

# Autonomous Acquisition of the Meaning of Sensory States Through Sensory-Invariance Driven Action

Yoonsuck Choe and S. Kumar Bhamidipati

Department of Computer Science  
Texas A&M University  
3112 TAMU  
College Station, TX 77843, USA,  
choe@tamu.edu, sarvanik@cs.tamu.edu

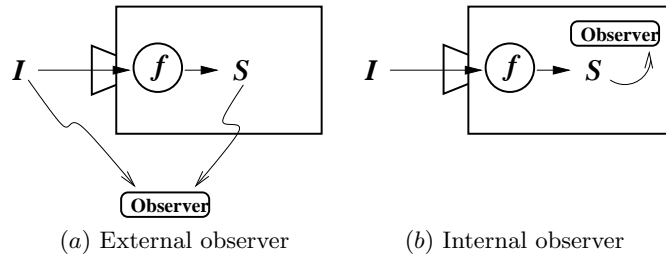
**Abstract.** How can artificial or natural agents autonomously gain understanding of its own internal (sensory) state? This is an important question not just for physically embodied agents but also for software agents in the information technology environment. In this paper, we investigate this issue in the context of a simple biologically motivated sensorimotor agent. We observe and acknowledge, as many other researchers do, that action plays a key role in providing *meaning* to the sensory state. However, our approach differs from the others: We propose a new learning criterion, that of *on-going* maintenance of *sensory invariance*. We show that action sequence resulting from reinforcement learning of this criterion accurately portrays the property of the input that triggered a certain sensory state. This way, the meaning of a sensory state can be firmly grounded on the choreographed action which maintains invariance in the internal state.

## 1 Introduction

Information technology has been largely driven by the growth in quantity, speed, and precision, but less by the improvement in quality and relevance of information. For that reason, while the amount of data being accumulated daily is staggering, our *understanding* of the data is not. The problem boils down to that of *meaning* [1, 2], i.e., what do these data mean and how can software systems understand the meaning of the data that they process? In this paper, we put this problem in the context of the brain, which is the only known device which naturally processes meaning, and find out what could be a potential solution.

The brain is made up of 100 billion neurons [3], which generate a complex pattern of activity in response to sensory stimuli from the external environment. A fundamental question in brain and cognitive sciences is, how do we understand what this pattern means? To make the question even simpler, we can ask what does a spike of a single neuron *mean*? [4]. Even this reduced problem is not trivial, and it took an enormous effort to come to the current state of understanding, beginning from muscle spindles [5] to cat visual cortical neurons [6] to sophisticated stimulus reconstruction methods developed lately (see, e.g., [4]).

*In Proceedings of the First International Workshop on Biologically Inspired Approaches to Advanced Information Technology (Bio-ADIT 2004), Lausanne, Switzerland, 2004.*



**Fig. 1. External vs. Internal Observer** The problem of decoding neural spikes is seen from the outside (a), and from the inside (b) of a perceptual agent. The neuron shown as a circle inside the box performs a input  $I$  to spike  $S$  transformation using function  $f : I \rightarrow S$ . The function  $f$  is supposed to be a highly complex one, and the neuron may be deeply embedded in the agent (i.e., it is not at the immediate sensory transduction stage, such as the photo-receptors). The task then is to find out the property of input  $I$  given just the spikes  $S$ .

A popular approach to this question is through associating the neural spikes with the stimulus that triggered those spikes [7] (see [8] for a review). Such methods have been successful in characterizing the neural spiking properties and accurately predicting the stimulus given just the spike train. This method involves the experimenter systematically varying the environmental stimulus while measuring the neural response (see, e.g., [4] chapter 2), so that at a later time, when only the spike train is observed, something can be said about the stimulus property. Mathematically, this is conveniently written using the Bayes theorem [4] (see Figure 1a):

$$P(I|S) = \frac{P(S|I)P(I)}{P(S)},$$

where  $I$  is the input stimulus and  $S$  is the spike train. Note that the likelihood term  $P(S|I)$  requires that we have either an empirical statistics or a reasonable model of the stimulus-to-spike translation. Thus, the interpretation of the current spike train  $P(I|S)$  seems to *depend* on direct knowledge about the stimulus properties, one way or another, which introduces the problem of circularity (cf. [9]).

Now suppose we ask the same question “what does a single spike mean?” to *another* neuron in the brain where such spikes are received (Figure 1b). Because this neuron does not have immediate knowledge about the environmental stimulus associated with the spike it received nor that of the receptive field properties (as it does not have the tools of an experimenter) the neuron cannot apply the technique mentioned above. (This problem can also be seen in the context of the Bayesian *theorist*, i.e., not merely an observer; an issue raised by Jepson and Feldman [10].) For example, consider a similar situation depicted in figure 2. Inside the agent, the only available information is the sensory array activity, so if we are trapped inside this agent without access to the visual environment outside, we can never figure out what the sensor activity means.

This paper begins by realizing this as a genuine issue. Our contention is that such a dilemma can be overcome through learning how to associate sensory activity to the motor actions the brain itself generates. The importance of sensorimotor learning has been emphasized in the past by many researchers:

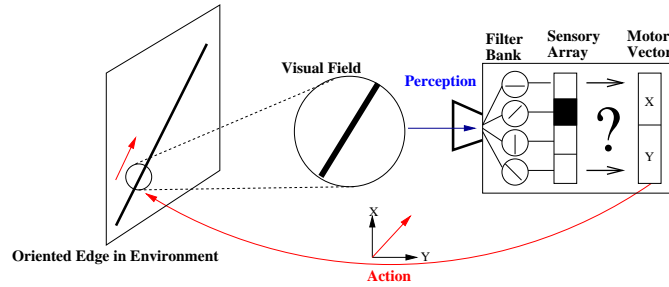
1. schema theory [11, 12];
2. learning of sensorimotor contingency [13, 14];
3. two-level interdependence of perception and action [15];
4. ecological perception of affordances [16];
5. subsumption architecture in robotics [17];
6. sensory-motor coordination in autonomous agents [18–21];
7. dynamical systems approach in agent-environment interaction (reviewed in [22]);
8. learning of natural semantics in robots [23];
9. mirror neurons and imitation in primates [24, 25];
10. motor learning enhancing perceptual performance [26];
11. the role of action in meaning and semantics in a dynamically coupled system [1] [2] (pp.13–17);
12. sensory substitution through active exploration [27, 28];
13. fixed action patterns (FAP) and thought as internalized action [29] (pp.134–141); and finally,
14. the role of action in consciousness [30] (pp.193–196), all recognize action as a key element in intelligent brain function.

Building upon these, we begin by examining how action can help in autonomous discovery of meaning in agents as shown in figure 1b. Our problem formulation is similar in spirit to Philipona et al. [14] and Pierce and Kuipers [20, 21], where a sensorimotor agent has to learn about its own raw sensors and actuators. The twist is that we provide a simple criterion that can exactly link the sensory states and the associated actions in a meaningful way.

Below, we first define the problem in terms of a sensorimotor agent we introduced in figure 2, and propose a learning algorithm based on *on-going sensory-invariance driven* motor action. The basic idea is that the agent has knowledge about its own movements, and the movements that it generates that reliably activate a particular sensor in the sensor array constitute the meaning of that sensor’s spike. The acquired meaning for each sensor and the resulting behavioral patterns are presented next, followed by discussion and conclusion.

## 2 Meaning of Spikes in a Sensorimotor Agent

To better illustrate our point, let us consider a small, concrete example as shown in figure 2, a simple sensorimotor agent. The agent has a limited visual field, and the incoming visual signal is transformed via the oriented filters (mimicking primary visual cortical neurons) into a spike pattern in the *sensory array*. Let us further assume that the agent does not have any knowledge (e.g., about the receptive field structure) of its oriented filters. The task of the agent then is to



**Fig. 2. A Sensorimotor Agent** An illustration of a simple sensorimotor agent is shown. The agent has a limited visual field where the input from the environment is projected. A set of orientation-tuned neurons receive that input and generate a pattern of activity in the sensory array (black marks active). In the situation here, the  $45^\circ$  sensor is turned on by the input. Based on the sensory array pattern, after some processing (signified by “?”), the  $x$  and  $y$  values of the motor vector is set, resulting in the movement of the visual field and a new input is projected to the agent.

attach meaning to its own sensory array activity pattern, i.e., to come to the understanding that each sensor represents a certain oriented visual input.

Imagine we are inside this agent, isolated from the world outside the box, sitting near the big “?” sign. It is questionable then whether we can ever be able to associate an orientated visual input stimulus with the spikes generated in the sensor array because we cannot peek outside, and we do not know the particular mechanism of the filters. The spike, in principle, could have been generated from any sensory modality, e.g., auditory or tactile input.

The only way we can see this issue resolved is through action, that is, the movement generated by the agent. This point is not entirely obvious at first, so let us elaborate a little bit what we mean. As shown in figure 2, we included the capability of action in the agent. The agent is able to gaze at different parts of the scene by moving around its visual field. The  $x$  and  $y$  variables correspond to the movement of the visual field in the  $x$  and the  $y$  direction, respectively. Thus, these two variables are like motor commands. We, sitting on that “?” sign, can generate different combinations of  $(x, y)$  values and observe the changing pattern in the sensory array. By relating the sensor activity and the motor command that was just generated, *certain aspects* of the sensor property can be recovered. We believe this is generally agreeable, but it is too general. It begs the question of *what* is that “certain aspects” of the sensory property and *how* can they be learned?

A crucial insight that occurred to us at this point was that certain kinds of action tend to keep the sensory activity pattern to remain unchanged (i.e., *invariant*) during vigorous movement, and this action exactly reflects the property of the sensory stimulus. For example, consider the state of the agent as shown in figure 2, where a  $45^\circ$  input is presented, and the corresponding sensor is activated in the agent. Now imagine we move the visual field according to the motor

vectors  $(1, 1), (1, 1), \dots, (1, 1), (-1, -1), (-1, -1), \dots, (-1, -1)$ , which corresponds to a back-and-forth movement along the  $45^\circ$  diagonal (i.e., aligned on the input). Such an action will keep only the  $45^\circ$  sensor turned on during the motor act, i.e., the sensory array stays invariant. We can see that this motion, generated while trying to keep the sensor array unchanged, has led the agent to perform an act which reflects the stimulus. Thus, we are led to conclude that associating this kind of sensory-invariance driven action with spikes can potentially serve as a *meaning* for each sensory neuron.<sup>1</sup>

To test this insight that ascribing meaning to sensory neuron activity is possible through learning the sensorimotor association based on sensory-invariance, we implemented a learning agent following the description in figure 2. The following sections describe the learning rule of the agent, followed by the results.

### 3 Learning of Sensory-Invariance Driven Action

Consider the agent described above (figure 2). We define a simple learning rule based on our idea of sensory-invariance driven action. The agent has the current state of its sensors  $S$  (the sensory array), and a set of actions  $D$  (possible combinations of the motor vector) that it can perform. For simplicity, we limit the sensor state set  $S$  to four different values

$$S \equiv \{0^\circ, 45^\circ, 90^\circ, 135^\circ\}, \quad (1)$$

which correspond to the four different orientation preference (note that  $0^\circ$  is the same as  $180^\circ$  etc.) of the sensors, and the action set  $D$  to eight different categories

$$D \equiv \{0^\circ, 45^\circ, 90^\circ, 135^\circ, 180^\circ, 225^\circ, 270^\circ, 315^\circ\}, \quad (2)$$

which are the possible *directions* of motion of the visual field with a limited distance of movement. Thus, the above corresponds to  $(1, 0), (1, 1), (0, 1), (-1, 1), \dots$  in terms of the motor vectors we mentioned earlier (the motion in  $x$  and  $y$  were either 0 or 1).

The learning task of the agent can then be treated as a standard reinforcement learning problem with a Markov assumption (see, e.g., [31, 32]). The goal of the agent is to select an action from the action set  $D$  that maintains the sensory array activity *invariant*. Thus, the reward is simply the degree of sensory-invariance in successive stages of action. More formally, the agent has to learn a policy function  $\pi$ ,

$$\pi : S \rightarrow D, \quad (3)$$

at step  $t$  which selects a direction of motion  $d_t \in D$  based on the previous state  $s_t \in S$  so that the resulting reward  $r_t$  is *maximized*. The execution of the policy at each state  $s_t$  results in reward:

$$r_t = r(s_t, d_t), \quad (4)$$

---

<sup>1</sup> Note that the invariance of this kind is different from that in Philipona et al. [14] where invariance is gained as a result of compensated motion, but not *during* the motion itself.

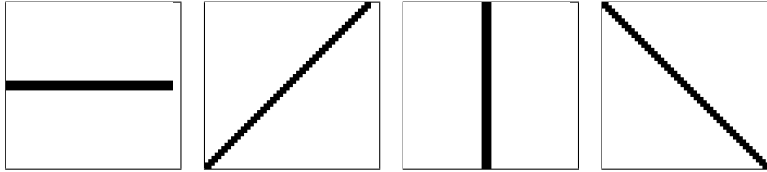
based on the reward function  $r(s, d)$  for  $s \in S, d \in D$ , and this function is updated as follows:

$$r_{t+1}(s, d) = \begin{cases} r_t(s, d) + \alpha * f_t & \text{if } s_t = s_{t-1}, \\ r_t(s, d) - \alpha * f_t & \text{if } s_t \neq s_{t-1}, \end{cases} \quad (5)$$

where  $r_{t+1}$  is the reward at step  $t + 1$ ;  $\alpha (= 0.01)$  is a fixed learning rate; and  $f_t$  is the number of action steps taken by the agent up till  $t$  which resulted in either (1) continuously maintaining the sensory array to be invariant or (2) the opposite (i.e., changing all the time). Thus, if  $s_t = s_{t-1}$  was true for the past  $n$  (= a large number) consecutive steps, then  $f_t = n$ , and this will increase the reward associated with  $(s, d)$ . On the other hand,  $n$  consecutive failures of maintaining sensory invariance will also lead to a high  $f_t$  value, but this time the reward for  $(s, d)$  will decrease. The reward function is simple but even such a simple rule is sufficient for the agent to learn sensorimotor associations.

In the following, we will present the learned policy  $\pi$  and the behavior of the agent which mimics the input stimulus.

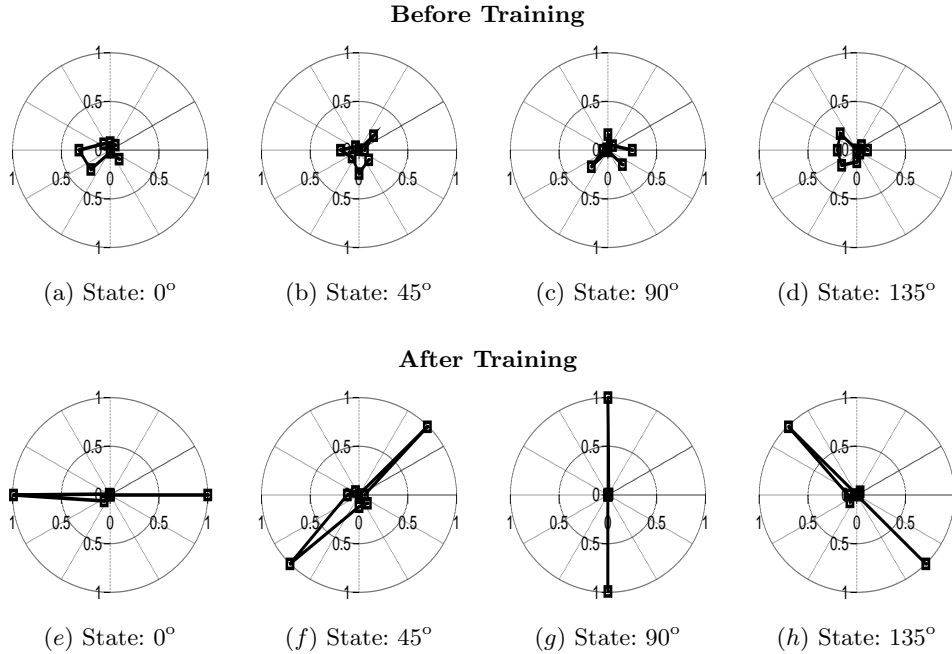
## 4 Experiments and Results



**Fig. 3. Inputs Used for Training and Testing** The agent was trained and tested on  $51 \times 51$  bitmap images each containing a 3-pixel wide oriented edge. Four inputs with four different orientations are used for the experiments (from the left:  $0^\circ$ ,  $45^\circ$ ,  $90^\circ$ , and  $135^\circ$ ).

In the learning process the agent interacted continuously with the visual environment in a series of episodes. During each episode, the agent was presented with a  $51 \times 51$  bitmap image containing an oriented edge (figure 3). The visual field of the agent was  $9 \times 9$  which can slide across the image. The visual field input was directly compared to each of the four sensory filters (also  $9 \times 9$  in size) and the sensory state  $s$  was set to a value  $\theta$  when there was an exact match with one of the four orientations  $\theta \in S$  (see equation 1).

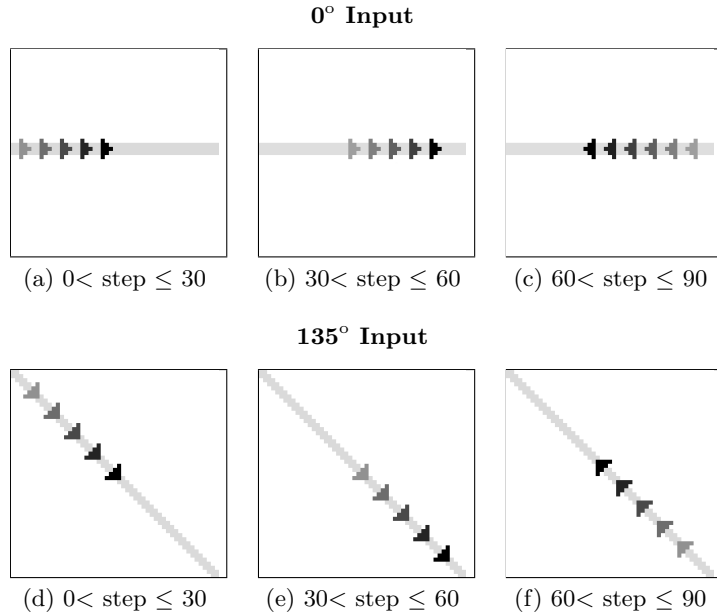
The agent was trained to learn the policy  $\pi : S \rightarrow D$  using equation 5 by going through the four different inputs. Since the size of the state and the action sets were  $|S| = 4$  and  $|D| = 8$ , the policy  $\pi$  and the associated rewards can be enumerated in a  $4 \times 8$  table. At each step, the next direction of motion  $d \in D$  (see equation 2) was determined based on the expected reward values stored



**Fig. 4. Reward Vector of Each Sensory State** The reward values of the four possible sensory states ( $0^\circ$ ,  $45^\circ$ ,  $90^\circ$ , and  $135^\circ$ ) are shown in polar coordinates. The top row from (a) to (d) are before training, and the bottom row from (e) to (h) are reward values after training. In each plot, for each point  $(\theta, \delta)$ , the angle  $\theta$  represents the direction  $d \in D$  of the visual field movement (there are 8 possible directions), and the distance  $\delta$  from the origin represents the associated reward value given the current sensory state (shown below each plot). The reward values were between 0 and 1. Initially, the rewards are randomly assigned for each direction of motion for each sensory state. After the agent is trained, the reward values become maximal for the movement along the orientations that correspond to the input that triggered that sensory state.

in such a reward table of the agent. The reward table was initialized to hold uniformly distributed random numbers between 0 and 1. Note that the reward was limited to the range  $0 \leq r_t \leq 1$ . Figure 4a-d shows the initial reward values where each plot corresponds to a state  $s \in S$ , and each polar plot shows the reward  $r$  (distance from origin) for each action  $d \in D$  (angle) for the given state  $s$ .

The training was carried out until the agent was able to learn to maximize the reward by consistently meeting the sensory-invariance criterion. The training usually lasted for up to 500 steps for each input. The reward table after training is visualized in figure 4e-h. The results clearly show that the agent learned to associate motion  $d$  which reflects (or mimics) the actual orientation of the environmental input to the current sensory state  $s$  triggered by that input. For



**Fig. 5. Behavior of the Agent after Training** Each plot shows a snapshot of 30 steps of movement of the agent’s visual field in the  $51 \times 51$  scene (only every 6 steps are shown). The triangles indicate the location of the visual field in the scene and their grayscale values represent the simulation step (black is the most recent step). The light gray lines in the background show the oriented input edges. Two simulation runs are shown here: (a) to (c) are for  $0^\circ$  input and (d) to (f) are for  $135^\circ$ . The trained agent successfully generates motion sequence to trace the input in both runs based on its sensor state and policy  $\pi$ . For example, in (b) the agent starts in the center and moves right, and bounces back when it reaches the end of the input (c).

example, in figure 4f, the maximum reward values associated with the sensory state  $s = 45^\circ$  are  $d = 45^\circ$  and  $d = 225^\circ$ , indicating a preference for a back-and-forth movement along the  $45^\circ$  axis which exactly mimics the visual input. The same is true for all other states (figure 4e, g, and h).

One thing to note from the actual numerical reward values (not shown) is that there is a slight difference ( $\leq 0.01$ ) between reward values for the two opposite directions separated by  $180^\circ$  (e.g.,  $d = 45^\circ$  and  $d = 225^\circ$ ). The minor difference helps the agent to have an initial bias in the selection of the first movement, and to maintain a momentum to continuously follow along an orientation instead of rapidly oscillating between two opposite directions. Note that this desirable effect was not explicitly built in by us, but rather, emerged from the sensory-invariance driven learning rule.

In order to verify if our analysis of the reward table is accurate, the trained agent was tested with fixed oriented inputs and the resulting motor behavior was observed. Figure 5 shows the action sequence generated by the agent for two



different inputs with orientations  $0^\circ$  and  $135^\circ$ . The plots show the movement of the visual field of the agent in response to the given input. The results show that the action of the agent based on the learned reward table exactly reflects our analysis above: The agent, upon activation of a single orientation sensor, performs a movement mimicking the external input that triggered that sensor, thus assigning (in our interpretation) a *meaning* to the sensory neuron’s spike in terms of its own action.

## 5 Discussion and Future Work

The main contribution of our work is the realization that a sensorimotor agent can find the meaning of its sensory state within its own actions, but more importantly, that the objective of maintaining *on-going* sensory-invariance plays a key role in allowing the agent to autonomously discover this semantic link.

An important message implicit in our work is that invariance can be seen from a totally different perspective. Usually, invariance is seen as something that needs to be detected or picked up from the environment by the perceptual system (e.g., invariant feature detection in vision). However, our approach differs in that invariance is sought-after in the internal activity pattern and it is *internally enforced* through a well-choreographed action. We speculate that there may be a link between this kind of action-based neural invariance and invariant sensory features in the conventional sense. For example, an approaching object will expand as time flows (turning on a certain neuron), and the same kind of effect can be achieved through a forward motion (again turning on the same neuron). Thus, the meaning of that neuron firing can be understood in terms of the action that would turn on that neuron reliably (cf. Gibson’s work on ecological perception and detection of environmental invariances [16]). Thus, even without action, when that neuron turns on (i.e., object is approaching), the brain can infer that it is analogous to moving forward toward the object.

Bell [33] posed an interesting question regarding the perception-action cycle. To quote, “What quantity should a perception-action cycle system maximize, as a feed-forward channel might maximize its capacity?”, which is relevant in our context. This is an important question, and we believe our reward criterion of *maximizing on-going sensory invariance* can serve as a potential answer. As we have seen, such a criterion can be used to internally learn the meaning of sensory state which may be a very important function for a “perception-action cycle system” to possess.

One criticism we anticipate is that if the agent had a rich array of sensors, such as a 2D matrix of RGB pixels, then the properties of the visual environment can be easily recovered within the agent through unsupervised learning even without direct access to the outside world. However, this does not help solve the problem, because this rich information is only available at the very first stage of sensory processing. The next stage, and the stage following that, etc. only receive a more and more *encoded* version from the previous stage, just like the sensory array in our agent which receives only encoded spikes from the orientation-tuned

filters. Thus, the same difficulty can remain. Another related criticism may be that the learning criterion we proposed may not be applicable to all stages of sensory processing. For example, the retinal ganglion cells and the lateral geniculate nucleus (LGN) in the early visual pathway show a center-surround receptive field property. It is not easy to imagine any kind of action sequence that would possibly keep these neurons' activities invariant. Our response to this is that sensorimotor coupling does not seem to exist at such an early stage of sensory processing (see e.g. [34]), and thus we do not claim that our approach will work in this stage.

One potential limitation of our account is that our model implicitly assumes that the agent has direct knowledge about its own movement, upon which the meaning of the sensors are grounded. The work by Philipona et al. [14] and Pierce and Kuipers [20, 21] point into a direction where a possible resolution can be found. They showed that without any knowledge of the external world, physical properties of the environment can be learned through sensorimotor learning. Especially, Philipona et al. [14] observe that there are two classes of sensors, exteroceptive and proprioceptive. They observed that agents have complete control over proprioceptive sensors (i.e., they can exactly predict the values based on their actions), whereas the same is not true for exteroceptive sensors. Thus, action, and the closely tied proprioceptive sensors may provide a more direct knowledge (as we proposed in this paper) to the agent than other common sensors. Another point is that unlike perception which is highly underconstrained (the problem of inverse-optics), but action or movement is strictly constrained by the bodily limits (e.g., we cannot stretch beyond a certain point). Such strong constraints may provide a learning problem to the brain which is significantly easier than perceptual learning.

Can our approach be extended into other sensory modalities such as audition, somatic sense, olfaction, etc.? Our approach is general enough to be easily extended into certain modalities such as somatic sense (see e.g. [35]), but it cannot work very well in domains where there is not much correlation between action and the perceived sensory state, e.g., olfaction.<sup>2</sup> Here, it would be useful to mention a different kind of meaning, those that are related to reinforcement signals such as gustatory rewards. Rolls proposed in [37] that semantics of brain states can be either related to (1) sensorimotor skills (as we also proposed) or to (2) environmental reinforcement. Olfaction, gustation, and other such sensory modalities convey signals that are highly related to survival values, and thus they may take on a different kind of meaning.

The model presented here is decidedly simple to convey the essence of the problem, and as such, it can be extended in several directions. We would like to note that sensory invariance does not always have to be defined on a single neuron's activity. Any kind of pattern, be that spatial or temporal, can be attempted to be maintained invariant while performing an action. Thus, meaning based on

---

<sup>2</sup> There is however some evidence that the act of sniffing can alter the perceived sense of smell [36], which indicates that our approach may also be applicable in the olfactory domain.

action can also be ascribed to a repeating pattern of activity, not just to a single spike. Also, invariance can be maintained in only one part of the global pattern, which can be seen as a relaxed version of the invariance criterion. (Attentional mechanisms may be necessary for this purpose [38].) We believe investigating in this direction will be most fruitful, and in fact we are currently steering our effort into this direction.

Finally, we would like to re-emphasize that the problem of meaning we raised in this paper is not only a central issue in autonomous agent or neuroscience research (cf. [2]), but also in information technology (IT) in general. The current information technology is mostly syntax-driven, and there is not much provision for autonomous semantics: At the end of the day, the entities that assign meaning to the meaningless symbols are us, humans [39]. This is becoming a serious problem because of the rapid growth in the amount and rate of data, since we humans no longer have sufficient time to attach meaning to the continuous stream of data. The problem seems to be that current IT systems are *passive* processors of information. As we have seen in this paper, *activeness* and *action* is key to autonomous understanding, thus, exploring how and in what manner can we make IT systems to be active may allow us to create major breakthroughs for the future IT.

## 6 Conclusion

From the realization that neural decoding methods requiring direct knowledge of the stimulus pose a problem when viewed from within the brain, we derived a novel solution to the problem of learning the meaning of sensory states, i.e., through sensorimotor learning based on on-going sensory invariance. We believe that the insight developed in this paper can help build a more autonomous agent with a semantics grounded on its own sensorimotor capacity, for its own sake. Such agents with autonomous understanding will be necessary for a major breakthrough in the future of information technology.

## Acknowledgments

We would like to thank Kuncara A. Suksadadi for his help with initial simulations. We also thank members of NIL and Cognoscenti at Texas A&M for lively discussions. Insightful comments by Benjamin Kuipers, Un Y. Nahm, Ricardo Gutierrez-Osuna, and Ronnie Ward are also gratefully acknowledged. The comments by two anonymous reviewers also helped clarify a couple of points, for which we are grateful. This research was supported in part by the Texas Higher Education Coordinating Board ARP/ATP program grant 000512-0217-2001.

## References

1. Freeman, W.J.: A neurobiological theory of meaning in perception. In: Proceedings of the International Joint Conference on Neural Networks, IEEE (2003) 1373–1378

2. Freeman, W.J.: How Brains Make Up Their Minds. Wiedenfeld and Nicolson Ltd., London, UK (1999) Reprinted by Columbia University Press (2001).
3. Williams, R.W., Herrup, K.: The control of neuron number. *Annual Review of Neuroscience* **11** (1988) 423–453
4. Rieke, F., Warland, D., de Ruter van Steveninck, R., Bialek, W.: *Spikes: Exploring the Neural Code*. 1st edn. MIT Press, Cambridge, MA (1997)
5. Adrian, E.D.: The impulses produced by sensory nerve endings. *Journal of Physiology (London)* **61** (1926) 49–72
6. Hubel, D.H., Wiesel, T.N.: Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *Journal of Physiology (London)* **160** (1962) 106–154
7. Warland, D., Meister, M.: Multi-neuronal firing patterns among retinal ganglion cells encode spatial information. *Investigative Ophthalmology Vision Science Supplement* **36** (1995) 932
8. Oram, M.W., Földiák, P., Perrett, D.I., Sengpiel, F.: The ‘ideal homunculus’: Decoding neural population signals. *Trends in Neuroscience* **21** (1998) 259–265
9. Hacker, P.: Languages, minds and brain. In Blakemore, C., Greenfield, S., eds.: *Mindwaves: Thoughts on Intelligence, Identity, and Consciousness*. Basil Blackwell, Oxford, UK (1987) 485–505
10. Jepson, A.D., Feldman, J.: A biased view of perceivers: Commentary on ‘observer theory, bayes theory, and psychophysics.’ In Knill, D.C., Richards, W., eds.: *Perception as Bayesian Inference*. Cambridge University Press (1996) 229–235
11. Arbib, M.A., Conklin, E.J., Hill, J.: *From Schema Theory to Language*. Oxford University Press, Oxford, UK; New York (1987)
12. Arbib, M.A.: *The Metaphorical Brain 2: Neural Networks and Beyond*. Wiley, New York (1989)
13. O’Regan, J.K., Noë, A.: A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* **24(5)** (2001) 883–917
14. Philipona, D., O’Regan, J.K., Nadal, J.P.: Is there something out there? Inferring space from sensorimotor dependencies. *Neural Computation* **15** (2003) 2029–2050
15. Hurley, S.: Perception and action: Alternative views. *Synthese* **129** (2001) 3–40
16. Gibson, J.J.: *The Perception of the Visual World*. Houghton Mifflin, Boston (1950)
17. Brooks, R.A.: Intelligence without representation. *Artificial Intelligence* **47** (1991) 139–159
18. Pfeifer, R., Scheier, C.: Sensory-motor coordination: The metaphor and beyond. *Robotics and Autonomous Systems* **20** (1997) 157–178
19. Pfeifer, R., Scheier, C.: *Understanding Intelligence*. The MIT Press, Cambridge, MA (1999)
20. Pierce, D.M.: *Map Learning with Uninterpreted Sensors and Effectors*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin, Austin, TX (1995)
21. Pierce, D.M., Kuipers, B.J.: Map learning with uninterpreted sensors and effectors. *Artificial Intelligence* **92** (1997) 162–227
22. Beer, R.D.: Dynamical approaches to cognitive science. *Trends in Cognitive Sciences* **4** (2000) 91–99
23. Cohen, P.R., Beal, C.R.: Natural semantics for a mobile robot. Technical Report 00-59, University of Massachusetts, Department of Computer Science (2000)
24. Arbib, M.A.: Language evolution: The mirror system hypothesis. In Arbib, M.A., ed.: *The Handbook of Brain Theory and Neural Networks*. 2nd edn. MIT Press, Cambridge, MA (2003) 606–611

25. Rizzolatti, G., Fogassi, L., Gallese, V.: Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature Reviews Neuroscience* **2** (2001) 661–670
26. Hecht, H., Vogt, S., Prinz, W.: Motor learning enhances perceptual judgment: A case for action-perception transfer. *Psychological Research* **65** (2001) 3–14
27. Bach y Rita, P.: *Brain Mechanisms in Sensory Substitution*. Academic Press, New York, NY (1972)
28. Bach y Rita, P.: Tactile vision substitution: Past and future. *International Journal of Neuroscience* **19** (1983) 29–36
29. Llinás, R.R.: *I of the Vortex*. The MIT Press, Cambridge, MA (2001)
30. Taylor, J.G.: *The Race for Consciousness*. MIT Press, Cambridge, MA (1999)
31. Sutton, R.S., Barto, A.G.: *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA (1998)
32. Mitchell, T.M.: *Machine Learning*. McGraw-Hill (1997)
33. Bell, A.J.: Levels and loops: The future of artificial intelligence and neuroscience. *Philosophical Transactions of the Royal Society of London* **354** (1999) 2013–2020
34. Fuster, J.M.: The prefrontal cortex – An update: Time is of the essence. *Neuron* **30** (2001) 319–333
35. Naito, E., Roland, P.E., Ehrsson, H.H.: I felt my hand moving: A new role of the primary motor cortex in somatic perception of limb movement. *Neuron* **36** (2002) 979–988
36. Bensafi, M., Porter, J., Pouliot, S., Mainland, J., Johnson, B., Zelano, C., Young, N., Bremner, E., Aframian, D., Khan, R., Sobel, N.: Olfactomotor activity during imagery mimics that during perception. *Nature Neuroscience* **6** (2003) 1142–1144
37. Rolls, E.T.: Representation in the brain. *Synthese* **129** (2001) 153–171
38. LaBerge, D.: *Attentional Processing*. Harvard University Press, Cambridge, MA (1995)
39. Searle, J.R.: Minds, brains and programs. *Behavioral and Brain Sciences* **3** (1980)