

Motion-Based Autonomous Grounding: Inferring External World Properties from Encoded Internal Sensory States Alone*

Yoonsuck Choe and Noah H. Smith

Department of Computer Science

Texas A&M University

College Station, TX 77843-3112

choe@tamu.edu, nhs6080@cs.tamu.edu

Abstract

How can we build artificial agents that can autonomously explore and understand their environments? An immediate requirement for such an agent is to learn how its own sensory state corresponds to the external world properties: It needs to learn the semantics of its internal state (i.e., grounding). In principle, we as programmers can provide the agents with the required semantics, but this will compromise the autonomy of the agent. To overcome this problem, we may fall back on natural agents and see how they acquire meaning of their own sensory states, their neural firing patterns. We can learn a lot about what certain neural spikes mean by carefully controlling the input stimulus while observing how the neurons fire. However, neurons embedded in the brain do not have direct access to the outside stimuli, so such a stimulus-to-spike association may not be learnable at all. How then can the brain solve this problem? (We know it does.) We propose that motor interaction with the environment is necessary to overcome this conundrum. Further, we provide a simple yet powerful criterion, sensory invariance, for learning the meaning of sensory states. The basic idea is that a particular form of action sequence that maintains invariance of a sensory state will express the key property of the environmental stimulus that gave rise to the sensory state. Our experiments with a sensorimotor agent trained on natural images show that sensory invariance can indeed serve as a powerful objective for semantic grounding.

Introduction

For an agent (natural or artificial) to be autonomous in the truest meaning of the word, it must be able to learn, on its own, about the external world and its basic properties. The very first obstacle here is that the agents have direct access only to its internals: It does not have direct access to the world, nor to the world properties, since it cannot step outside of itself to observe the world (cf. Pierce & Kuipers 1997 and Philipona et al. 2003; 2004). This is in short the problem of *grounding* (Harnad 1990): How can we make the meaning (or semantics) of the internal state intrinsic to the system (Freeman 1999), rather than it being provided by a third party? (The problem is also related to the concept of

*Preliminary results were presented at a workshop (Choe & Bhamidipati 2004). Thanks to S. Heinrich for discussions. Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

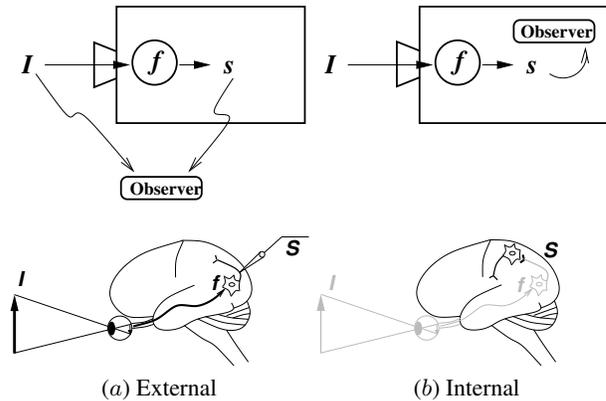


Figure 1: **External vs. Internal Observer.** The problem of decoding internal sensory state is seen from (a) the outside, and (b) from the inside of the brain. The sensory neurons shown inside the brain perform a transformation from input I to spikes s using the function f (i.e., s encodes properties of I). The task is to find out the property of the input I given the internal state s . (a) With full access to the input I and the state s , or if the property of f is known, what s stands for can be inferred. (b) However, with the lack of those knowledge, such an inference may be impossible.

natural semantics by Cohen & Beal (1999).) For example, consider the situation depicted in figure 1. If we have access to both the external and the encoded internal state of an agent (or the brain), then we can infer what external properties are represented by the internal state, but this involves a third party (i.e., we, as external observers; figure 1a). On the other hand, if we can only observe the encoded internal state while trapped inside the agent (i.e., intrinsic to the system), then trying to infer external world properties may seem futile (figure 1b; also see the discussion on the limitations of isomorphism in Edelman 1999). However, we know that autonomous agents like us are fully capable of such an inference. How can that be the case?

Let us consider what should (or could) be the minimal set of things given to an agent at birth (or switch-on time; Weng 2004). From a biological perspective, the set will include raw sensors and actuators, and rudimentary initial processes built on top of those, such as the orientation-tuned neurons

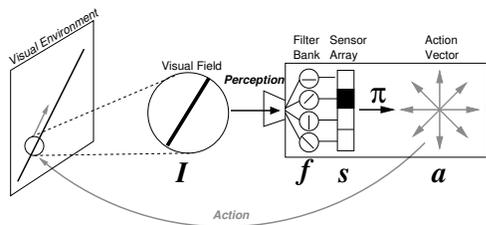


Figure 2: **Visual Agent Model.** An illustration of a simple sensorimotor agent is shown. The agent has a limited field of view where part of the input from the environment (I) is projected. A set of orientation-tuned units f (sensory primitives) receive that input and transform it to generate a pattern of activity in the sensory array s (black marks active). In the example shown here, the 45° unit is turned on by the input. Based on the sensory array pattern, a mapping π to motor action vector a is determined, resulting in the movement of the visual field in that direction, and then a new input is projected to the agent. Note that the agent is assumed to be aware of only its internal sensory state s , thus it has no knowledge of I , nor that of f .

in the primary visual cortex (Daw 1995), or simple motor primitives in the central pattern generator (CPG) (Marder & Calabrese 1996). Note that CPGs can generate these motor patterns in the absence of any sensory input (Yuste *et al.* 2005). One thing to notice here is that this minimal set includes motor capabilities, which is missing from the picture in figure 1b, and it turns out that the motor side is key to solving the grounding problem. In this paper, we will show that these motor primitives are central in associating external stimulus properties with internal sensory states. We will also propose a simple yet powerful learning criterion for grounding: sensory invariance. In previous work, we demonstrated that the idea basically works on a toy problem using synthetic images, with a simple ad hoc learning rule (Choe & Bhamidipati 2004). In this paper, we employ a learning rule based on standard reinforcement learning (Q learning), and present results and analyses on natural images.

The remainder of the paper is organized as follows. First, we will provide a sketch of the general framework. In the section that follows, we will present details of our model and the learning rule. The next section provides the experimental procedure, along with results and analysis. Finally, we will discuss the contribution of our work and its relation to other works, and conclude the paper with a brief outlook.

Model of the Agent and the Environment

Let us consider a simple sensorimotor agent (figure 2), with a limited field of view. The visual input is transformed by an orientation filter (modeling primary visual cortical neurons) into a spike pattern in the sensory array. The sensory array forms the sensory primitive s that the agent must consider and infer the stimulus property as related to the external visual environment. The agent has no access to the input I , nor to the functional properties of filter f .

The agent is internal to this model, and it is not clear

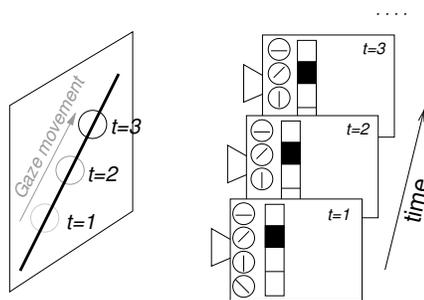


Figure 3: **Sensory Invariance Driven Action.** The state of the agent during movement of its gaze is shown. Under the condition shown here, the internal state does not change, and the stimulus property represented by the unchanging state (45° orientation) is congruent with the property of the motion (45° diagonal motion). Thus, by generating a motor behavior while trying to maintain internal state invariance results in that behavior *mirroring* the sensory property conveyed by that internal state.

that it will ever be capable of associating a visual orientation with the sensory spikes. If we consider the agent as in figure 1b there is no solution; however, with the addition of motor primitives as represented in figure 2, a solution can be found. The addition of the motor primitives is critical; by relating sensor activity and motor command, certain aspects of the sensor properties can be inferred. A crucial insight that occurred to us at this point was that certain kinds of action tend to keep the sensory activity pattern stable (i.e. invariant) during on-going movement, and the property of this motion reflects that of the sensory stimulus.

Consider the state of the agent in figure 3: the sensory unit for 45° orientation is activated at time $t = 1$. Now move the gaze diagonally along the 45° input: This will keep the 45° sensor stable during the motor act. This action would directly reflect the property of the input stimulus, and leads us to conclude that association of internal sensory states to sensory-invariance driven action can serve as the “meaning” for the sensory primitive. This way, even without any prior knowledge, or direct access to the external world, agents can learn about the key environmental property (here, orientation) conveyed by the sensory state (cf. “proximal representation” in second-order isomorphism; Edelman 1999). Further, the generated motor output is behaviorally relevant, meaning that the sensory state and the action are congruent (cf. Bowling, Ghodsi, & Wilkinson 2005). Thus, *sensory invariance* can serve as a simple, yet powerful criterion for enforcing this mapping while serving as a basis of grounding for internal representational states. In this paper, we will use a standard reinforcement learning algorithm, together with the invariance criterion, to infer external world properties from internal state information.

Model Architecture

The general model of the agent is shown in figure 2. We will first describe input preprocessing, then response generation, and finally the learning rule that will allow the agent to map

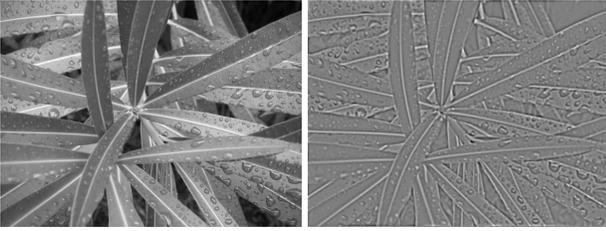


Figure 4: **Raw (I_R) and Difference-of-Gaussian (DoG) Filtered Input (I_D)**. A natural image input (left) and its DoG-filtered version (right) used in the training simulations are shown. The image was 640×480 in size.

the sensory state to meaningful motor pattern using sensory-invariance as the criterion.

Initial input preprocessing

The input image is first convolved by a difference-of-Gaussian (DoG) filter to simulate the preprocessing done in the lateral geniculate nucleus. The filter is defined as follows for each pixel (x, y) :

$$D(x, y) = g_{\sigma/2}(x, y) - g_{\sigma}(x, y), \text{ where} \quad (1)$$

$$g_b(x, y) = \exp\left(-\frac{(x-x_c)^2 + (y-y_c)^2}{b^2}\right) \quad (2)$$

is a Gaussian function with width b and center (x_c, y_c) . The parameter σ was $k/4$ for filters of size $k \times k$ ($k = 15$ for all experiments), and $x_c = y_c = 8$. The original raw image I_R is convolved with the DoG filter to generate I_D :

$$I_D = I_R * D, \quad (3)$$

where “*” is the convolution operator. I_D is then subtracted by its pixel-wise mean, and normalized:

$$I_D(x, y) := \frac{I_D(x, y) - \mu_D}{\max_{u,v} |I_D(u, v)|}, \quad (4)$$

where (x, y) is the pixel location, and μ_D the mean for all $I_D(x, y)$. An example of I_R and I_D are given in figure 4. The input I was a 31×31 square area sampled from I_D , centered at the agent’s gaze.

Sensorimotor primitives

The sensory state is determined by an array of oriented Gabor filters. Each Gabor filter G_i is defined at each location (x, y) as:

$$G_{\theta, \phi, \sigma, \omega}(x, y) = \exp^{-\frac{x'^2 + y'^2}{\sigma^2}} \cos(2\pi\omega x' + \phi), \quad (5)$$

where θ is the orientation, ϕ the phase, σ the standard deviation of the Gaussian envelope, and ω the spatial frequency. The values x' and y' are calculated as:

$$x' = x \cos(\theta) + y \sin(\theta), \quad (6)$$

$$y' = -x \sin(\theta) + y \cos(\theta). \quad (7)$$

All $m \times m$ -sized filters G_i shared the same width ($\sigma = m/2$), phase ($\phi = -\pi/2$), and spatial frequency ($\omega =$

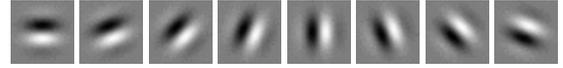


Figure 5: **Oriented Gabor Filters**.

$2/m$). Orientation θ varied over the index i , where $\theta_i = \lfloor (i-1)\pi/n \rfloor$, given the number of filters n . Figure 5 shows eight Gabor filters G_i ($i = 1..n$, for $n = 8$).

The filter response is a column vector s' of elements s'_i , corresponding to the vectorized dot-product of the input and the Gabor filter:

$$s'_i = \sum_{x,y} G_i(x, y) I(x, y). \quad (8)$$

The vector s' is then normalized by its l_2 -norm $|s'|$:

$$s' := \frac{s'}{|s'|}. \quad (9)$$

The current sensory state index s is determined by:

$$s = \arg \max_{\theta_i, i=1..n} s'_i, \quad (10)$$

where each s corresponds to a unique orientation of θ as described above.

For each orientation, there are two matching gaze directions, such as 0° and 180° motion for $\theta = 0^\circ$. Thus, for n orientation filters, the motion direction set contains $2n$ movement-directions (figure 2):

$$A = \left\{ (d \cos(\theta), d \sin(\theta)) \mid \theta = \frac{(i-1)\pi}{n}, i = 1..2n \right\}, \quad (11)$$

where d is the travel distance of each movement ($d = 7$), θ the direction, and n the number of orientation filters. Each action vector $(a_x, a_y) \in A$ changes the agent’s center of gaze from (x, y) to $(x + a_x, y + a_y)$. When the location of the gaze reached the boundary of I_D , the gaze location was wrapped around and continued on the opposite edge of I_D .

Learning algorithm

Consider a particular sensory state s_{t-1} at time $t-1$, taking action a_{t-1} takes the agent to sensory state s_t . The state transition depends on the particular edge feature in the visual scene, and is probabilistic due to the statistical nature of natural images. The reward is simply the degree of invariance in the sensory states across the response vectors s'_{t-1} and s'_t :

$$r_t = s'_t \cdot s'_{t-1}, \quad (12)$$

where “ \cdot ” represents the dot-product. When filter response is invariant, reward is maximized ($r_t = 1$), and in the opposite case minimized ($r_t = -1$). This provides a graded measure of invariance rather than a hard “Yes” or “No” response.

The task is to form state-to-action mapping that maximizes reward r_t at time t . This is basically a reinforcement learning problem, and here we use the standard Q -learning algorithm (Watkins & Dayan 1992). (Note that other reinforcement learning algorithms such as that of Cassandra, Kaelbling, & Littman (1994) may be used without loss of

		A: direction of motion							
		→	↗	↑	↖	←	↘	↓	↙
S: sensory state (orientation)	○	0.5	0	0	0	0.5	0	0	0
	○/	0	0.5	0	0	0	0.5	0	0
	○\	0	0	$Q(s_i, a_j)$	0	0	0	0.5	0
	○	0	0	0	0.5	0	0	0	0.5

Figure 6: **Q table.** An illustration of the Q table with four orientation states is shown. Since for each orientation there can be two optimal directions of motion, there are eight columns representing the eight possible directions of motion. An ideal case is shown above, where $Q(s_i, a_j)$ is max in the two optimal directions for a given sensory state, and hence the diagonal structure. (Note that the actual values can differ in the real Q table.)

generality.) The agent determines its action at time t using the learned state-action function $Q(s_t, a_t)$ for state s_t and action a_t (see figure 6). Assuming that the Q function is known, the agent executes the following stochastic policy π at each time step t :

1. Given the current state s_t , randomly pick action a_t .
2. If a_t equals $\arg \max_{a \in A} Q(s_t, a)$,
 - (a) then perform action a_t ,
 - (b) else perform action a_t with probability proportional to $Q(s_t, a_t)$.
3. Repeat steps 1 to 3 until one action is performed.

To mimic smooth eye movement, momentum was added to the policy so that $a_t = a_{t-1}$ with probability 0.3. Also, the probability of accepting a move in step 2(b) above was controlled by a randomness factor c (set to 1.8 during training) so that if c is large, the chance of accepting a low-probability action is increased.

Up to this point, we assumed that the true Q function is known, upon which the policy is executed. However, the Q function itself needs to be learned. Following Mitchell (1997), we used Q -learning for nondeterministic rewards and actions to update the Q table:

$$Q_t(s, a) := (1 - \alpha_t)Q_{t-1}(s, a) + \alpha_t \left(r_t + \gamma \max_{a' \in A} Q_{t-1}(s', a') \right), \quad (13)$$

where s' is the state reached from s via action a , γ the discount rate ($= 0.85$ in all experiments), and α_t defined as: $\alpha_t = \frac{1}{1 + v_t(s, a)}$, where $v_t(s, a)$ is the number of visits to the state-action pair (s, a) up to time t (initial α was 1.0).

Because the design of the agent implies two optimal actions for each input state, each round of Q learning was restricted to the current action's cardinal block (left or right half of the Q table). For policy execution, one of the two halves of the Q table was randomly selected with each half having equal chance of being selected.

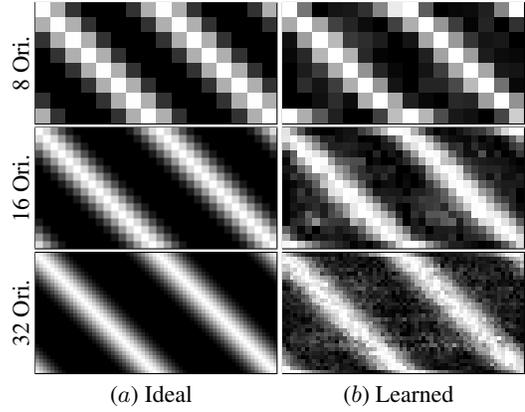


Figure 7: **Ideal and Learned $Q(s, a)$ Values.** The grayscale representation of the (a) ideal and the (b) learned Q table are shown for the three experiments (8, 16, and 32 orientation sensors). Black represents the minimum value, and white the maximum value. Note that the true ideal values in (a) should be two identity matrices appended side-by-side. However, given the fairly broad tuning in orientated Gabor filters, a small Gaussian fall-off from the central diagonal was assumed to form an approximately ideal baseline.

Experiments and Results

In order to test the effectiveness of the learning algorithm in the previous section for maximizing invariance, and to observe the resulting state-action mapping, we conducted experiments on the natural image shown in figure 4. We tested three agents with different number of sensory primitives: 8, 16, and 32. The corresponding motor primitives were 16, 32, and 64. The agents were trained for 100,000, 200,000, and 400,000 iterations, respectively. As the Q table grows in size, visiting a particular (s, a) pair becomes less probable, so the training time was lengthened for the simulations with more sensory primitives to assure an approximately equal number of visits to each grid (s, a) .

Figure 7 shows the ideal vs. the learned Q table. The learned Q tables show close correspondence to the ideal case. As shown in figure 6, the ideal case would show two optimal directions of motion for a particular sensory state. For example, 0° sensory state would be maintained if the gaze moved either in the 0° or the 180° direction, under ideal circumstances (the input is an infinite line with orientation 0° : cf. O'Regan & Noë 2001). Thus, in principle, the true ideal case should be two (appropriately scaled) identity matrices pasted side-by-side, as in figure 6. However, as the Gabor filters are broadly tuned (i.e., not very thin and sharp), there needs to be some tolerance since moving in a close-enough direction will still result in a high degree of invariance as measured by equation 12. Thus, in practice, those tables shown in figure 7a would be a good approximation of the ideal case (these tables were obtained by convolving the true ideal case with Gaussian kernels).

To measure quantitatively the performance of the learning algorithm, we adopted an error measure by comparing the ideal Q table and the learned Q table. The root mean squared error (RMSE) in the Q tables was obtained as:

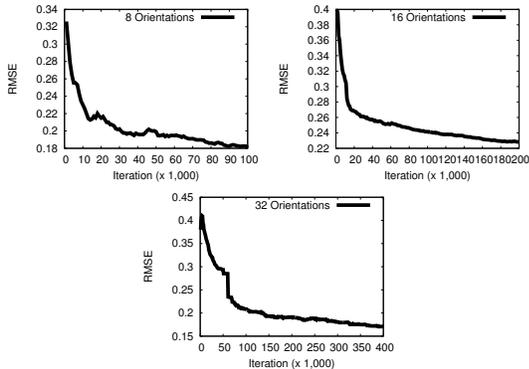


Figure 8: **Quantitative Measure of Performance.** The root mean squared error (RMSE) in the normalized Q table (figure 7b) compared to that of the ideal case (figure 7a) is shown for the three experiments. Each data point was gathered from every 1,000 training iterations. In all cases, convergence to a low error level is observed.

$$E = \sqrt{\frac{\sum_{s,a} (Q'_I(s,a) - Q'_L(s,a))^2}{n}}, \quad (14)$$

where E is the error in Q ; $Q'_I(s,a)$ and $Q'_L(s,a)$ are the normalized, ideal and learned, Q values, for state-action pair (s,a) ; and n the total number of state-action pairs in the Q table. The normalized Q values were obtained with:

$$Q'(s,a) = \frac{Q(s,a) - \mu_Q}{\max_{s,a} Q(s,a) - \min_{s,a} Q(s,a)}, \quad (15)$$

where μ_Q is the mean Q value over all (s,a) pairs.

Figure 8 shows the evolution of the error in Q estimate over time for the three experiments. All three experiments show that the error converges to a stable level.

Finally, we observed the gaze trajectory before and after learning, to see how well the expressed motor pattern corresponds to the world properties signaled by the internal sensory state. Figure 9 shows the results for the three experiments. Initially, the gaze trajectories are fairly random. Also, there seems to be no apparent relation between the gaze trajectory and the underlying image structure. (The input image (figure 4) has a set of radially arranged leaves centered near the middle, slightly to the left.) On the contrary, the gaze trajectory after learning shows a striking pattern. The trajectories are more directed (longer stretches in the same direction), and they show close correspondence to the image structure. For example, for all rows in figure 9b, the trajectories emerge from and converge to the center of radiation of the leaves. What is more important is that, at any point in time, the agent will have an internal sensory state corresponding to the local orientation in the input, and the generated motor pattern will be in the direction that is congruent with that external orientation property. Thus, through learning to maximize invariance in the internal sensory state through action, the agent can very well infer the external world property signaled by its internal, encoded, sensory state. Here, the basis of inference for the agent is its own motor act, thus grounding is based on the “behavior” of the motor system, not from the sensory system.

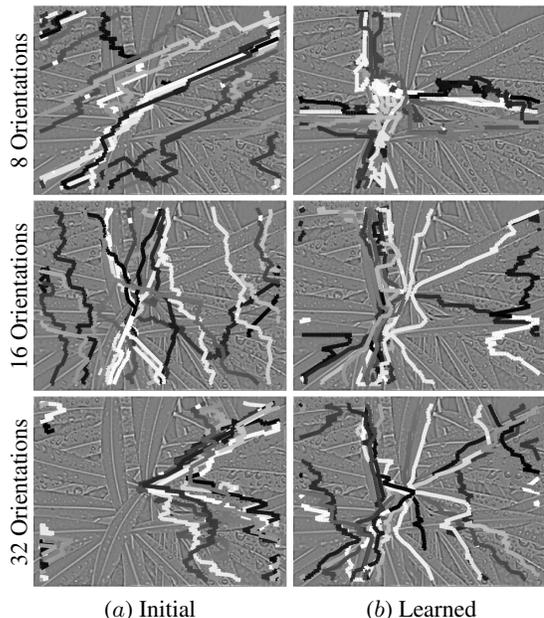


Figure 9: **Gaze Trajectory Before and After Learning.** The gaze trajectory generated using (a) initial Q and (b) learned Q are shown for the three different experiments. In each plot, 1,000 steps are shown. The trajectory was colored black→gray→white over time so that it is easier to see the time-course, especially where there is a large overlap in trajectories. (The color was repeated after 768 steps.) For these experiments, the randomness factor c was reduced to 1.0 to encourage “exploitation” over “exploration”. (a) Initially, the gaze trajectory is fairly random, and does not show high correlation with the underlying image structure. (b) However, after training, the trajectory closely follows the prominent edges in the natural image (around the radial edges coming out from the center). The oriented property of the motor pattern resulting from the sensory-invariance criterion is congruent with the underlying image structure. (The trajectory is a bit random due to the stochastic policy.)

Discussion

The main contribution of this work is in the demonstration of how a sensorimotor agent can infer about external world properties based on its internal state information alone. We showed that even without any prior knowledge about external world properties or any direct access to the environment, an agent can learn to express the sensory properties through its own actions. Our key concept was *invariance*, which served as a simple (simple enough for biology) yet powerful criterion for grounding. Note that our contribution is not in saying that action matters (others have successfully argued for the importance of action: Arbib 2003; Linás 2001; Brooks 1991). Rather, it can be found in *how* action can provide grounding for the internal sensory state and *what criterion* is to be used. We presented early conceptual work in Choe & Bhamidipati (2004), but it had theoretical (use of an ad hoc learning rule) and experimental shortcomings (synthetic inputs), both of which are overcome in this paper.

Our work is similar in spirit to those of Pierce &

Kuipers (1997), Philipona, O'Regan, & Nadal (2003), and Philipona *et al.* (2004), where they addressed the same problem of building up an understanding of the external world based on uninterpreted sensors and actuators. However, Pierce & Kuipers focused more on how basic primitives can be constructed and used, and Philipona *et al.* took a different route in linking environmental properties and internal understanding (the concept of *compensability*).

One apparent limitation of our model is that it seems unclear how the approach can be extended to grounding of complex object concepts. The answer is partly present in figure 9: If we look at the motor trajectory, we can already see the overall structure in the environment. However, for this to work, memory is needed. Our agent currently does not have any form of long-term memory, so it cannot remember the long trajectory it traced in the past. If memory is made available, the agent can in principle memorize more complex motor patterns, and based on that ground complex object concepts. A parallel work in our lab showed preliminary results on how such an approach can be advantageous compared to straight-forward spatial memory of the image (Misra 2005). These results suggest that motor primitives may be an ideal basis for object recognition and generalization (cf. *motor equivalence* of Lashley (1951)).

It is not surprising that certain neurons in the brain are found to associate sensory and motor patterns in a direct manner: Rizzolatti *et al.* (1996) discovered “mirror neurons” in monkey prefrontal cortex, which are activated not only by visually observed gestures, but also by the motor expression of the same gesture. The role of these neurons have been associated with imitation, but in our perspective, these neurons may be playing a deeper role of semantic grounding.

Conclusion

In this paper we analyzed how agents can infer external world properties based on its encoded internal state information alone. We showed that action is necessary, and motor pattern that maintains invariance in the internal state results in that motor pattern expressing properties of the sensory state. The sensory state can thus be grounded on this particular motor pattern. We expect our framework and approach to provide deeper insights into the role of action in autonomous grounding in artificial and natural agents.

References

- Arbib, M. A. 2003. Schema theory. In Arbib, M. A., ed., *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA: MIT Press, 2nd edition. 993–998.
- Bowling, M.; Ghods, A.; and Wilkinson, D. 2005. Action respecting embedding. In *Proc. of the 22nd Int. Conf. on Machine Learning*, 65–72.
- Brooks, R. A. 1991. Intelligence without representation. *Artificial Intelligence* 47:139–159.
- Cassandra, A. R.; Kaelbling, L. P.; and Littman, M. L. 1994. Acting optimally in partially observable stochastic domains. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 1023–1028. AAAI Press.
- Choe, Y., and Bhamidipati, S. K. 2004. Autonomous acquisition of the meaning of sensory states through sensory-invariance driven action. In Ijspeert, A. J.; Murata, M.; and Wakamiya, N., eds., *Biologically Inspired Approaches to Advanced Information Technology*, Lecture Notes in Computer Science 3141, 176–188. Berlin: Springer.
- Cohen, P. R., and Beal, C. R. 1999. Natural semantics for a mobile robot. In *Proceedings of the European Conference on Cognitive Science*.
- Daw, N. 1995. *Visual Development*. New York: Plenum.
- Edelman, S. 1999. *Representation and Recognition in Vision*. Cambridge, MA: MIT Press.
- Freeman, W. J. 1999. *How Brains Make Up Their Minds*. London, UK: Wiedenfeld and Nicolson Ltd. Reprinted by Columbia University Press (2001).
- Harnad, S. 1990. The symbol grounding problem. *Physica D* 42:335–346.
- Lashley, K. S. 1951. The problem of serial order in behavior. In Jeffress, L. A., ed., *Cerebral Mechanisms in Behavior*. New York: Wiley. 112–146.
- Llinás, R. R. 2001. *I of the Vortex*. Cambridge, MA: MIT Press.
- Marder, E., and Calabrese, R. L. 1996. Principles of rhythmic motor pattern production. *Physiological Reviews* 76:687–717.
- Misra, N. 2005. Comparison of motor-based versus visual representations in object recognition tasks. Master's thesis, Department of Computer Science, Texas A&M University, College Station, Texas.
- Mitchell, T. M. 1997. *Machine Learning*. McGraw-Hill.
- O'Regan, J. K., and Noë, A. 2001. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* 24(5):883–917.
- Philipona, D.; O'Regan, J. K.; Nadal, J.-P.; and Coenen, O. J.-M. D. 2004. Perception of the structure of the physical world using unknown multimodal sensors and effectors. In Thrun, S.; Saul, L.; and Schölkopf, B., eds., *Advances in Neural Information Processing Systems 16*, 945–952. Cambridge, MA: MIT Press.
- Philipona, D.; O'Regan, J. K.; and Nadal, J.-P. 2003. Is there something out there? Inferring space from sensorimotor dependencies. *Neural Computation* 15:2029–2050.
- Pierce, D. M., and Kuipers, B. J. 1997. Map learning with uninterpreted sensors and effectors. *Artificial Intelligence* 92:162–227.
- Rizzolatti, G.; Fadiga, L.; Gallese, V.; and Fogassi, L. 1996. Premotor cortex and the recognition of motor neurons. *Cognitive Brain Research* 3:131–141.
- Watkins, C. J. C. H., and Dayan, P. 1992. Q-learning. *Machine Learning* 8(3):279–292.
- Weng, J. 2004. Developmental robotics: Theory and experiments. *Int. J. of Humanoid Robotics* 1:199–236.
- Yuste, R.; MacLean, J. N.; Smith, J.; and Lansner, A. 2005. The cortex as a central pattern generator. *Nature Reviews: Neuroscience* 2:477–483.