavy," like any property name, is useful since it lets Ripley pick out the object that the human wants. In this respect, the connotation of heavy is positive – we might say, *positive with respect to the verbal response drive*. In contrast, any object that is heavy will be negative with respect to Ripley's other habit executions, since manipulating heavy objects accelerates motor heating, and heavy objects often slip. Although Ripley's curiosity drive is currently a strictly visually grounded process, it would be relatively straightforward to extend curiosity to include the weight of objects in which case Ripley would be compelled to lift all novel objects to characterize their weights. In such a setting, the negative relation of heaviness to motor heat would also emerge during the activity of satisfying curiosity, which would now involve lifting all novel objects.

To summarize, our emerging view of connotation is that the subjective meaning of a word and its underlying concept is a summary of the agent's sensory-motor experiences "projected" through the lens of the particular concept. We believe this nascent idea may lead to interesting new insights for NLP and, more generally, linguistic theory.

## Discussion and Next Steps

Our pragmatic goal is to create interactive robots capable of semi-autonomously assisting humans in various tasks. To this end, a voice interface capable of simple dialogue and command execution is only a first important step. The addition of a behavioral framework within which interaction with humans is just one habit gives our robot the ability to homeostatically prolong its own "lifespan" (by protecting its physical resources) as well as to update its knowledge of the world for efficient execution of future commands.

The two primary ways we expect to extend this work are 1) to add a goal system that sits atop the habit system, and 2) to move the processing of human interaction into the goal system.

The most important extension to this work is to add the goal system. As mentioned, human behavior is heavily governed by automatic habits, which are inhibited in a goal-directed manner in order to select only habits that support intentional actions. Having built a simple habit system that acts based on motivational contexts, the next step is to implement goal-directed inhibition by making use of the biased competition model which governs habit execution.

The typical functionality of the goal system is not to arbitrarily select and pursue random goals; rather, the goals represented in the goal system are the explicit versions of the implicit goals that habits are geared to automatically pursue. For instance, since high motor heat triggers a habit for resting, the automatic nature of this resting action implicitly leads towards a goal of reduced motor heat. The purpose of the goal system is to explicitly represent this goal, and to use this representation to accomplish two functions.

First, the goal system should also provide an explicit representation of expected results for each habit. Expectation violations should cause the goal system to take more control away from the habit system to diagnose the cause of the failure. Second, the goal system should provide a certain additional level of goal persistence, to further reduce dithering. This happens because the goal system adds more inhibition in the habit competition stage to habits that fail to contribute to the currently represented goal. However, other habits can still overpower the habit being supported by the goal system. If another habit has the power to seize attention, then the explicit version of its goal will enter the goal system and replace the previous goal.

One other effect of the goal system is that, unlike the habit system's single-valued attention, the goal system can represent multiple goals. One goal may come from the current habit, but other goals may come from human interaction or other knowledge. The interaction of the goal and habit system should thus be a mix of stimulus-driven actions that make sense in the current context, biased according to the currently represented goals of the system, along with a limited amount of explicit planning that uses knowledge about the results of actions.

The other important extension to the habit system is to move the processing of human interactions out of the habit system and into the goal system. Currently, human interaction is handled by a separate subsystem that allows interactions to be treated as if they were a simple habit execution. More plausibly, human interaction should be handled by the goal system as a source of explicit goals for the system to pursue. This is made possible by the goal system's ability to handle multiple goals at once, evaluating all conducive actions and biasing habit execution according to all the represented goals.

## Conclusion

Studies show that human and primate behavior are governed dually by the automatic execution of stimulus-driven habits, mediated by goal-directed plan-based inhibition. We have set out to implement a computational model of this phenomenon into a robot platform, with the ultimate goal of building a conversational robot with very rich grounded structures for processing interactions. To date, we have constructed a simple habit system that intersperses habits based on motivational contexts that minimize motor heat, maintain the mental model, and address user interactions.

## Acknowledgments

We would like to thank Peter Gorniak for building the Bishop parsing and referent system, and Nikolaos Mavridis for building the mental model.

## References

Blumberg, B. M. 1994. Action-selection in hamsterdam: Lessons from ethology. In Cliff, D.; Husbands, P.; Meyer, J.-A.; and Wilson, S. W., eds., *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press.

Carnegie Mellon University, Sun Microsystems Laboratories, Mitsubishi Electric Research Laboratories. 2004. Sphinx 4 Java Speech Recognizer. http://www.speech.cs.cmu.edu/cgi-bin/cmusphinx/twiki/view/Sphinx4/WebHome.

Desimone, R., and Duncan, J. 1995. Neural mechanisms of selective visual attention. *Ann. Rev. Neurosci.* 18:193–222.

Gorniak, P., and Roy, D. 2004. Grounded semantic composition for visual scenes. *Journal of Artificial Intelligence Research* 21:429–470.

Hsiao, K.; Mavridis, N.; and Roy, D. 2003. Coupling perception and simulation: steps towards conversational robotics. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 928–933.

Miller, E., and Wallis, J. 2004. Volition and the prefrontal cortex. In Chalupa, L., and Werner, J., eds., *The Visual Neurosciences*. MIT Press. 1546–1560.

Minsky, M. 1986. *Society of Mind*. Simon and Schuster.

Pratt, J.; Krupp, B.; and Morse, C. 2002. Series elastic actuators for high fidelity force control. *Industrial Robot* 29(3):234–241.

Roy, D.; Hsiao, K.; and Mavridis, N. 2004. Mental imagery for a conversational robot. *IEEE Transactions on Systems, Man, and Cybernetics* 34:1374–1383.

Roy, D. in press. Grounding language in the world: schema theory meets semiotics. *Artificial Intelligence*.

Tyrrell, T. 1993. *Computational Mechanisms for Action Selection*. Ph.D. Dissertation, Edinburgh University.

Viola, P., and Jones, M. 2001. Rapid object detection using a boosted cascade of simple features. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 511–518.