

## 18 Robotic Models of Self

Justin Hart and Brian Scassellati

Why do puppies chase their tails? Folk wisdom would tell us that a puppy is being playful or is seeking attention. Veterinarians would say that the puppy chases, and sometimes even bites, its own tail because it does not realize that this fascinatingly evasive object is actually part of its own body. While this behavior in a puppy is nothing to be concerned with, tail chasing in an older dog is often a sign of dementia, skin irritation, or anxiety. Whereas an older dog is expected to understand that its own tail is not an object that should be chased or bitten, a young puppy is still coming to understand the boundaries of its own body. We assume that through its experiences, a puppy is able to learn that its tail does belong to itself, perhaps by observing that its tail is a constant companion or that catching and biting its tail results in pain.

Like puppies, human infants are not born with a complete sense of themselves. During the first few months of life, infants must learn to discriminate between their own bodies, the movements of parents and others who are responsive to the child, and the movement of objects on television or of wind-blown leaves, that is, items that are unresponsive to the child's actions (Rochat, 2003). They must come to understand that the flailing fingers and arms that they often see in their cribs are part of themselves, that they will be able to control the movements of these strange-looking appendages, and eventually be able to effect desired changes in the world using them.

Traditionally, robots have had no sense of self, nor did they need it. In factory automation, or even in traditional task-based robotic systems, the robot carried out a specific goal by selecting between appropriate behaviors or by tuning the parameters of a fixed behavioral repertoire. These robots could not perceive their own mechanical bodies and did not need to discriminate between different types of activity within their environment. As robots become more complex (involving richer sensing and more degrees of freedom) and as they move out of the factory and into environments like our homes, schools, and hospitals, the need for these machines to be more aware of their own limits, their own capabilities, and the results of their own actions becomes critical. Robots should not chase their own mechanical tails.

To highlight this point, consider two robotic applications. First, consider a robot operating in a factory that constructs automobiles. This robot consists of a camera system that looks down onto a conveyor belt, a mechanical arm that can maneuver parts from one position on the belt to another, and the computational resources to recognize defects in these parts. As parts slide into view on the belt, the robot must decide if the part is defective. The robot must orient itself to defective parts, grasp them, and remove them from the line. In this case, distinguishing the robot's arm from other objects in the field of view can be accomplished simply in multiple ways. One solution would be to paint the arm and gripper a distinctive color that is not used elsewhere in the vicinity of the robot, allowing the robot to spot the gripper and the part by identifying the unique color. Another would be to preprogram the kinematics (the body structure and motion capabilities) of the robot. The kinematic equations could then be solved algebraically to identify how to move the arm to grasp particular parts. When wear and tear on the robot's parts slowly degrade the accuracy of these precoded equations, trained technicians can be on hand to recalibrate or reprogram the equations as needed. Faults can be detected simply by establishing boundaries on the robot's behavior. If the gripper moves too far away from the assembly line, or parts are not picked up as frequently as expected, a fault can be signaled. The system could then be stopped until a technician repairs the equipment.

Now compare this factory automation system to a robot designed to aid elderly homeowners carrying groceries or other supplies from their car to their kitchens. Perhaps with two arms and a wheeled base, this robot would need to perform grasping, lifting, and carrying of arbitrary packages under the direction of its owner. This home assistant robot requires many of the same behavioral capabilities as the factory robot; it too must recognize important components in the environment, grasp them, and maneuver them into appropriate positions. However, none of the easy-to-construct systems that were used in the factory robot are likely to be successful in the home assistance robot. We cannot count on selecting a unique color for the robot that completely distinguishes it from all homes and from all shopping packages. We also cannot rely on maintaining a perfect kinematic model to predict the locations of the robot's limbs—without the constant supervision of trained technicians, these equations are likely to be useful for only a short time. Instead, our robot requires some more flexible way of identifying itself, identifying when faults occur, and adapting to new configurations (such as when it is carrying a large shopping bag).

The real-world requirements of robotics add a dimension to the self-model not directly considered in other chapters of this book. To our formalism, this chapter will add the capability to reason about the robot's physical presence in the world, its construction, its sensory capabilities, and its interactions with its environment. As adaptive and self-trained kinematic and sensory self-models are introduced, we will also observe that lower-level processes, traditionally hard-coded into the system and buried

beneath convenient abstractions, will become first-class cognitive models accessible directly to the system (as discussed in this vol., chap. 1).

### Self-Identification

In ethology, the traditional test of self-awareness in animals is the mirror rouge test (Gallup, 1970). Figure 18.1 shows a chimpanzee participating in this test. First, a mirror is placed into the habitat of an animal and the animal is given time to acclimate to its presence. During this phase, many animals will engage their own reflection with either social or aggressive behaviors, as they do not recognize the reflection as themselves. After acclimating to the mirror, some animals, such as chimpanzees, will begin to use the mirror to groom themselves, in a recognizable self-directed behavior. The animal is then anesthetized, and a section of the body that can only be seen in the mirror (such as the forehead) is dyed. If the animal inspects the mark through use of the mirror, it is considered to have recognized its own appearance in the mirror.

Gallup's (1982) model supposes that the animal must have a self-concept, typically in the form of an image that resembles the animal. This supposition leads to a model of how to perform self-recognition based on similarity of appearance. An appearance-based model stores an explicit representation of appearance that is then matched against a current sensory state to determine if the animal (or robot) currently perceives itself. Though this technique has a simplicity that is appealing, there are many difficulties in implementing this solution. First, the perspective of the image is often seen as third person (as one would appear in a photograph), though this is clearly not easily



**Figure 18.1**

A chimpanzee subject of the mark test. (Photo used by permission of Daniel Povinelli.)

matched to a first-person perspective (as the animal might observe itself). Second, the complexity of the matching process makes a complete implementation of this strategy infeasible. To identify the difficulty inherent in this process, imagine looking down at your hand and attempting to catalog all of the possible shapes that your hand might form. The range of possible appearances from a fist to an open palm to a peace sign provides an endless variety of physical appearances.

A second methodology focuses not on visual appearance matching but rather on some matching process between the movement of the body and the visual scene. Mitchell (1997) supports the idea of kinesthetic-visual matching in which the only knowledge required to recognize oneself in the mirror is the relationship between the visual scene and proprioception. In our own research, we have demonstrated the effectiveness of this alternative explanation by constructing a robot that can distinguish between self and other. An early version of this system (Michel, Gold, & Scassellati, 2004) used temporal contingency to learn timing parameters that distinguished the movement of the robot's own arm (seen in the camera's field of view) from the movement of people in the environment. The robot estimated the delay between sending a motor command and observing a visual change. Though this approach had some advantages, it was limited in its extensibility and by sensor noise.


A more recent version of this system (Gold & Scassellati, 2007) uses a Bayesian kinesthetic-visual matching model to allow a humanoid robot to perform self-other discrimination and mirror self-recognition without social understanding and without an explicit kinematic model. A humanoid robot named Nico learned the relationship between its own motor activity and perceived motion by observing the movements of its arm for four minutes. Each new observation was used to update three models for each object in its visual field. The first model is that of random noise, generated with no structure over time. The second model consists of an observed internal state of motor activity that generates the external feedback of motion; thus, the consistency of the match between motor activity and motion dictates the likelihood of this model. The third model is that of motion generated by somebody else; it is identical to its own self-motion model, only the motor state is hidden and must be reasoned about probabilistically. Presented with a mirror, the robot then judged its mirror image to match its "self" model, while people were judged to be "animate others." Figure 18.2 shows the scene through Nico's cameras during this test. In this picture, we can see Kevin Gold in front of Nico, a mirror that Nico can see himself in, and to the right we see Nico's finger. In figure 18.3, we see that Nico has segmented out Kevin as an animate other, marking him in purple. Nico has marked himself both in the mirror and directly in his visual field in green. Other moving objects determined to be noise are marked in red.

Why pursue such research for a robotic system? What advantage does the ability to recognize oneself provide to a robot? One answer is that the modeling effort itself



**Figure 18.2**

Nico looking at himself in the mirror, with experimenter Kevin Gold behind it (Gold & Scassellati, 2007).



**Figure 18.3**

Nico's software segmenting himself and an animate other, Kevin Gold, from the scene (Gold & Scassellati, 2007).

has value, as it provides insight into potential methodologies and algorithms that may be occurring in biological systems. Though the fact that a robot performs a task in a certain way is never proof that a biological system also necessarily utilizes the same solution, the computational model can both provide a proof-of-concept for a particular solution and potentially provide insights into the nature of the problem itself (Webb, 2001). In this case, the fact that a kinesthetic-visual matching algorithm can successfully solve the self-identification problem leads us to question the necessity of purely visual appearance-based methods.

A second answer is that these self-identification algorithms are the first step toward a more comprehensive robotic model of self. Current research in our lab focuses on developing robotic self-models that integrate the kinematic and sensory systems of the robot (Hart, Scassellati, & Zucker, 2008). Kinematic self-models such as ours and others (Hersch, Sauser, & Billard, 2008) enable robots to learn through experience the structure of their bodies and how they move through space. We will argue that a robot that had a more comprehensive model of its body schema and of its own capabilities would provide connections to other areas that have been traditionally disparate areas of research in robotics: fault recognition and recovery, causal learning, and tool use.

### **Fault Detection and Recovery**

Though the majority of robotic systems operate with no fault-detection mechanism, the detection, identification, and diagnosis of faults in machinery is an active area of interest in both research and industry. Systems used to perform this in an automatic fashion offer both the capability to assist human technicians in diagnostic tasks as well as to allow machinery to automatically diagnose and recover from faults.

In industry, the dominant method to accomplish this task is rule-based diagnosis (Darwiche, 2000). These systems use hand-crafted sets of rules written by domain experts that are checked against the system's status. More popular in research is model-based diagnosis, in which a model of the system is developed using symbolic logic (de Kleer & Williams, 1987, 1989; Darwiche, 2000; Hofbaur & Williams, 2002). An automated theorem prover then uses this model along with status reports from devices in the system in order to perform diagnosis.

Rule-based diagnosis systems are favored in industry because they have a lower computational overhead and do not require a background in symbolic logic and artificial intelligence to understand (Darwiche, 2000). Model-based diagnosis systems offer a number of advantages including being easier to update and modify and allowing developers to mathematically prove properties of the model.

Perhaps the most intriguing use of model-based diagnosis to date has been the Livingstone system, which was employed in the Remote Agent software aboard NASA's *Deep Space One* probe (Muscettola, Nayak, Pell, & Williams, 1998). *Deep Space One* was

the first spacecraft to be controlled by artificial intelligence without human supervision. Though *Deep Space One's* self-model was built by scientists and engineers on the ground, it did use a logical model of itself while operating in space in order to adapt its control policy to systems reporting faults.

Fault detection as it is currently envisioned follows either rule-based or model-based techniques, both of which require a constant detection and recovery system to be preconstructed by the programmers when the system is initially deployed. These techniques cannot adapt to online changes in the system's hardware configuration, nor can they adapt to changes in the control architecture. For a robot that can construct its own model of its physical extent, its kinematic structure, and its capabilities, fault detection takes on a somewhat different role; fault detection becomes an ongoing process of comparing current short-term models of the robot's self with a more stable longer-term model of the robot's self. An adaptive model thus allows for a more flexible recognition process that is based on the perception of the robot's current capabilities that also allows for long-term modifications.

### Causal Learning

Causal learning is a research area concerned with the sequences of events that link causes with effects. Often modeled by the forward algorithm (Rabiner, 1989), which asserts that prior time steps have a causal relationship to future ones, causal learning often operates over symbolic descriptions of the world (which at times makes it difficult to apply in robotic systems). These symbols are linked together in either predetermined or statistically salient sequences to create causal chains that indicate the prevalence at which a particular event (the cause) results in the production of a secondary event (the effect). Though this learning is often symbolic in nature, there have been many attempts to ground these symbols in perceptually salient cues (Yoshikawa, Hosoda, & Asada, 2004; Yoshikawa, Tsuji, Hosoda, & Asada, 2004).

Notice that this process, by which a causal learning system searches for pairings of events separated in time, is very similar to the process of kinesthetic-visual matching described above for self-identification. Rather than seeking a visual stimulus to match an earlier motor command from the robot, we instead initially match any motor command from the robot with a later-occurring event. If these two events recur under similar actions and situations, we can imagine that the robot could learn to produce particular actions (causes) to create a certain desired result (the effect). While this process by itself may provide interesting evidence and goal-directed behavior to the robot system, the most common application for this type of learning is tool use, which we discuss as a special case below. Causal learning has also been studied in the context of fault detection and diagnosis, as implemented in the OCCAM system (Pazzani, 1990a,b).

## Tool Use

When a person uses a tool, that tool becomes causally tied to him or her. An interesting example of this incorporation into the self-model comes from Yamamoto, Moizumi, and Kitazawa (2005), in which it is demonstrated that when a person touches something with the tip of a tool, the sensory experience attached to that action is the sensation of feeling the tip of the tool contacting the object being touched rather than the feeling of the tool providing increased tactile resistance in the hand in response to the touch. In other words, while the person is using that tool, they perceive it as an extension of his- or herself.

Experiments to allow robots to build better models of this boundary between themselves and the rest of the world also mark the crucial difference between a robot that must be programmed to grip an object in its gripper and one that can learn to grip an object on its own. The current state-of-the-art is to preprogram robots with such capabilities. A robot with a causal model, however, can learn its own optimal gripping strategies. By modeling the relationship of objects in the environment to the self, rather than programming in grasping behaviors, future robotic systems may be able learn things such as tool use without needing to be programmed to use individual tools.

## Conclusion

We often think of metareasoning as a high-level component that can be added to existing agent architectures to oversee or monitor typical activity. This internal critic offers suggestions, monitors progress, or infers higher-level information from the mundane activities of the agent. Perhaps the most salient lesson from our work on self-modeling in robotic systems is that metareasoning can be built into some of the most basic components of these systems in order to solve real-world problems. This may include basic components that are often considered to be complete and beyond need for revision (such as low-level control algorithms and kinematic models). As part of the basic construction of an agent, metareasoning and self-modeling systems can serve to unify a range of problem domains under a single system-wide design. This integration also allows for problems (like self-recognition) that on the surface appear to be high-level cognitive tasks to become part of the moment-to-moment operation that is critical to agent behavior. Perhaps we should not consider metareasoning systems as an additional module that can be added late in the design process but rather as central guiding principles to self-governed behavior.

In this chapter, we have promoted a viewpoint that unifies a few subfields of robotics that until now were studied in isolation. By recasting the primary questions of these fields as part of the continual process for constructing an accurate model of self,





**Figure 18.4**

Nico, the humanoid robot, used by the authors in robotic self-modeling experiments.

we demonstrate that each of these questions can be characterized as part of a larger domain. Adaptive, self-taught models of self provide a framework for studying causal learning, tool use, kinematic analysis, and fault detection and recovery. While the study of these fields independently will continue to advance the state of the art, it is our belief that the study of these as part of an integrated self-model will allow for even more fundamental insights into how to build useful, adaptive, and practical robotic systems and may cast light onto the underlying processes of self-identification that biological systems must also solve.

### **Acknowledgments**

Support for this work was provided by National Science Foundation awards 0534610 (Quantitative Measures of Social Response in Autism), 0835767 (Understanding Regulation of Visual Attention in Autism through Computational and Robotic Modeling), and CAREER award 0238334 (Social Robots and Human Social Development). Some

parts of the architecture used in this work were constructed under the DARPA Computer Science Futures II program. This research was supported in part by a software grant from QNX Software Systems Ltd., hardware grants by Ugohe Inc., and generous support from Microsoft and the Sloan Foundation. Research supported by AFOSR and NGA.

## References

- Darwiche, A. (2000). Model-based diagnosis under real-world constraints. *AI Magazine*, 21, 57–73.
- de Kleer, J., & Williams, B. (1987). Diagnosing multiple faults. *Artificial Intelligence*, 32(1), 97–130.
- de Kleer, J., & Williams, B. (1989). Diagnosis with behavioral modes. In N. S. Sridharan (Ed.), *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI '89)* (pp. 124–130). San Francisco: Morgan Kaufmann.
- Gallup, G. (1970). Chimpanzees: Self-recognition. *Science*, 67(3914), 86–87.
- Gallup, G. (1982). Self-awareness and the emergence of mind in primates. *American Journal of Primatology*, 2, 237–248.
- Gold, K., & Scassellati, B. (2007). A Bayesian robot that distinguishes “self” from “other.” In *Proceedings of the Twenty-Ninth Annual Meeting of the Cognitive Science Society (CogSci2007)* (pp. 384–392). Mahwah, NJ: Lawrence Erlbaum.
- Hart, J., Scassellati, B., & Zucker, S. W. (2008, May). Epipolar geometry for humanoid robotic heads. In B. Caputo & M. Vincze (Eds.), *Proceedings of the 4th International Cognitive Vision Workshop (ICVW 2008)* (pp. 24–36). Berlin: Springer.
- Hersch, M., Sauser, E., & Billard, A. (2008). Online learning of the body schema. *International Journal of Humanoid Robotics*, 5, 161–181.
- Hofbaur, M., & Williams, B. C. (2002). Mode estimation of probabilistic hybrid systems. In C.J. Tomlin & M.R. Greenstreet (Eds.), *Proceedings of the International Conference on Hybrid Systems, Computation, and Control*. Lecture Notes in Computer Science 2289 (pp. 253–266). Berlin: Springer.
- Michel, P., Gold, K., & Scassellati, B. (2004, September). Motion-based Robotic self-recognition. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Piscataway, NJ: IEEE Press.
- Mitchell, R. (1997). Kinesthetic-visual matching and the self-concept as explanations of mirror self-recognition. *Journal for the Theory of Social Behaviour*, 27(1), 17–39.
- Muscettola, N., Nayak, P., Pell, B., & Williams, B. (1998). Remote agent: To boldly go where no AI system has gone before. *Artificial Intelligence*, 103(1–2), 5–47.

Pazzani, M. (1990a). *Creating a memory of causal relationships: An integration of empirical and explanation-based learning methods*. Hillsdale, NJ: Lawrence Erlbaum.

Pazzani, M. (1990b). Learning fault diagnosis heuristics from device descriptions. In Y. Kodratoff & R. S. Michalski (Eds.), *Machine learning III: An artificial intelligence approach* (pp. 214–234). San Mateo, CA: Morgan Kaufmann.

Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2), 257–296.

Rochat, P. (2003). Five levels of self-awareness as they unfold early in life. *Consciousness and Cognition*, 12, 717–731.

Webb, B. (2001). Can robots make good models of biological behaviour? *Behavioral and Brain Sciences*, 24(6), 1033–1050.

Yamamoto, S., Moizumi, S., & Kitazawa, S. (2005). Referral of tactile sensation to the tips of L-shaped sticks. *Journal of Neurophysiology*, 93(5), 2856–2863.

Yoshikawa, Y., Hosoda, K., & Asada, M. (2004). Cross-anchoring for binding tactile and visual sensations via unique association through self-perception. In *Proceedings of the Fourth International Conference on Learning and Development*. Piscataway, NJ: IEEE Press.

Yoshikawa, Y., Tsuji, K., Hosoda, K., & Asada, M. (2004, October). Is it my body? Body extraction from uninterpreted sensory data based on the invariance of multiple sensory attributes. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robotics and Systems*. Piscataway, NJ: IEEE Press.