Robotic Self-Models Inspired by Human Development

Justin W. Hart and **Brian Scassellati**
Department of Computer Science
Yale University, New Haven, CT 06511, USA
justin.hart@yale.edu, scaz@cs.yale.edu

Motivation

Traditionally, in the fields of artificial intelligence and robotics, representations of the self have been conspicuously absent. Capabilities of systems are listed explicitly by developers during construction and choices between behavioral options are decided based on search, inference, and planning. In robotics, while knowledge of the external world has often been acquired through experience, knowledge about the robot itself has generally been built in by the designer. Built-in models of the robot's kinematics, physical and sensory capabilities, and other equipment have stood in the place of self-knowledge, but none of these representations offer the flexibility, robustness, and functionality that are present in people. In this work, we seek to emulate forms of self-awareness developed during human infancy in our humanoid robot, Nico. In particular, we are interested in the ability to reason about the robot's embodiment and physical capabilities, with the robot building a model of itself through its experiences.

Grounding in Developmental Psychology

Traditionally, developmental theorists viewed infants as being born into a state of confusion, devoid of an overall organization of their perception and action. Infants' actions at birth, under this view, are entirely random and without intent, self-centered and detached from perception of their environment (Mahler, Pine, and Bergman 1975; Gergely 2000). Another view of this phase is that infants are unable to differentiate themselves from their environment. The basic knowledge is not yet available to the infant to allow them to know where they end and the rest of the world begins. In spite of the state of confusion that they experience in either scenario, they are able to learn by continuously making observations about the impact of their actions, and will develop their sense of self along with other cognitive capabilities.

Though there are many aspects of this self-understanding which have been addressed in the literature, we concentrate on two that encompass the agent's sensorimotor capabilities and ability to impact its environment, the *ecological self* and *self-efficacy*. We choose these because the skills represented by these facets of the self represent a basic model of perception and action in which the self is an integral component.

Rochat describes the ecological self as a cohesive, egocentric model of the body and its relationship to the environment arising from coordination between the senses (Rochat 2001). It provides a coherent intersensory model, which has been demonstrated to exist at birth through a number of recent experiments. Infants have been demonstrated to be more likely to exhibit the "rooting reflex", which assists in breast-feeding, when touched on the cheek, rather than when they touch themselves on the cheek (Rochat and Hespos 1997), demonstrating that they understand that it is their own hand, rather than an external stimulus. They have also been shown to open their mouth in preparation to receive their fist, rather than accidentally placing their fist in their mouth and leaving it there (Rochat, Blass, and Hoffmeyer 1988), premeditating the act of sucking the fist. Though these abilities are present at birth, that does not mean that they are innate. There is evidence that the capabilities required to support them are learned in the womb (Mahler, Pine, and Bergman 1975; Gergely 2000; Rochat, Blass, and Hoffmeyer 1988). We are able, however, to witness visual and motor capabilities develop after birth due to lack of visual stimulation in the womb and skeletomuscular development after birth.

If the ecological self describes the body, then knowledge of self-efficacy describes its capabilities. The sense of self-efficacy describes the causal relationship between motor actions and the body and objects in the environment (Rochat 2001). Due to their limited motor skills, it was once thought that infants did not develop a sense of self-efficacy until much later than is now believed. Experiments that cater to the limited sensorimotor repertoire of the neonate have now revealed that infants have the capacity to understand causal relationships from the time that they are born. Using pressure-sensitive pacifiers to detect subtle differences in rate and intensity, newborn infants have been demonstrated to change the patterns and pressure of their sucking on a pacifier in order to have their mother's face, rather than that of another female appear on a screen (Walton, Bower, and Bower 1992). Infants have also been demonstrated to move their heads (Papousek 1992) and legs (Rovee-Collier 2005) in order to cause mobiles mounted over their cribs to move, and have demonstrated pre-reaching behaviors that emerge

---

Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

before they are strong enough to perform reaching and grasping behaviors (Rochat and Senders 1991). These results indicate that infants have an ability to comprehend and learn causal relationships even within the first few months of life. It is this understanding of the causal relationship between their motor behavior, their bodies, and objects in their environment that their sense of self-efficacy arises from.

Bringing these two concepts into robotics changes the way that we look at many problems. Traditionally, roboticists have looked at tasks, such as manipulating objects, in terms of the goal to be carried out. That kind of reasoning still plays a role here, but in this work, the robot learns about the form of its body, and the nature of causal relationships between its motor actions and the motion of objects. This knowledge can later be used to reason about such tasks, rather than learning about tasks directly, or programming primitive actions to be chained together by a higher-level reasoning system.

Self-modeling in Robotic Systems

The traditional test of self-awareness in humans and animals is the mirror test (Gallup 1970). In this test, first, a mirror is placed in the cage with the animal to be tested for a habituation period. After habituation, the animal is sedated. While asleep, an odorless, non-tactile mark is placed on its forehead. If the animal uses self-directed behavior to inspect the mark in the mirror, it is considered to be self-aware, if not, it is not. According to Gallup's hypothesis (Gallup 1982), in order to perform this feat, the animal must have a self-concept against which to compare the image in the mirror. It is the difference between what the animal expects to see in the mirror, and what it does see, that prompts the self-inspection behavior. In a competing theory, Mitchell (1997) states that mirror self-recognition can be accomplished simply through a model of kinesthetic-visual matching, and that no self-concept is necessary. Prior work in our lab (Gold and Scassellati 2007) explores this subject, allowing our upper-torso humanoid robot, Nico, Figure 1, to recognize itself in a mirror using models reminiscent of Mitchell's hypothesis.

The plan for our current work is to construct a self-taught robotic model of self, composed of both an ecological model of self and a model of self-efficacy. The ecological model of self encompasses the body's shape, structure, and orientation. It indicates where body parts exist with respect to each other, including end-effectors such as hands and sensors such as the eyes. The self-efficacy model is a causal model of the relationship of motor actions to the body, and the body to its environment. It indicates how motor actions can change the body's configuration and how this interacts with objects in the environment. Together, these form a complete sensorimotor model of the robot. Rather than the typical approach of having human designers model kinematics up front, calibrate sensors, and design libraries of manipulation strategies, our model will allow the robot to learn the relationship of its end-effectors and sensors to each other and the environment, and the underlying principles of manipulation and tool use through a process of self-exploration and interaction with the environment.

Robots provide a unique context for the study of self-aware reasoning processes in that their embodiment and interactions with the environment are fundamental to the tasks that they perform. We are able to bring work from developmental psychology to bear on this problem, modeling human developmental processes in our hardware. These low-level self-modeling abilities will serve as a building-block for later self-referential social cognitive abilities.

The Ecological Self-Model The ecological self-model will be a combined kinematic-sensory model that represents the robot's physical embodiment. In order to do this, we will develop a model of the robot's kinematics online, through its experience. These systems will not rely on external calibration rigs, nor will they require explicit calibration routines to be run by a trained operator. Instead, using techniques inspired by human development, the robot will learn the relevant calibration parameters in an online and self-supervised fashion. Because it is the most well-developed robotic sense that is analogous to a human sense, we will concentrate on the use of vision in this study.

We can think of the process of learning an ecological self-model as involving three components, though we hope for these components to be learned concurrently, in a unified fashion, sensor calibration, kinematic calibration, and development of a joint kinematic-sensory model. To put this in perspective, by sensor calibration we mean the classical problem of camera calibration, which can be broken into two components, the *intrinsic parameters* which describe the camera itself in terms of the projection that it imposes on the world, and the *extrinsic parameters* which describe its position in space. Similarly, a kinematic model is a model of the geometry and motion of the robot. Generally, in a stereo-vision system, the extrinsic parameters of the cameras are measured with respect to each other, and considered to be static. In the development of our ecological self-model, we consider the camera to be a kinematic endpoint whose motion can be measured along with other objects in the system, such as the robot's hands. This motion correlates to the motion of the head and eyes in a human. This combined kinematic-sensory model gives robots the benefits that the ecological self provides to humans. It allows the robot to continue to see in 3D after the motion of its cameras, is a predictive model of the position of end-effectors in the visual field, and is a continuously-updated kinematic self-model.

In (Hart, Scassellati, and Zucker 2008), we discuss a method for measuring the motion of the cameras in a stereo active-vision system with respect to each other in terms of eye motion. Later in this paper, we discuss our recent work in which the robot learns a kinematic self-model through experience. Other current work includes updating the model of camera motion to be learned similarly, allowing us to combine the two models and measure the motion of the cameras in a more general way, allowing for head motion, and calibrating the intrinsic parameters of the cameras as part of a general self-calibration process.

The Model of Self-Efficacy Rochat describes self-efficacy as knowledge of how one's actions can enact changes upon their environment (Rochat 2001). We will instantiate and implement a concrete model of this capability in a robotic system. We view self-efficacy as beginning with the motor actions that one can take. These actions change the configuration of the body, which, in turn, can move objects, which in turn, can move other objects. We will model self-efficacy by inferring the causal Bayesian relationship between the robot's motor actions, its body, and objects in the environment. This will be the basis for the robot's motor skills, including improvements on self-other discrimination, manipulating objects, and tool use. These capabilities will be evaluated through tasks requiring the robot to move itself and objects in its environment.

In our model of self-efficacy, we plan to start with the causal relationship between the robot's motor actions and its motion. The model from (Gold and Scassellati 2007) is a simple model of this that doesn't account for the robot's kinematic structure. This model will explicitly account for the robot's physical presence in terms of its ecological self, and will build a causal Bayesian model of the change, in terms of joint angle, that a motor action enacts on the ecological self. In these terms, it can be viewed as an improvement on the self-other discrimination algorithm from (Gold and Scassellati 2007), but we will view it as a starting point for the motor actions that the robot can enact on the world.

If the robot's kinematic endpoints are objects that are directly under its control, then the rest of the objects in the world are only under its control when in contact with these endpoints. Consider, for instance, the interaction between a block and the robot's gripper. While the robot's gripper is always under its control, the block is temporarily under the robot's control when in the gripper. By modeling the causal relationship between the gripper and the block, the robot will be able to learn how it is able to manipulate objects in its environment. Additionally, objects have causal relationships between each other. Blocks can be stacked on each other, if a ball in motion collides with another ball, it will cause that ball to move. By learning these object to object relations, the robot can learn not only what manipulations are possible on its environment, but also, through this chain of causal relations, tool use.

Kinematic Learning

In prior work, we have learned a simple kinematic model of the motion of robotic eye-cameras in an active stereo-vision system (Hart, Scassellati, and Zucker 2008). In more recent work, (Hart et al. 2010), we have learned the kinematic structure of our robot's arm by witnessing its motion in its visual field. Starting with a calibrated stereo vision system, we reconstruct the 3D position of the robot's end effector. Tracking its motion in space, we are able to reconstruct the kinematic chain that produced it.

We present two basic methods for performing this task. In the first, we observe that the motion of the end-effector attached to a revolute joint will move through a circular path in space. By reconstructing this circle from points observed to lie along it, by moving a single joint, we are able to recover the Denavit-Hartenberg parameterization of the joint, which is a conventional description of its kinematics. Having more than one joint in the kinematic chain complicates this relationship, but the solution to this problem is beyond the scope of this paper. A thorough treatment of the topic is currently under review (Hart et al. 2010).

The second method for learning this kinematic model begins with a candidate model of the kinematic chain. Moving the robot's joints into a random position, we compare the position that the robot expects the end-effector to appear in with the position that it actually appears in. We perform non-linear optimization on the distance between the position predicted by the candidate forward kinematic model and the observed position of the end-effector.

One way of dealing with tool use which has been mentioned in the kinematic learning literature (Hersch, Sauser, and Billard 2008), is to model the end-point of a tool in use by the robot as a change in the robot's kinematic structure. Temporarily incorporating a tool into an agent's kinematic structure is a method of dealing with tool use that is supported in the neurophysiology literature (Yamamoto, Moizumi, and Kitazawa 2005). While the approach in (Hersch, Sauser, and Billard 2008) is to simply retrain the kinematic model, treating the tool tip as though the position of the robot's hand has moved, our model is capable of both this approach and an alternative approach of modeling the tool as an object that is held in the robot's hand. The latter model correlates more closely to our model of self-efficacy, since we view the causal relationship between the hand and the tool as a mechanism for the extension of the self.

Knowledge of the camera's intrinsic and extrinsic parameters, which have been obtained through classic camera-calibration techniques, allows us to project our 3D model of the robot's end-effector position down into the 2D images obtained through the robot's cameras. This is an example of an intersensory capability provided by our model of the ecological self, which correlates to the sorts of intersensory phenomena described by Rochat (2001), by combining the visual and kinesthetic senses.

We have implemented this model on our robot Nico. To determine the position of the end-effector as witnessed in the visual field, we take the centroid of a color segmented region, which has been marked in red tape on either the robot's end-effector or the tool-tip of the held tool. The projection of the end-effector position from the trained kinematic model in the visual field is computed. In our experiment, we first train on the hand, then retrain with either the pen or the hammer, arriving at the model for tool use. See Figure 1. Results to be reported in an upcoming paper demonstrate the effectiveness of this model. In current work, we are extending this model to utilize an inverse-kinematic solution that will allow for the robot to reach for objects in space as it trains this model. Under this updated model, Nico will attempt to reach objects in space, refining its ecological self-model as it does so.

Conclusion

In this paper, we have detailed our current research program of robotic self-modeling. In this program, we seek to build models of the ecological self and self-efficacy, early forms of self-awareness that arise in infants. Our work is inspired by developmental psychology and neuroscience, and seeks

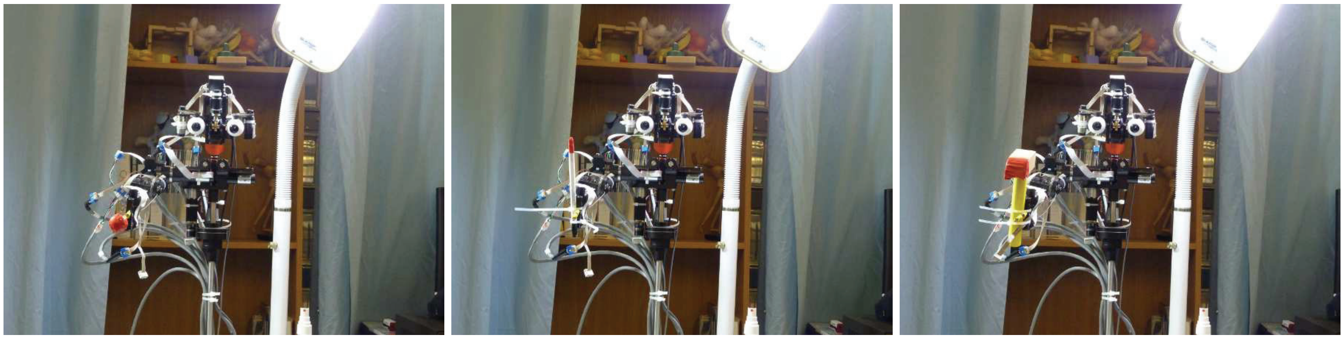
(a) Hand (b) Pen (c) Hammer

Figure 1: Images of Nico. The red electrical tape is the marker used to track the motion of the end effector. We replaced the hand with a ball to track in the first test, cable tied a pen to the arm in the second, and a toy hammer in the third.

to both improve the state-of-the-art in robotics by incorporating the "self" into robotic reasoning processes, as well as further our knowledge of metacognition by modeling these forms that are found in humans. It represents a significant departure from traditional robotics practice in that, rather than reasoning only about the task at hand, it reflectively attempts to learn and reason about the machinery carrying out the task. Our current progress is promising, and we believe that this work has the potential to significantly impact both the robotics and metacognition communities.

Acknowledgements

Support for this work was provided by National Science Foundation award #0835767 (Understanding Regulation of Visual Attention in Autism through Computational and Robotic Modeling). The authors also acknowledge the generous support of Microsoft and the Sloan Foundation. This work was made possible through a software grant from QNX Software Systems Ltd.

References

Gallup, G. G. 1970. Chimpanzees: self-recognition. *Science* 167:86–87.

Gallup, G. G. 1982. Self-awareness and the emergence of mind in primates. *American Journal of Primatology* 2:237–248.

Gergely, G. 2000. Reapproaching Mahler: New perspectives on normal autism, symbiosis, splitting and libidinal object constancy from cognitive developmental theory. *Journal of the American Psychoanalytic Association* 48(4):1197–1228.

Gold, K., and Scassellati, B. 2007. A bayesian robot that distinguishes "self" from "other". In *Proceedings of the 29th Annual Meeting of the Cognitive Science Society (CogSci2007)*.

Hart, J.; Golub, D.; Avrunin, E.; Hayes, B.; and Scassellati, B. 2010. Vision-driven kinematic calibration of the ecological self. In *Proceedings of the IEEE International Conference on Development and Learning (ICDL)*. In submission.

Hart, J.; Scassellati, B.; and Zucker, S. W. 2008. Epipolar geometry for humanoid robotic heads. In *Proceedings of the 4th International Cognitive Vision Workshop (ICVW)*.

Hersch, M.; Sauser, E.; and Billard, A. 2008. Online learning of the body schema. *International Journal of Humanoid Robotics* 5(2):161–181.

Mahler, M. S.; Pine, F.; and Bergman, A. 1975. *The Psychological Birth of the Human Infant: Symbiosis and Individuation*. Basic Books.

Mitchell, R. W. 1997. Kinesthetic-visual matching and the self-concept as explanations of mirror-self-recognition. *Journal for the Theory of Social Behavior* 27(1):86–87.

Papousek, H. 1992. Experimental studies of appetitional behavior in human newborns and infants. *Advances in Infancy Research* 7(xix–liii).

Rochat, P., and Hespos, S. 1997. Differential rooting response by neonates: Evidence for an early sense of self. *Early Development and Parenting* 6(3–4):105–112.

Rochat, P., and Senders, S. 1991. Active touch in infancy: Action systems in development. In Weiss, M., and Zelano, P., eds., *Newborn attention: Biological constraints and influence of experience*. Norwood, N.J.: Ablex, 2nd edition. 412–442.

Rochat, P.; Blass, E. M.; and Hoffmeyer, L. B. 1988. Oropharyngeal control of hand-mouth coordination in newborn infants. *Developmental Psychology* 24(4):459–463.

Rochat, P. 2001. *The Infant's World*. Cambridge, Massachusetts and London, England: Harvard University Press.

Rovee-Collier, C. 2005. Learning and memory in infancy. In Osofsky, J., ed., *Handbook of Infant Development*. New York: John Wiley and Sons, 2nd edition. 98–148.

Walton, G.; Bower, N.; and Bower, T. 1992. Recognition of familiar faces by newborns. *Infant Behavior and Development* 15(2):265–269.

Yamamoto, S.; Moizumi, S.; and Kitazawa, S. 2005. Referral of tactile sensation to the tips of l-shaped sticks. *Journal of Neurophysiology* 93:2865–2863.