SENSORY INVARIANCE DRIVEN ACTION (SIDA) FRAMEWORK FOR UNDERSTANDING THE MEANING OF NEURAL SPIKES

A Thesis

by

SARVANI KUMAR BHAMIDIPATI

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

May 2004

Major Subject: Computer Science

SENSORY INVARIANCE DRIVEN ACTION (SIDA) FRAMEWORK FOR UNDERSTANDING THE MEANING OF NEURAL SPIKES

A Thesis

by

SARVANI KUMAR BHAMIDIPATI

Submitted to Texas A&M University
in partial fulfillment of the requirements
for the degree of

MASTER OF SCIENCE

Approved as to style and content by:

Yoonsuck Choe (Chair of Committee)	Donald K. Friesen (Member)
Reza Langari (Member)	Valerie E. Taylor (Head of Department)

May 2004

Major Subject: Computer Science

ABSTRACT

Sensory Invariance Driven Action (SIDA) Framework

for Understanding the Meaning of Neural Spikes. (May 2004)

Sarvani Kumar Bhamidipati, B.E.,

Gujarat University, India

Chair of Advisory Committee: Dr.Yoonsuck Choe

What does the spike of a sensory neuron mean? This is a fundamental question in computational neuroscience. Conventional approaches provide an answer based on correlation between spike pattern and the stimulus that caused it. However, these approaches do not satisfactorily explain how the brain, which does not have direct knowledge of the world or the stimuli, can achieve this task. This thesis frames the problem in terms of a task for a simulated agent and provides a solution based on an approach which regards action as necessary for acquiring the meaning of neural spikes. This approach differs from some others in that it proposes a new criterion called the sensory invariance criterion, which can be used to associate meaning to spike patterns in terms of action sequences the agent generates. This criterion forms the basis of the Sensory Invariance Driven Action (SIDA) framework presented in this thesis. This framework is implemented in a reinforcement learning agent and the results indicate that the agent can successfully learn to associate meaning to the sensor activity in terms of specific actions which reflect the properties of the stimulus. Further behavioral experiments on the agent show that this framework allows the agent to learn the meaning of complex (spatio-temporal) spike patterns. The successful learning exhibited by the agent raises hopes that SIDA can be used to build agents with natural semantics.

To my parents Mohan Sarma and Padmavathi

ACKNOWLEDGMENTS

Thanks are foremost due to my advisor, Dr. Yoonsuck Choe. His insight and guidance have been instrumental in the successful completion of this thesis. I will be forever grateful for his enormous patience and support during my research and for treating me more as a friend than his student.

I would also like to thank Dr. Donald K. Friesen and Dr. Reza Langari, for their insightful comments during the presentation of my research and on the thesis draft, which made this thesis better.

Thanks are also due to Kuncara A. Sukasdadi for his help with initial simulations in this thesis.

I would further like to thank Subru, Yingwei and Sejong for helpful discussions during various stages of this research.

Most of all, I want to thank my parents and brother for their love and support throughout my life.

This research was supported in part by the Texas Higher Education Coordinating Board ARP/ATP program grant 000512-0217-2001.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

INTRODUCTION

Our perception of the world around us is dependent on how the brain interprets the neuronal spikes [25]. Spikes are electrochemical pulses that are generated in the sensory neurons in response to external stimuli. Spikes are used as a means by which the brain is able to perform its tasks, e.g., sending signals to motor neurons or to neurons in other parts of the brain that perform diverse functions. With over 100 billion neurons in the brain, a huge number of spikes are continuously and simultaneously generated in the brain. This huge number of spikes form complex spike patterns. The brain is continuously interpreting, or associating meaning to, these spike patterns to understand the world. The task faced by the brain is enormous, and understanding this process in the brain is a major challenge in computational neuroscience.

To understand how the brain associates meanings to spike patterns, general neuroscience approaches begin by asking a much simpler and fundamental question: How does the brain understand the meaning of a single spike? [25]. Even this simple question took a lot of time to answer, beginning with understanding of how spikes are generated [1], to dynamics of spike propagation [15], to how the spikes are generated in response to individual components of sensory input [16], to techniques which can reconstruct stimulus given only spikes [25]. In more recent years, mathematical and statistical techniques have been used to provide an explanation of how spikes are interpreted, e.g., Bayes methods [25], population vector method [11], etc.

This thesis will look at some of these traditional approaches and models and identify the drawbacks in them. Then it will look at an alternative approach called the

---

The journal model is *IEEE Transactions on Neural Networks.*

sensorimotor approach [17], [21], which can be applied to understand the meaning of spikes. Based on this approach this thesis will provide a framework for understanding the meaning of spikes by the brain through actions. This framework can be used by natural as well as artificial agents to obtain meaning of their own internal spikes in terms of actions they perform.

In the rest of the chapter, a general overview of the problem of understanding the meaning of the spikes is first provided (section 1.A), followed by an introduction to sensorimotor approach used to solve the problem (section 1.B).

A. Problem Overview

For a better understanding of the problem addressed in this thesis, a quick review of a basic nervous system is first provided. Later the concept of internal and external observers are introduced which will explain the difference in focus of the traditional approaches and the alternative approach taken in this thesis.

1. How Does the Nervous System Work?

To better understand the problem addressed in this thesis, let us take a quick review of how the nervous system functions. All the systems in an organism need to interact and coordinate with each other for the proper functioning of the organism. The nervous system provides a mechanism for such an interaction. It is made of neurons connected to each other carrying signals which coordinate the functioning of various parts of the body. The neuron (Fig. 1) constitutes the basic unit of the nervous system. They are like any other cell in the body except that they are specialized to conduct signals. For this purpose they have special extensions called dendrites and axons. Dendrites are short hair like extensions on the neuron body through which the neuron receives signals from other neurons. On the other hand, axons are long hair-like extensions that carry the signal to other neurons. These signals are in the form of electrochemical pulses and are called *spikes* (or action potentials). The cells that

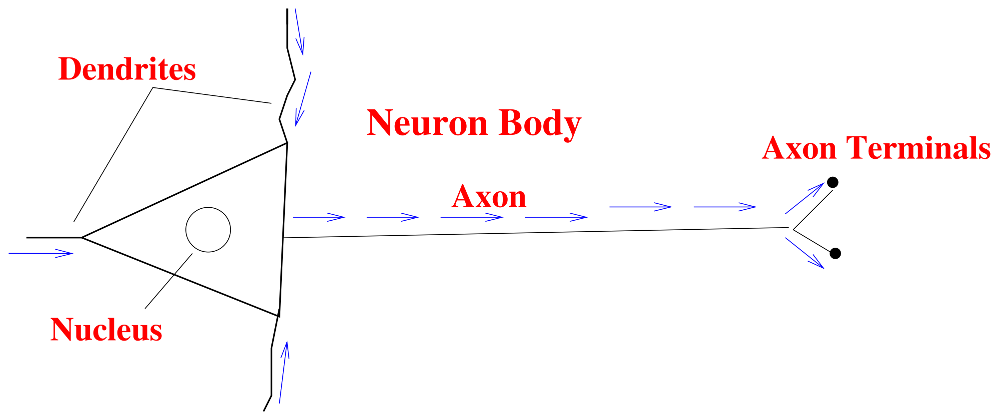

Fig. 1. Neuron. The neuron is the basic functional unit of the brain. The main cell body of the neuron has hair like appendages called the dendrites. The dendrites receive input spikes from other neurons. The neuron also has a long extension called an axon, through which the spikes are sent out of the neuron, which will be picked up by latter neurons in the neural pathway. The neuron generates a spike (action potential) when it reaches a firing threshold. The arrows in the figure indicate the direction of propagation of spikes through the neuron.

make up the receptors of our sensory faculties such as the photoreceptors in the retina of the eyes, or the mechanorecepters below the epithelial tissue (i.e., skin), generate spikes in response to external stimuli like light or touch, respectively [1]. These spikes travel to the downstream neurons thus causing internal stimulation in the latter. This stimulation generates spikes in these neurons which in turn, travel to the next neurons, thus setting off a chain of spikes. In this manner, the spikes originating from the sensory receptors keep on reproducing across neurons until receptors reach the brain. In the brain, the spikes originating from various sensory neurons combine together to form complex spike patterns. The brain interprets these spike patterns to understand the stimulus. The understanding of this process of interpretation forms the basis of the problem addressed in this thesis.

2. Analysis of the Problem

As already mentioned, a fundamental question in neuroscience is understanding what is the meaning of a neural spike. The traditional methods revolve around associating the spike patterns with the stimulus causing them. These methods have been successful in characterizing neural spike properties and predicting the stimulus responsible for a spike, to a high degree of accuracy. In these methods, the experimenter systematically changes the input stimulus and observes the spike patterns generated. By examining the spike and stimulus pairs over a period of time, the experimenter is able to establish the correlation between the spike patterns and the stimulus, so that the stimulus can be attached as a meaning to the spike patterns. Using Bayes method mentioned earlier [25], this process can be written as,

$$P(s|t) = \frac{P(t|s)P(s)}{P(t)},$$

where s is the stimulus and t is the spike train. The likelihood term $P(t|s)$ requires that the brain already has some previous knowledge of the stimulus-to-spike translation. In order to acquire this knowledge, previous knowledge of the stimulus properties is required [13], [8].

However, the question is how did the neurons in the brain acquire this previous knowledge. Suppose we pose the question to a neuron "How do you know what meaning to attach to a spike that you receive from your preceding neuron? After all as a neuron you have no knowledge about the stimulus in the outside world." Since the neuron, unlike the experimenter in the traditional approaches, cannot have any direct knowledge of the outside stimulus, it cannot associate a meaning to the spike pattern using these techniques.

To better illustrate the distinction between the two views, let us examine the concept of external vs. internal observer. Consider a natural agent who is able to view a stimulus. As a result of this viewing some spike patterns are generated in it. In order to understand the spike patterns, the agent will attempt to interpret these spike patterns, and we would like to understand how these are interpreted. To accomplish this task, traditional approaches use the method of external observer (see Fig. 2(a)).

In this figure, the natural agent is represented as a closed box. This agent is able to view the input I, which is encoded through the nervous system (function f) before it reaches the brain. The brain thus receives a heavily encoded spike pattern S, which is the result of the stimulus I. In the case of the external observer, the experimenter is sitting outside and is able to view both the external input I and the spike S generated in the brain. The external observer now correlates the spike patterns and the input stimuli and is able to attach the meaning of a particular spike to be a given stimuli. This method, indeed allows the observer to know the meaning of spike.

Contrast this with the problem faced by an internal observer (Fig. 2(b)). The experimenter is sitting inside the box thus becoming an internal observer. This observer is never directly able to view the external input I. This observer can only view the spikes S, which are heavily encoded through the function f. Since this internal observer does not have any knowledge of the external world, there is no data to correlate the stimulus and the spike to understand the meaning of the spike. So how can this observer be able to associate a meaning to the spike pattern?

In the real world, the brain is the internal observer and it faces the same problem described above. However, the brain seems to be fully capable of associating the

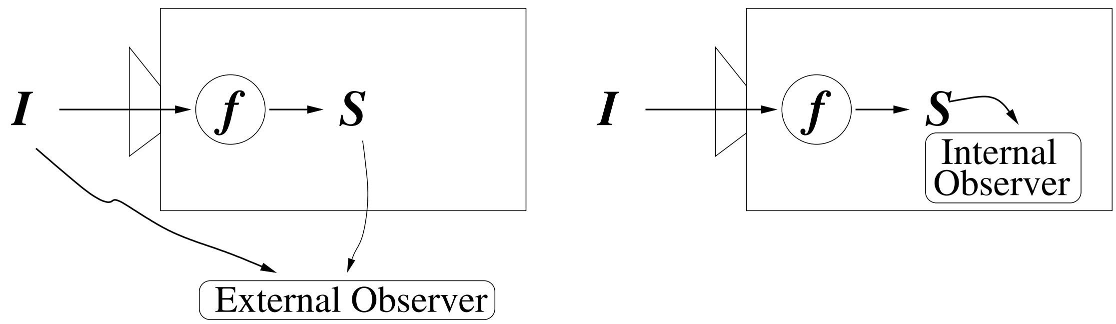

(a) External observer

(b) Internal observer

Fig. 2. 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 photoreceptors). The task then is to find out the property of input I given just the spikes S.

meaning to the spikes, based on its observation of spikes alone. How is the brain making this possible? This thesis proposes a solution to this problem of understanding how the brain, as an internal observer, is able to associate meaning to the spike patterns, which will be introduced next.

B. Approach

To explain how the internal observer is able to understand what the spike stands for, we use the sensorimotor approach. Such an approach used by Philipona et al. [22] and Pierce and Kuipers [23], [24] is based on the relationship between sensory activity in the organism and the motor actions performed by it. The sensorimotor approach suggests that action is inseparable from perception because the act of perception on the part of an organism cannot be independent of the behavior of the organism (e.g., [17], [19]). For example, if we want to perceive the smoothness of a surface, intuitively we slide our hand across it. Thus this action of moving our hand works together with the tactile perception and can be viewed as aiding the process of perception. There is enough experimental evidence to show that action influences the perception process [3], [26], some of which are discussed later.

The inclusion of action into the process of understanding the world gives rise to questions like: how to relate action and perception? What criterion can be used to establish the relationship between motor activity of the organism and its sensory activity arising due to perception and so on. In this thesis, we explore the sensorimotor approach to provide a framework called the Sensory Invariance Driven Action (SIDA) framework which underlies the mechanism for relating the sensory activity and the motor activity. This framework uses sensory invariance as criterion for establishing relationship between sensory activity and motor activity. The framework and the criterion are discussed in detail in later chapters.

In summary, the goals of this thesis are to (1) suggest an alternative approach which can be used to explain how the brain is able to associate a meaning to the neuronal spike patterns, (2) provide a framework and a criterion for implementing such a approach into a natural agent, (3) prove the validity and applications of such an approach by performing computational experiments.

C. Outline of the Thesis

So far, we have identified that there has to be a mechanism through which the brain understands the meaning of spikes it receives. In chapter 2, we look at some of the previous methods and approaches to understand the meaning of spikes. In chapter 3, I propose the Sensory Invariance Driven Action (SIDA) framework to understand the meaning of spikes [8]. Here I will show how the sensory invariance can be used as a criterion for SIDA framework and present the algorithm to implement SIDA in a sensorimotor agent. Chapter 4 will present the results of the training and behavioral experiments performed using SIDA. Chapter 5 will contain discussion about the SIDA framework in relation to other approaches in the current literature and future work. The final chapter will will be the conclusion to this thesis.

CHAPTER II

BACKGROUND

The brain is the controlling unit of the organism and it is responsible for activities like perception, cognition, etc., in the organism. Understanding how it works is an important goal of computational neuroscience. The discovery that the brain was made up of neurons [7], and further experiments revealed that there is continuous neuronal activity in the brain. This led to the question, how does the neuronal activity help us in perception? [25]. Möller, Helmholtz, and others (see [25]), suggested that neuronal activity was generated in response to the stimulus in the world based on their observations of neuronal activity. Later, Adrian [1] tested these ideas through his experiments and firmly established that *spikes* were generated in sensory neurons in response to stimuli and went ahead to show that the nature of the spikes was dependent on the features of the stimulus. At a microscopic level, researchers were able to learn how spikes are generated [15], how spikes depend on neuron properties (e.g., feature selectivity: Hartline and Ratliff [14], Barlow [5], Hubel and Wiesel [16]), etc. However, at higher levels of understanding e.g., perception, cognition, etc., researchers are still trying to understand the process of how the spikes are interpreted by the brain.

In this chapter, some of the approaches and ideas which have been suggested to explain the process are analyzed. For the purpose of this thesis, these approaches have been categorized as conventional and sensorimotor approaches. This chapter will highlight (1) the limitations of conventional approaches in explaining the process of how the spikes are interpreted by the brain and, (2) how the sensorimotor approach has been used to overcome some of these limitations.

A. Conventional Approach

The conventional approaches evolve from the idea that perception and generation of behavior are functions of two separate subsystems [19]. During perception the brain interprets the sensory spikes into some internal representation which is then used by a separate subsystem to generate behavior. As a result, the process of action generation is regarded to be driven by perception process. This idea is based on the work of Marr on vision (as discussed in [19]). Because of this idea of independence of perception and action, the neuroscientists seeking to understand the process of perception performed experiments which did not take into account the behavior of the organism in response to stimulus. The drawbacks of this approach are discussed later in the section. Here we will briefly introduce two popular methods which follow this approach.

1. Population Vector (PV) Method

Georgopoulos et al. introduced the first method called the PV method [11], [20] (see Fig. 3) in the early 1980s. They were trying to understand representation of movement directions in primate motor cortex. They conducted experiments to understand neuronal activity in M1 cortex of the monkey brain as the monkey performed some specific motor actions e.g., moving its hand. Based on the results of the these experiments, Georgopoulos et al. suggested the population vector method, which took into account the contribution made by each individual neuron to the spike as it propagated through the neuronal pathway. For example, if the hand of the monkey was moved in a particular, say left to right direction, several sensory neurons in the body of the monkey get activated. These group of neurons are termed as population of neurons The spikes from this population of neurons propagate through the neural pathways to reach the brain where they result in a complex spike activity. According to Georgopoulos et al. the PV method can be used to calculate this activity by taking into account the contribution made by each of these neurons. The contribution of each neuron is related to the amount of activity in that neuron. The PV method calculates the resultant spike activity by doing a vector addition of each neuron's activity. Thus the PV method provided a mathematical model to explain how the spikes can be interpreted.

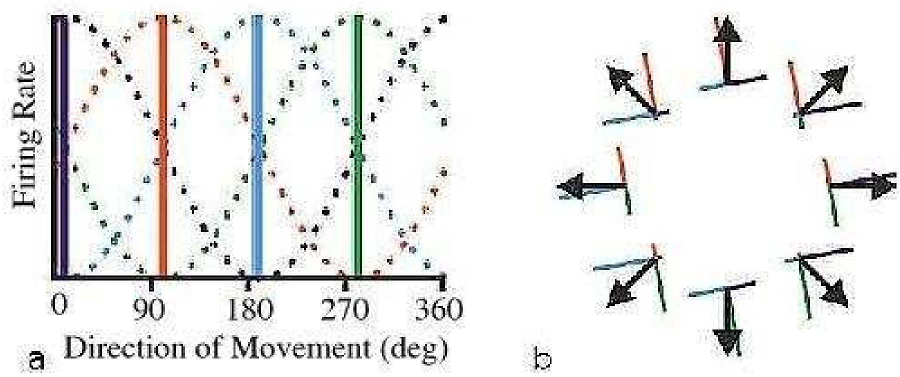

Fig. 3. Population vector method. The construction of Population Vector from the firing rates of cells. (a) The activities of motor cortical neurons tend to be broadly tuned to the direction of movement (dashed lines). The solid vertical lines denote the preferred direction (PD) of each cell, the movement direction for which that cell is maximally active. (b) For each movement direction, the activity of each cell scales the length of vector aligned with its PD, and the vector sum of all cells defines a population vector. It has been shown that this population vector will predict the direction of motion of movement if three conditions are met. First, neural activity is broadly tuned to the direction of movement; second, the PDs of the cells are uniformly distributed in space; and third, there is no systematic relationship between a cell's discharge rate and its preferred direction. (Adapted from [27]).

2. Bayesian Method

In the brain, the spike patterns in response to a stimulus are seldom the same as they were when the same stimulus was presented some time ago. In fact the spike patterns can vary by small amounts for the same stimulus. As a result, a certain amount of randomness is introduced into the observations of spikes. In order to take into account these fluctuations, probability concepts were introduced into the models explaining the spike activities. Several models based on various probability techniques exist, however here we will discuss models based the popular Bayes method. The description in this section follows that of [25].

In Bayes method, the experimenter who seeks to understand the relationship between the spike and stimulus begins by presenting a known stimulus $s(t)$ to the natural agent over a length of time. In response to the stimuli spike patterns are generated in the brain of the natural agent. However, the resulting spike patterns are slightly different from each other. In order to identify specific spike patterns the experimenter builds a probability distribution of the spike patterns generated in response to stimulus $s(t)$. This distribution is given as $P[\{t_i\}]$ where $\{t_i\}$ is a set of arrival times of spikes.

In the real world however, the stimulus is not always constant and varies with the world around. This varying stimuli results in spike trains that are indiscernible from each other. This means we have a varying stimuli $s(t)$, but the same spike pattern. Hence, to find any association of stimuli $s(t)$ and neural response, we must first know the probability of a stimuli occurring. This is given by the probability distribution $P[\{s(t)\}]$.

The experimenter who wants to build a unique relationship between the spike and stimulus has to rely on two conditional probabilities to achieve his goal: (1)

$P[\{t_i\}, s(t)]$, the probability of specific spike occurring given a fixed stimulus, and (2) $P[s(t)|\{t_i\}]$, the probability of given spike arising in response to a specific stimulus. In the probabilistic context these two conditional probabilities are decompositions of the same joint probability distribution and hence related to each other:

$$P[s(t)|\{t_i\}] * P[\{t_i\}] = P[\{t_i\}|s(t)] * P[s(t)] \tag{2.1}$$

$$\Rightarrow P[s(t)|\{t_i\}] = P[\{t_i\}|s(t)] * P[s(t)]/P[\{t_i\}] \tag{2.2}$$

This relation is the Bayes theorem for inferring the stimulus when a particular spike is observed. (For detailed review of the Bayes method refer to Rieke et al. [25]).

Bayes methods depend on conditional probability distribution estimates of spikes for a stimulus, and require prior probability distribution of the stimuli to estimate the stimuli. When the prior probabilities of the stimuli are known, Bayes methods are very useful in calculating the probability of the stimulus for a given spike. For example, Oram et al. [20] have shown that Bayes method can be effectively used to decode neural population signals effectively to very high degree of accuracy.

3. Evaluation of Conventional Approaches

The PV method and the Bayes method described above are based on concrete mathematical models and they are supported by several experimental results which establish that these methods are successful in characterizing spike patterns, and predicting the stimulus responsible for spikes. However, in spite of these advantages, there are certain requirements in these methods which limit their usage while explaining how the spikes are understood by the brain. If we take a closer look at either of the methods described above, it is obvious that in order to achieve the goal of understanding the meaning of a spike, some form of prior knowledge about the world is required. For example, in the case of PV method, in order to calculate the contribution of each neuron, prior knowledge about the properties of the neurons need to be known so that the contribution made by the neuron to a particular stimulus (e.g., movement of had in a particular direction) is known. Similarly in the case of the Bayes method, the prior probability distributions of the stimuli need to be known in advance to build the conditional probability statistics required in Bayes relation.

An *experimenter* can learn of these prior requirements by studying the stimuli and spike patterns because he has access to both the internal and external world, where the stimuli and the spike patterns are. The experimenter's view is thus the view of the external observer (refer to section 1.A.2, Fig. 2). However, in direct contrast to the experimenter is the brain, which does not have any knowledge of the outside world. All it has access to is the spike activity. Without access to the stimuli and any direct knowledge of the world, like the experimenter does, the brain is unable to use the techniques above to interpret the spikes.

However, the brain is able to associate the meaning of spikes in such a way that we are aware of world around us. Obviously, there is a method by which the brain is able to perform the task of associating meaning to spikes. Many researchers recognize the limitation of conventional approaches and are actively seeking alternative approaches which can explain this ability of the brain. In the following section we look at some of the alternative approaches called the sensorimotor approaches that offer an explanation to how meaning can be attached to spikes.

B. Sensorimotor Approach

Conventional approaches are based on the idea that perception and behavior of the organism, i.e., the actions generated by the organism, are independent and governed by separate subsystems. As a result they disregard the role of action while trying to understand the perception process. The Bayes method and the PV method described above are based on such approaches. We have seen that, owing to the dependency of these methods on prior information about the world, they cannot provide satisfactory explanation as to how the brain, which does not have direct knowledge of the world can interpret neuronal spikes. As a result, researchers started questioning the idea of perception and action being separate. Marr's idea which is representative of conventional approaches (see [19]), claims that brain has to make correct inference of the spikes so that the proper behavior of the agent can be generated, but researchers supporting the interdependence of action and perception ask, how can the inference be determined as correct if the behavior cannot influence perception (Dreyfus [10], [19])? Hence, many researchers have started investigating the influence of action on perception on each other and see if this approach can better explain how the meaning of spikes can be interpreted.

This approach linked the sensory activity in the brain to the actions generated by the agent. The approach realizes that just as the motor actions performed by a natural agent are driven by the sensory activity it experiences, the sensory activity it experiences as part of the perception process is dependent on the motor activity performed by the agent. For example, if the influence of the action on perception is not taken into account, the perception of a train approaching a stationary observer will be the same as the perception of an observer moving towards a stationary train. However, when the action of the observer is taken into account, i.e., whether he is stationary or moving, accurate perception of whether the train is moving or stationary can be made. This approach which takes into account the interdependence of perception and action is generally called the *sensorimotor approach* [22]. In the rest of the chapter I will describe experiments and theories which support the sensorimotor approach to explain the perception process.

1. Experimental Evidence of Sensorimotor Interactions

This section describes the 'Sensory substitution' experiment which clearly demonstrated the influence of action on the perception process. Later in the section I will discuss 'mirror neurons' which provides biological evidence supporting such an influence.

In 1983, Bach y Rita performed an experiment called the 'Sensory substitution experiment' [3], [4], which strongly suggested that the action performed by an individual can alter the perceptual experience. In the 'Sensory substitution experiment', the influence of action by a subject on his own perceptual process was analyzed. The subject was a blind man from birth and hence did not have any idea of the visual world. A video camera was attached to his head, which could record the visual world around the subject. The output of the camera was attached to a dot-matrix tactile array, which is a device analogous to a dot-matrix printer, that generates a tactile pattern on to the subject's back. This tactile pattern that is generated on the subject's body gives him the feeling of touch. In the first part of the experiment, the subject was made to sit passively. He was not allowed to move while an object was moved around in the visual scene. The camera recorded this activity and sent signals to the tactile array. The tactile array in turn generated a tactile pattern onto the back of the subject. At this stage of the experiment, the subject's description of his sensory feeling was that of touch, which was expected.

However, during the second part of the experiment, the subject was allowed to be active, i.e., he was allowed to move his head around. Once again an object was placed in front of him. As he moved his head, the video camera comes across the object, and sends a signal to the tactile array. This tactile array generated a tactile pattern on the back of the subject. Surprisingly, this time around the subject's sensory experience was different from touch. His description of the feeling lead the researchers to believe that he was experiencing a sense of sight. For some reason, when the subject was able to perform an action by moving his head, the sensory patterns which were generated as a result of touch, were interpreted in a manner which was similar to how an normal subject would interpret sensory patterns generated when he sees the object in the visual scene.

This experiment strongly suggests that the action generated by the subject himself aids in his perceptual experience. It means that the action performed by the subject influenced his brain or a part of the brain that is responsible for perception, into interpreting the sensory patterns as something different from tactile sense. Thus, from this experiment it can be inferred that action somehow influences the perception process.

Another evidence which supports the claim that action influences perception is the discovery of mirror neurons in the brains of primates. Mirror neurons are special neurons which fire when an action is performed by the subject itself or when the actions are observed in some other subject [26], [2]. They generate a similar type of activity during the perception of action, performed by another subject as they do when the subject performs the action itself. Because of this behavior, mirror neurons are thought to fill the gap in the relationship between action and perception. This biological mechanism relating action and perception is viewed to support the idea that action and perception are closely related and they should be treated as a whole while understanding the meaning of spikes in the brain.

2. Theories on Sensorimotor Interaction

There are several theories which support the idea that action and perception work together. These theories may vary in opinion regarding how the action and perception work together, however all of them support the sensorimotor approach, i.e., perception of a natural agent and the behavior of the agent are related and influence each other. Here we will briefly present a few of the famous approaches.

Gibson, in his book on ecological perception [12] explained how action can serve as a feedback mechanism for the perception process. He explains that the process of perception and generation of action is not one-way, but two-way, i.e., the generated action guides the process of perception. Thus action is both a cause and an effect of perception [17]. This action, however may not always be active; even a passive action is sufficient for the process of perception. It should be noted that though Gibson recognized the interdependence of action and perception, he still regarded them as separate subsystems which work together [17].

Susan Hurley [17], presented a new idea on interdependence of action and perception. She argues that action and perception are functions of interactions between input and output. She suggests that there is a strong interplay between action, perception, and the environment through 'dynamic singularity', i.e., through a sequence of causal process that continuously influence each other. The same causal inputs can generate different perceptions of the same environment owing to varying intentions. This interdependence of perception and action can be explained if they are considered to be just a subset on input and output interactions. This view is called the 'Two-level interdependence view'.

Möeller [19] presented an approach to behavior based perception called 'perception through anticipation', which claims that perception and action are two aspects of one neural process. In this approach the natural agent is able to perceive space and shape based on anticipation of the 'sensory consequences' of the actions performed by it, i.e., given a sensory state, the agent anticipates the influence of taking an action on its sensory state and generates further actions accordingly. In this manner, the action influences the perception of the agent.

The ‘sensorimotor contingency theory’ by O’Regan and Noe [21] views the process of visual perception as an act of exploration. This approach does away with the representational approach and suggests that the perception process is based on a set of rules called the sensorimotor contingencies. These contingencies are framed based on the knowledge of the natural agent about its actions. This approach claims that the ability of the agent to perceive the world is based on its ability to master these rules.

The approaches and theories described here, though might vary in explaining how the action and perception work together, they all support the idea of action and perception working together. This forms the basis of sensorimotor approach which is used in this thesis, to explain how the meaning of sensory spikes can be acquired by the brain.

C. Computational Systems Using Sensorimotor Approaches

In this thesis I will propose a framework by which a natural agent can understand the meaning of neuronal spikes. By understanding we mean the agent will be able to acquire the meaning and maintain the meaning of the spike patterns by itself. Such systems are said to be capable of learning natural semantics. Cohen and Beal [9] and Pierce and Kuipers [24, 23] have described similar systems which are able to learn about themselves and their world.

Cohen and Beal developed a robot that is able to learn about concepts and meanings by interacting with the environment. This robot was not only able to identify the objects like a chair, pit, etc. in the world but actually understood the significance of such objects. Through a process called *perceptual redescription* the robot was able to assign ‘tokens’ to describe its experiences with the objects, e.g., when it experiences a crash when it runs into a wall. The robot uses these tokens as meanings of the objects and used them to understand the world it is in. Another system using sensorimotor approach was built by Pierce and Kuipers. They developed a learning agent that was able to obtain the knowledge of the world it is in by learning about its own sensorimotor abilities. The learning agent achieved this task by learning to control a robot which has raw, uninterpreted sensors and effectors. The learning agent has no knowledge of these devices in the robot. Initially the agent generated commands unaware that they generated motion in the robot, and obtained feedback from the sensors unaware of their purpose. However, over a period of time the agent was able to learn about, (1) the motor actions in the robot which it could control, and (2) the sensors which were providing some form of feedback. With these abilities the agent was able to construct a hierarchical representation of the world around it.

D. Summary

In this chapter, I reviewed the body of work that gave motivation for the current thesis. Limitations in the conventional approach have been identified, and an alternative approach based on sensorimotor integration has been reviewed.

In the following chapters, a system that uses a sensorimotor approach not unlike the above, but based on the SIDA framework that can understand the meaning of spikes and thus the world will be presented.

CHAPTER III

APPROACH

In this work, I propose a framework that is based on the sensorimotor approach, called Sensory Invariance Driven Action (SIDA) for the understanding of the meaning of neuronal spikes. This framework suggests that *sensory invariance* can be used as a criterion for generating action that reflects sensory stimulus properties. Sensory invariance refers to a state during which the internal sensor pattern generated in response to a stimulus remains unaltered. According to SIDA, the action that is generated when sensory invariance is maintained can be viewed as the meaning of the sensor pattern or the neuronal spike pattern. In order to better explain the SIDA framework, a sensorimotor agent is introduced in the following section. This sensorimotor agent is an abstraction of a natural agent and lays out the basic components for the understanding of neuronal spikes.

A. A Sensorimotor Agent

A basic sensorimotor agent possesses at least one sensory faculty, e.g., vision, tactile, auditory, etc. In addition, sensorimotor agents also possess motor abilities that can be used to generate actions. Using the SIDA framework, a sensorimotor agent can be trained to associate a sensory pattern with an action. This association can be achieved using sensory invariance as a criterion and can help the agent to obtain the meaning of the sensory pattern. In the rest of the section, I will describe such a sensorimotor agent in detail.

A simple sensorimotor is shown in the Fig. 4. The agent is provided with a limited visual field, which acts as the *eye* of the agent. Using this visual field the agent can view various parts of the input scene presented to it. The agent has a rudimentary

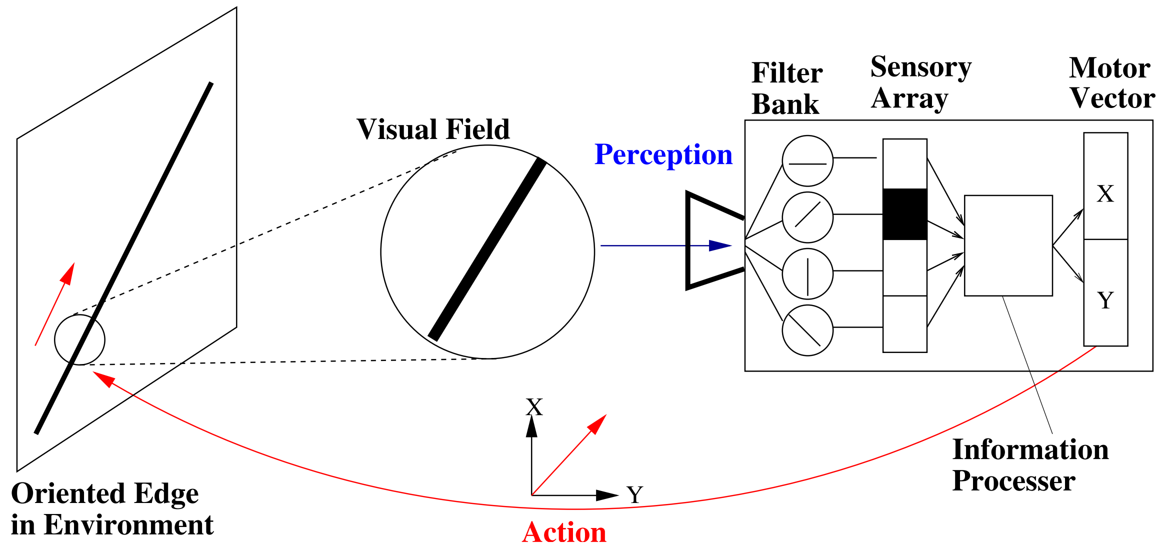

Fig. 4. A sensorimotor agent. An illustration of a simple sensorimotor agent is shown. The agent has a limited visual field where the input from the environment is projected. A set of orientation-tuned neurons receive that input and generate a pattern of activity in the sensory array (black marks active). In the situation here, the 45° sensor is turned on by the input. Based on the sensory array pattern, after some processing in the information processor, 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.

visual input processing ability: The agent has a bank of orientation filters and a sensory array. The filter bank and sensory array are connected back to back, so that each filter in the filter bank is connected to one sensor unit. These *orientation* filters in the filter bank are selective to specific orientations in the input, and when a preferred input is encountered, an activity is generated in the corresponding sensor unit. Such a combination of the filter bank and the sensor array is analogous to the circuitry in the primary visual cortex (V1) in the brain, such that the filter bank corresponds to the receptive fields of the visual cortical neurons and the sensory array activity to the V1 activity. It is known that the neurons in the V1 cortex are tuned to specific orientations of input [16] and spikes are generated by the V1 neurons, when the orientation of the input stimulus corresponds to their tuned orientation. Analogous to the functioning of these neurons, the sensory units in the sensor array generate activity when the orientation of the input matches the orientation of the attached filter in the filter bank. In Fig. 4 the sensory unit which corresponds to the filter whose orientation matches that of the input is shown to be active.

The sensorimotor agent is also equipped with motor abilities, which permit it to move its own visual field. By moving its visual field, the agent is able to view different areas of the input scene. The x and y variables in the figure correspond to the motor vectors which the agent can generate to move in the horizontal and vertical directions respectively. From a combination of these variables, the agent can move its visual field in a various directions in the x-y plane. The agent is aware of its ability to move its visual field and can generate motion independent of or in response to the activity in the sensory array.

For the rest of the chapter, this agent is used as a subject for better explanation of the problem being addressed in this thesis. The agent will then use the SIDA framework described later to solve the problem.

B. Meaning of Spikes in a Sensorimotor Agent

A simple visual stimulus is presented to the agent, like the oriented edge shown in the Fig. 4. The visual field of the agent conveys the input information to the combination of filter bank and sensory array, where it is processed. After processing, depending on the orientation of the input, a sensory pattern is generated in the sensory array. The task of the agent is to attach a meaning to this sensory pattern. In the case shown in Fig. 4, the agent has to understand that a spike pattern in which a single unit of sensory array is active corresponds to a certain orientation of the input. This task which the agent has to perform is modeled on the internal observer concept which was presented in chapter 1. Like the internal observer, the sensorimotor agent here does not have any direct knowledge of the stimulus presented to it. The agent is solely dependent on the spike patterns it can observe to understand the meaning of spikes.

Given this situation where the agent does not have any knowledge of the stimulus, it seems that the only option that the agent has, to understand the meaning of spike patterns is to use action. As described earlier, the agent is incorporated with action ability, which permits it to move its visual field. Through this ability, the agent generates different actions to move his visual field in various directions. This will cause a change in the sensory array activity in the agent because, as the visual field moves, the visual input changes. As the visual input changes, the corresponding sensory pattern changes. Hence, it can be inferred that there is a relationship between the sensory pattern and the action being performed by the agent. Hence, in simple terms, the task of the agent is well defined: By observing the relationship between the changing sensory patterns and the generated motor actions, the agent has to recover certain properties of the input stimulus [8].

An interesting observation about the sensorimotor relationship can be made when the agent attempts to find a relationship between the sensory patterns and the generated action. The observation is that, from among the wide range of actions that can be performed by the agent for a given stimulus, a subset of the actions will result in identical sensory array activity. When this subset of actions are performed repeatedly, the sensory array pattern will virtually remain the same. A closer examination of this subset of actions reveals that these actions seem to imitate a property of the visual stimulus, i.e., the orientation property of the input.

To make things clear, let us consider the example in Fig. 4, where the input is in the form of an edge oriented at 45°. When the visual field of the agent comes across this edge, a unit of the sensory array is activated. This represents a 'state' of the sensory array. Now, the agent, in an attempt to understand the sensory array pattern, performs the actions it is equipped to do, i.e., move the visual field. Let us assume, for simplicity, that the agent can move the visual field at intervals of 45° which means, the motor vectors defined by (x, y) (as shown in the figure), can assume values from among the following,

$$(x, y) = \{(x, y) : (x, y) \in (1, 0).(1, 1), (0, 1), (-1, 1), (-1, 0), (-1, -1), (0, -1), (1, -1)\} \tag{3.1}$$

When the motor vector is randomly generated from among the values above, the sequence of corresponding motor actions will result in varying sensory array states. However, when the sequence of motor vectors is like

$$(1, 1), (1, 1), (1, 1), (1, 1), ...(-1, -1), (-1, -1), (-1, -1), (-1, -1), \tag{3.2}$$

it means that the visual field of the agent is moved along a direction which is at 45° w.r.t. the horizontal. This direction of motion, in effect, follows the visual input edge as shown in the figure. Since the movement in this direction always maintains the oriented edge in the visual field of the agent, the same unit in the sensory array continuously remains active i.e., the sensory array pattern remains *invariant*.

This observation about sensory invariance, is the key to the solution provided in the thesis. To understand why it is the key, once again consider the above agent generating the action. This time however, the agent is generating action with the aim of maintaining the sensory pattern invariant. In such a scenario, the sequence of motor vectors that are generated by the agent will be similar to those in (3.2). These motor vectors will move the visual field at 45° with the horizontal, which in turn, reflects the 45° inclination property of the input. Hence, it can be said that sensory invariance can be used as a criterion for generating action which will reflect the input properties.

Now, it should be recalled that the task of the agent was to attach a meaning to a sensory pattern. From the example above, it can be seen that a certain action sequence, that is generated by maintaining the sensory pattern invariant, reflects the input stimuli. So if the generated action sequence can be attached as a meaning to the sensory pattern, the goal of the agent, i.e., to understanding the meaning of spike patterns can be achieved. This is the SIDA framework's solution to understanding the meaning of neuronal spikes.

Formally, it is proposed that *sensory invariance driven action*, generated by the agent, can be viewed as the meaning of the sensory pattern. Moreover, it is claimed that in order to associate such a meaning to the spike pattern, sensory invariance can be used as a criterion.

In order to verify the claim, a learning sensorimotor agent is simulated on the lines of the sensorimotor agent described here. The experiments and results are described in the next chapter. In the following section, an algorithm for learning the meaning of neuronal spikes using SIDA is presented.

C. Learning Algorithm

The algorithm is based on reinforcement learning algorithm with Markov assumption (see, e.g., [18]). It is generally used to train an agent, that can sense the state of its environment and act in it. By repeatedly providing the agent with a reward at each state, the agent is trained to choose an optimum action at each state. The optimum action is defined as the one, which will allow the agent to reach a defined goal. Our agent makes an ideal candidate for applying such a learning method. In its case, the states of the environment correspond to states of the sensory array in the agent. The motor vectors generated by the agent that move the visual field in the input scene correspond to the selection of optimum action at each stage. The movement of the visual field will cause a change in sensory states. This corresponds to the change in environment states in reinforcement learning algorithm perspective. The goal of the agent, in terms of reinforcement learning algorithm, is to learn to generate an action sequence which will maintain the sensory array activity constant.

In our sensorimotor agent, there are four sensory states. Each sensory state corresponds to a different orientation preference of the sensors. Based on the sensor which is active, the states are said to correspond to

$$S \equiv \{0°, 45°, 90°, 135°\}. \tag{3.3}$$

It should be noted that 0° refers to orientations of both 0° and 180° in the real world, and likewise for the other three states.

For simplicity, the set of actions D, that can be performed by the agent are limited to eight i.e., the agent can move the visual field in eight possible directions. These possible movements are in the following directions,

$$D \equiv \{0°, 45°, 90°, 135°, 180°, 225°, 270°, 315°\}. \tag{3.4}$$

In terms of the motor vectors x and y, the eight actions correspond to set of motor vectors as defined in (3.1). The goal of the agent is to select an action from $d \in D$, at each state $s \in S$, which maintains the sensory array activity invariant [8]. The agent needs to be provided to with a reward for achieving this goal at each state. To retain the simplicity of SIDA, the reward function is defined in terms of sensory activity itself i.e., reward is linearly proportional to degree of *sensory invariance* in successive stages of action.

In reinforcement learning terms, the agent has to learn a policy function π,

$$\pi : S \rightarrow D. \tag{3.5}$$

at each step t, that 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 such a policy π at each stage (s_t) provides a reward. The reward value is dependent on the state of the sensory array s_t and the selected action d_t. Mathematically, the reward is calculated as,

$$r_t = r(s_t, d_t), \tag{3.6}$$

based on the reward function $r(s, d)$ for $s \in S, d \in D$. This function is updated for each action taken by the agent as follows :

$$r_{t+1}(s,d) = \begin{cases} r_t(s,d) + \alpha * f_t & \text{if } s_t = s_{t-1}, \\ r_t(s,d) - \alpha * f_t & \text{if } s_t \neq s_{t-1}, \end{cases} \tag{3.7}$$

where,

r_{t+1} = reward at step $t + 1$,

$\alpha = 0.01$ i.e., 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 activity to be invariant or (2) the converse (i.e., continuously varying the the sensory array activity). 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 at maintaining sensory invariance will also lead to a high f_t value, but this time the reward for (s, d) will negative.

This simple rule for updating of rewards is used for two reasons, (1) Any sensorimotor agent with a simple processing capability can be trained using SIDA, and (2) To prove that even a simple reward can be a powerful motivation for the agent to learn sensory motor association.

Based on the algorithm presented here, SIDA experiments and simulations have been carried out. The following chapter will describe the experiments performed using SIDA and discuss the results with respect to the claims made about SIDA here.

D. Summary

The SIDA framework uses the sensorimotor approach as a basis to explain how the brain interprets the meaning of sensory spikes. It is proposed that using SIDA, a sensorimotor agent can learn the meaning of spikes in terms of its own actions. It is further proposed that such actions can be generated using the sensory invariance criterion that is introduced in this thesis. Furthermore, a reinforcement learning algorithm that implements SIDA into a sensorimotor agent was described. Based on this algorithm, SIDA experiments and simulations have been carried out to evaluate the effectiveness of SIDA framework. The following chapter will describe these experiments and discuss the results with respect to the claims made about SIDA here.

CHAPTER IV

EXPERIMENTS AND RESULTS

The sensorimotor agent and its environment were simulated using a computer program written in C++. The initial experiments were designed to train the agent to learn meaningful sensory motor association. The agent was then tested in three different behavioral tasks which allow us to see how the agent can understand relatively complex concepts such as two dimensional shapes.

A. Training of the Agent

The training of the agent using SIDA provided the first computational support for the framework. During the training process, the agent tries to learn the association between the sensory patterns and the motor actions that meaningfully reflect the input property. During the process, the agent continuously interacted with the visual environment through a series of 'episodes'. During each episode, the agent was presented with a 51×51 bitmap image of an oriented edge as shown in Fig. 5.

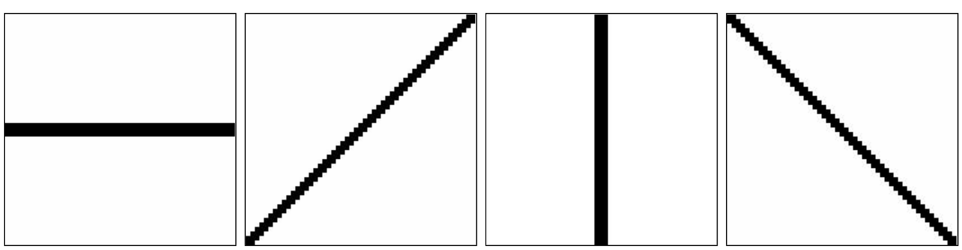

Fig. 5. Inputs used for training and testing. The agent was trained and tested on 51×51 bitmap images each containing a 3-pixel wide oriented edge. Four inputs with four different orientations were used for the experiments (from the left: 0°, 45°, 90°, and 135°).

When these inputs were presented to the agent, the visual field of the agent (9×9 size), begins sliding across the image. The movement of this visual field was controlled by the motor commands generated by the agent. As the visual field slides across the input image, the visual input continuously changes. This continuously changing visual field input was directly compared with each of the four filters (each of size 9×9) in the filter bank. When the orientation of the input was an exact match of the preferred orientation of the input, the sensory state s of the agent was set to value θ, where $\theta \in S$ (see 3.3). If there was a change in the orientation of the input, it caused a corresponding change of the sensory state s.

During this experiment, the agent was trained to learn the policy $\pi : S \rightarrow D$ as described in (3.5) for each of the four oriented edge inputs. The agent was allowed to perform a maximum of eight types of movements as described by D(see 3.4) for each of the four sensory states given by S (see 3.3). Hence, the policy π and the rewards for the agent can be enumerated in a 4×8 matrix. As the agent moves the visual field step by step by generating action $d \in D$, the direction of the next move was determined by the expected reward value stored in the 4×8 'expected rewards' table.

Before the training began, the rewards table was initialized to uniformly distributed random values between 0 and 1. The reward values were restricted to the range $0 \leq r_t \leq 1$ to simplify calculations. Fig. $6(a-d)$ show the initial reward values, where each polar plot corresponds to state $s \in S$, and each plot shows the reward r (distance from origin) for each action $d \in D$ (angle) for a given state s. In an exceptional case, when the agent selected an action which might have resulted in the visual field moving beyond the image boundary, the agent was given a fixed negative reward.

The agent for as long as it was required by the agent to learn the policy, i.e., maximize the reward by consistently meeting the sensory invariance criterion. The average training period usually lasted up to 500 action steps of the agent for each input.

B. Learning the Policy

Fig. 6 shows the results of the training procedure. The polar plots show the reward values for selecting a particular action, for each state of the sensory array. The reward values are denoted by each point in the polar plot: the angle θ represents the direction of motion, and the distance from the origin, the reward value. The plots in fig. $6(a-d)$ show the initial reward values before training and the bottom row plots, i.e., fig. $6(a-d)$ show the reward values after training.

Before the training, the agent did not have obvious preferred direction of motion, for any of its states, which could have resulted in sensory invariance. Hence, the rewards were randomly assigned among the various directions along which the agent can move. The resulting direction of motion due to random nature of rewards could not reflect the properties of the input stimuli. However, we can see that after training, the agent is able to associate a motion d which exactly reflects the orientation of the visual input, with a sensory state s which arises because of the input. For example, in the case of Fig. $6(e)$, the maximum reward value can be gained by the agent, when the current sensory state $s = 0°$ and it moves in a direction which will retain the state, i.e., 0° and 180°. The high reward values in these directions indicate that, given that sensory state, the agent will move in the direction that exactly corresponds to the input orientation. The same rationale can explain the high reward values in certain directions for other states shown in Fig. 6.

On a careful examination of the reward values, we observe that there is a slight difference ($\leq$ 0.01) between the reward values in the preferred directions for each

Before Training

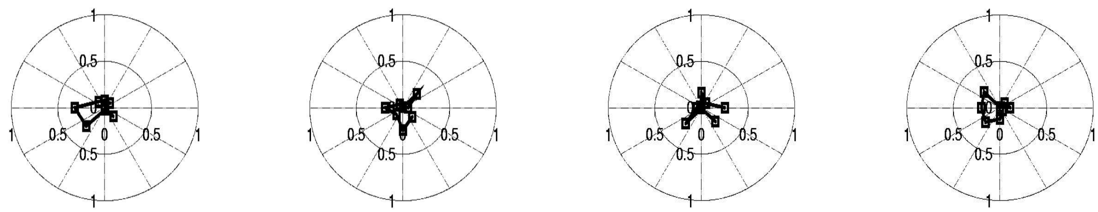

(a) State: 0°　(b) State: 45°　(c) State: 90°　(d) State: 135°

After Training

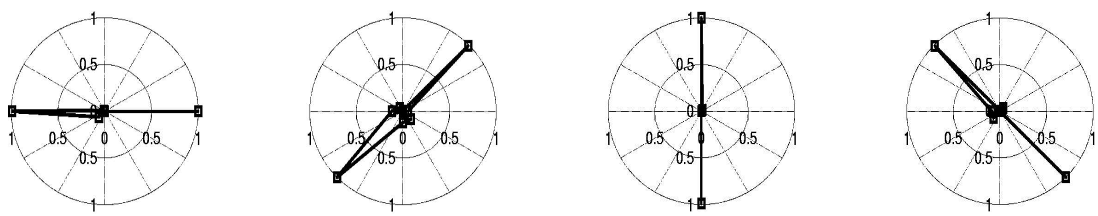

(e) State: 0°　(f) State: 45°　(g) State: 90°　(h) State: 135°

Fig. 6. Reward vectors sensory states. The reward values of the four possible sensory states (0°, 45°, 90°, and 135°) are shown in polar coordinates. The top row from (a) to (d) are before training of the agent, and the bottom row from (e) to (h) are reward values after training. In each plot, for each point (θ, δ), the angle θ represents the direction $d \in D$ of the visual field movement (there are 8 possible directions) and the distance δ from the origin represents the associated reward value given the current sensory state. The reward values were between 0 and 1. Initially, the rewards are randomly assigned for each direction of motion for each sensory state (a to d). After the agent is trained, the reward values become maximal for the movement along the orientations that correspond to the orientation of the input that triggered that sensory state (e to h).

sensory state. For example, in the case of sensory state $s = 0°$, the reward values for traveling in $0°$ and $180°$ directions differ by 0.01. This small difference helps the agent to avoid the dilemma of choosing between the two preferred directions of motion for a sensory state because, the sensory state of an agent can be maintained invariant by moving in either of the two directions. Moreover, this allows the agent to maintain momentum and travel continuously in a particular direction rather than oscillating between two adjacent locations which would also result in sensory invariance. This desirable bias emerged in the agent by itself during the sensory invariance driven learning process.

In the following section, the trained agent will be allowed to perform behavioral tasks so that we can assess the potential of the SIDA framework.

C. Behavioral Experiments

Once the agent was trained, experiments were performed to observe its behavior when presented with simple and later complex visual tasks. The tasks range from simple demonstration on the part of the agent to generate motion which can follow the input to recognition of simple input shapes. The gradual increment in the complexity of the tasks selected and the corresponding behavior of the SIDA trained agent will demonstrate how the SIDA framework can be potentially used by the agent to perform complex tasks, the kind of which, it encounters in real world situations.

1. Generation of Motion to Reflect Input by a SIDA Agent

In this experiment, the agent was made to demonstrate its learning of association between action and sensory state when a straight line was presented in the visual scene. The agent will then look up the rewards table it generated during training.

Using this table, it selects an appropriate action to generate based on the sensory state that arises due to the input stimulus.

The agent is provided with four input images, each of 51×51 in size containing an oriented edge (Fig. 5). Unlike in the training process, the agent is not allowed to learn: The behavior is based on previously learned policy. When the agent starts at some point on the edge, a sensory pattern is generated. The task of the agent is now cutout: Based on the sensory pattern the agent has to look up the rewards table and select an action which will allow it gain maximum sensory invariance.

Fig. 7, shows the action sequences generated by the agent for two sample inputs, i.e., edges with orientations of $0°$ and $135°$. The gray strip in the background is the input, and the directional triangles show the direction of action given the current sensory state. The plots show the sequence of steps taken by the agent while trying to maintain sensory invariance.

The results show that the agent has learned the action sequence that reflects important stimulus properties, i.e., (the orientation) of the input that triggered the sensory state, thus recovering the meaning of the sensory pattern in terms of its own action.

2. Generalization of Action Sequences

The previous experiment demonstrated that the SIDA criterion can be used by the agent to generate action sequences that reflect important input stimulus properties. These experiments involved inputs which were simple in structure, e.g., images of lines at various orientations. In the real world however, the agents are able to acquire the meaning of much more complex shapes and perform complex tasks. Moreover, agents are able to perform such complex tasks based on their learning of simple concepts like the concept of oriented edges discussed in the previous experiment. The following set

0° Input

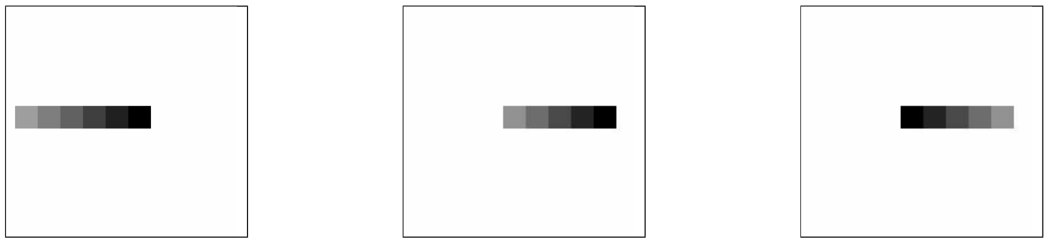

(a) 0< step ≤ 30 (b) 30< step ≤ 60 (c) 60< step ≤ 90

135° Input

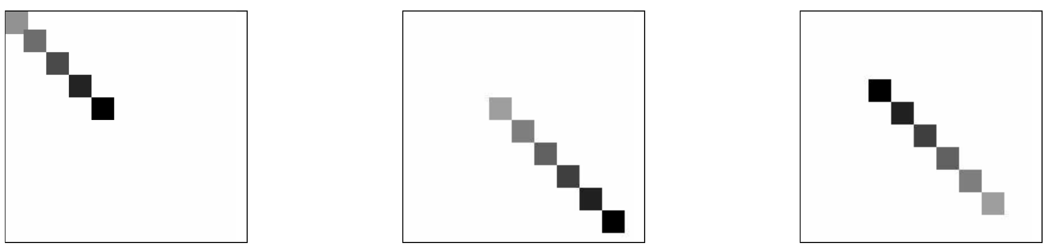

(d) 0< step ≤ 30 (e) 30< step ≤ 60 (f) 60< step ≤ 90

Fig. 7. Behavior of the agent after training. Each plot shows a snapshot of 30 steps of movement of the agent's visual field in the 51×51 scene (only every 5 steps are shown). The shaded boxes indicate the location of the visual field in the scene and their gray scale 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° input and (d) to (f) are for 135°. The trained agent successfully generates motion sequences to trace the input in both runs based on its sensory state and policy π of maximizing sensory invariance. For example, the agent starts in the center and moves right (b), and bounces back when it reaches the end of the input (c).

of experiments demonstrate that the SIDA-trained agent too can generate complex behavior when presented with a complex task. Later, after describing the experiment, I will present an explanation for the behavior of the agent and how such behavior can be be useful for the agent.

In the experiment the agent was presented with two sets of 255×255 images of two dimensional shapes as shown in Fig. 8. One set contained a pair of squares and the other a pair of octagons. The sizes and positions of the shapes in the input scene varied (Fig. 8). The visual field size of the agent was 9×9 and the agent was already trained using SIDA as explained in the previous section.

When an input is provided to the agent, the agent scans through the input till it reaches an edge belonging to the shape. This gives rise to a sensory (activity) pattern, though the sensory pattern may not be exactly the same as the one that the agent was trained on. However, the agent should be able to adapt to variations in input and still make correct choices. In this experiment too, the agent is able to perform robustly to slightly varying patterns. This is achieved by allowing the sensory array to have a graded response. By graded response I mean that even if the input is not of an orientation to which the sensory neuron is exactly tuned, it will still generate a fraction of the activity depending on the angular separation of the orientation. For example, if the input is oriented at $75°$, the neuron which is tuned to $90°$ orientation will generate an activity slightly less than its maximum. Similarly other neurons will also generate some activity.

The agent is generally confronted with two variations from standard sensory patterns. (1) The input edge may be in the periphery of the visual field of the agent. In that case, the activity in the neuron, which is tuned to that particular orientation, may not have as high a magnitude. (2) Another case could be when there are overlapping patterns in the input, such as at corners of the shapes. Here

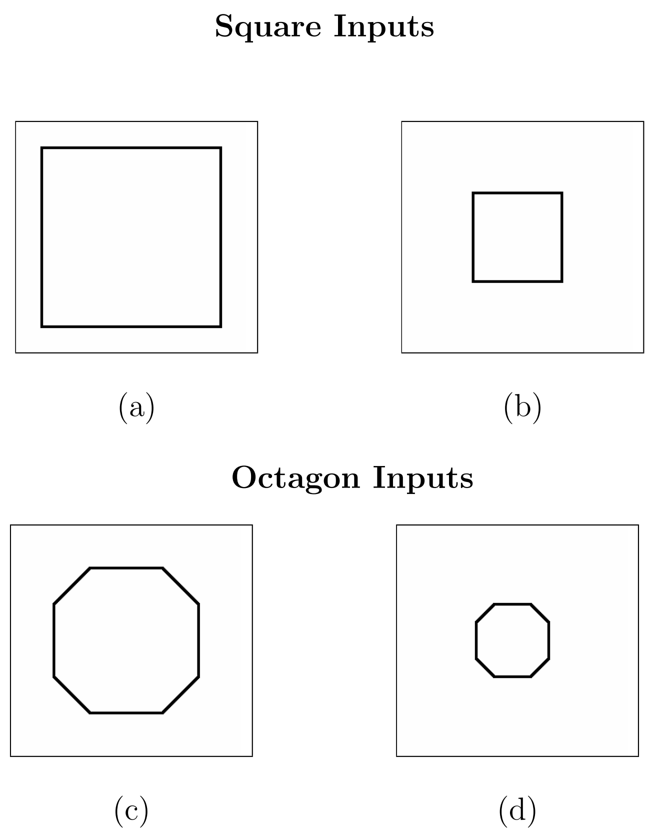

Fig. 8. Behavioral experiment: Inputs provided to the agent. The agent was provided with 255×255 bitmap images each containing a shape. The thickness of each edge was approximately 3-pixels wide. Two pairs of inputs were provided. One pair was squares of different sizes, and the other was octagons.

there will be at least two oriented inputs and it will result in two units of the sensory array being active simultaneously. In the first case the agent treats the less than maximum sensory activity as standard sensory activity and proceeds to perform the corresponding action. In the latter case, the agent tries to use the dominant pattern as the causal pattern. In the case where both patterns are equally dominant, the agent will decide based on the momentum and exploration bias. To do this the agent will recall the direction it traveled so far and venture in a different direction. In case it hits a dead end, it reverts to the original direction of travel. This ability is incorporated into the agent through the use of short term memory.

When the two dimensional shape was presented to the agent, the agent tried to trace the shape by following each side of the shape. It did so by treating each side as an independent oriented edge and attempted to maintain the sensory invariance for that edge. Fig. 9 shows behavioral trajectory of the agent. In the figures, we can notice that at some corners, the agent moved in a direction different from the intended path, although eventually it gets back on to the path. The number of such steps on the unintended path are negligible: They can be explained based on the bias in the agent as discussed in chapter 4. The corners of the shape can be regarded as overlapping of oriented edges. When the agent decided to follow an adjacent oriented edge, i.e., another side of the shape, it used its initial bias for one of the two possible directions (recall that each sensory state has two preferred directions of motion, one opposite of the other), to randomly select the direction to take. Based on whether the selected direction is correct or not, the agent eventually (if not immediately) continues to travel along the intended path.

To maintain the sensory invariance for each edge, the agent generated a series of actions that correspond to motor vectors (x_1, y_1),(x_2, y_2),(x_3, y_3),....(x_n, y_n). These groups of consecutive motor vectors that are generated by the agent to maintain an

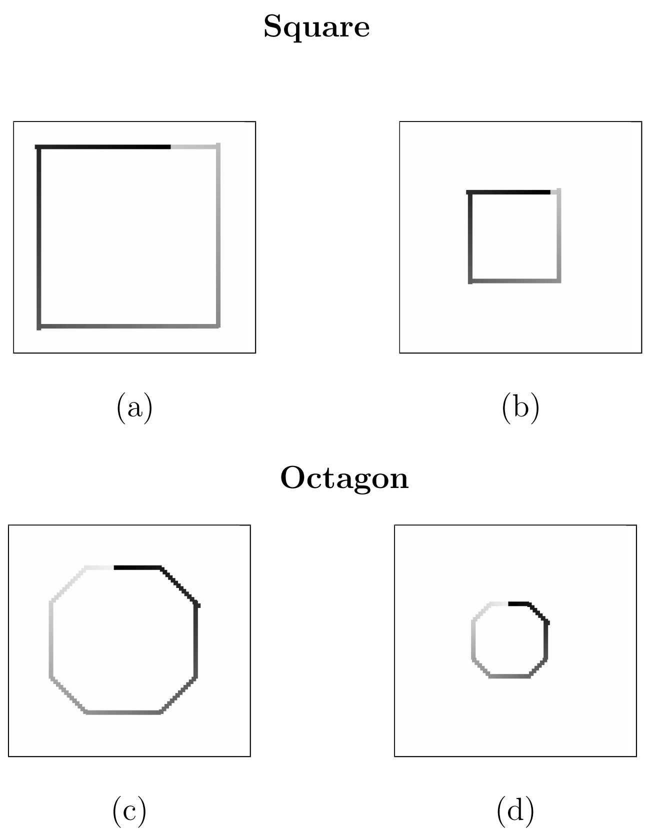

Fig. 9. Behavioral experiment: Paths traced by the agent. The agent traces the 255×255 inputs presented to it using sensory invariance as a criterion. The grayscale trace shows the agent's trajectory where the darker areas indicate the latest positions and the light ones, the preceding positions. The direction of the path traced is from light to dark grayscale. The intersection of the light and dark edges is the starting point of the trajectory.

invariant sensory state are termed as *motion vectors*. So, whenever a new sensory state arises in the agent, a new motion vector is generated. Ideally, since the sensory states are in response to input stimuli, i.e., the oriented edges in the input, the number of unique sensory states should not exceed number of sides of the input shape. For example, consider that the input is an octagon (Fig. 8(c)). In that case, the agent will encounter eight oriented edges as inputs, each of which give rise to a sensory state. For each of these eight sensory states, the agent generates a motion vector. Hence, for an octagon, a total of eight motion vectors are generated.

Fig. 10 plots the motion vectors that are generated by the agent, while traversing the paths as shown in Fig. 9. While plotting these motion vectors, the motor vectors that result in *negligent* movement of the agent along unintended paths are not used. In the agent, this is achieved by setting a lower limit on number of consecutive motor vectors required to form a motion vector.

From Fig. 10, it is evident that when the motion vectors are put together, and compared with the inputs they look very similar. This observation is very much in line with the earlier observation in case of oriented edge inputs, where the action generated by the agent reflected the input stimulus properties. Similarly the motion vectors generated by the agent in this experiment also reflect the input stimulus properties of the shape as a whole. This group of motion vectors that together reflect the input stimulus properties of a two dimensional shape is termed as *motion vector sequence* (MVS)(Fig. 11). Thus, MVS can be said to reflect the properties of the input stimulus. Sample MVSs for square and an octagon are shown in Tables I. and II. Each column in the tables is a motion vector for an input shape. Each row is a Motion Vector Sequence (MVS). The top row shows the direction of each vector N (North), E (East), W (West), S (South), NE (Northeast), SE (Southeast), NW (Northwest), SW (Southwest). The other rows show the length of each vector. Each

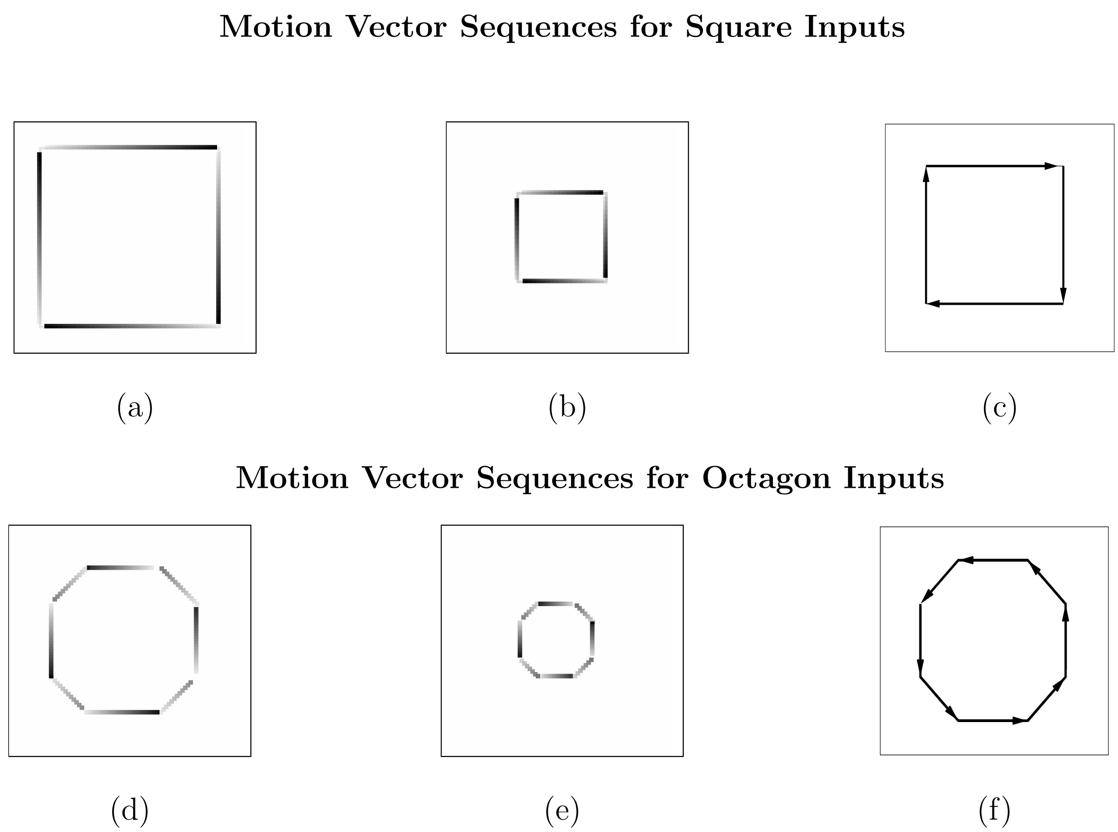

Fig. 10. Motion Vector Sequences (MVS). The agent generates motion vectors based on sensory patterns. The set of figures (a-b),(d-e) show the motion vector sequences generated by the agent for the input shapes in Fig. 8. Each side of the MVS plot indicates one vector. The direction of the vector is from the light to the dark shades. These plots clearly demonstrate the similarity of vectors for similar shapes irrespective of the sizes. The figures (c)and (f) show the normalized vectors for the two class of shapes namely squares and octagons.

motion vector denotes the number of steps moved by the agent's visual field and the direction of movement.

Table I. Sample MVS table for a square

Direction	E	S	W	N
Length	189	189	187	189

Table II. Sample MVS table for an octagon

Direction	NE	E	SE	S	SW	W	NW	N
Length	28	71	28	71	28	71	30	74

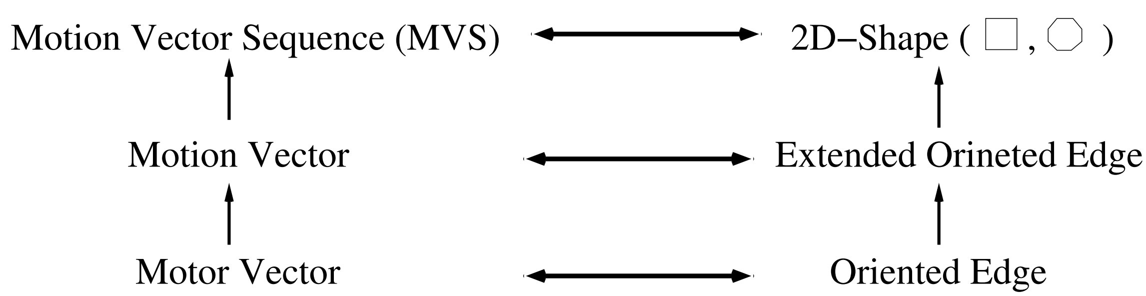

Fig. 11. Hierarchy of vectors.

Another observation from the Fig. 10 is that the MVS for inputs of the same shape but different sizes appear to be similar. For example, consider the two motion vector sequences for the pair of octagons presented as inputs in Fig. 10(d) and Fig. 10(e). Each of these MVSs constitute eight motion vectors where the angle between corresponding adjacent vectors is the same. The only difference is the length of the corresponding motion vectors. It is interesting to note that this observation about the difference in lengths was possible when MVSs were used rather than when sensory patterns themselves were used. This leads us to believe that that MVS can be potentially used to provide more information about the input in a manner so far not exploited. A potential use of MVS is for *generalization* of shapes, i.e., using MVS to reflect a whole class of shapes, rather than independent shapes of different sizes and positions. For example, MVS can be used to categorize all squares irrespective of size and position in the input scene as belonging to the 'square' class of shapes based on the motion vectors. The actual comparison needed for such categorization can be achieved by normalizing the motion vectors to a standard size. Then these normalized motion vectors can be said to reflect the general property of that shape class.

There are many advantages for using MVS in the sensorimotor agent and the SIDA framework in general. The most important feature of MVS is that it reflects the properties of a complex shape. Though the sensorimotor agent was trained only on very simple oriented edges, the agent was able to generate MVS for complex shapes. Hence, it is fair to assume that MVS are easy to generate and very simple agents can generate them based on the same sensory invariance criteria. Another advantage of MVS is that they are easy to compare. As can be seen from the Fig. 10, MVS for different sizes are very similar except for the length of the component motion vectors. Based on this comparison, as explained before, the MVS can also be used to reflect properties of a class of shapes. In addition, MVS can be potentially used across different modalities. This aspect of MVS vectors will be discussed in the next chapter.

By understanding the advantages of the MVS, it will be easier to understand the next set of experiments. We have shown that MVS generated by the agent can be used to generally reflect properties of a class of shapes. In the next set of experiments the agent will be able use these MVSs, generated using the SIDA framework, for understanding realistic spike patterns and use them for stimulus recognition.

3. Spatio-Temporal Patterns and Recognition

So far, we have seen that when a sensorimotor agent was presented with a simple oriented edge, a sensory pattern arises in the agent. These patterns are spatial in nature, and henceforth, these sensory patterns will be referred to as *spatial sensory patterns* or simply *spatial patterns*. A spatial pattern reflects the distribution of activity across a group of neurons in space (as opposed to time). In the case of our sensory agent there are four units, each corresponding to an orientation tuned neuron. Depending on the orientation of the input, a particular unit is maximally active and others are not. Any such state of the sensory array is a spatial pattern. Fig. 12 below shows the four basic spatial patterns that can arise in the sensorimotor agent. Of course, owing to graded response in the agent (refer to section C.2), the activity in each unit of the sensory array may not be always on or off but vary over a range, which might make the spatial patterns not so easy to distinguish. Moreover, in some cases the spatial patterns may overlap, when more than two oriented edges exist simultaneously in the visual field of the agent.

In the real world, the sensorimotor agent seldom comes across shapes which are simple oriented edges. The visual stimuli will always be a complex shape with a combination of oriented edges. These shapes give rise to spatial patterns which continuously vary over a period of time. In fact an individual spatial pattern will provide only partial information about the stimuli. This may not allow the agent to extract proper meaning from the sensory states. Hence, the agent has to look at the varying patterns over a period of time. Such sensory patterns which vary over time and together over time are in response to a complex input are called *spatio-temporal patterns*. For example, in this experiment, when an input was provided to the agent

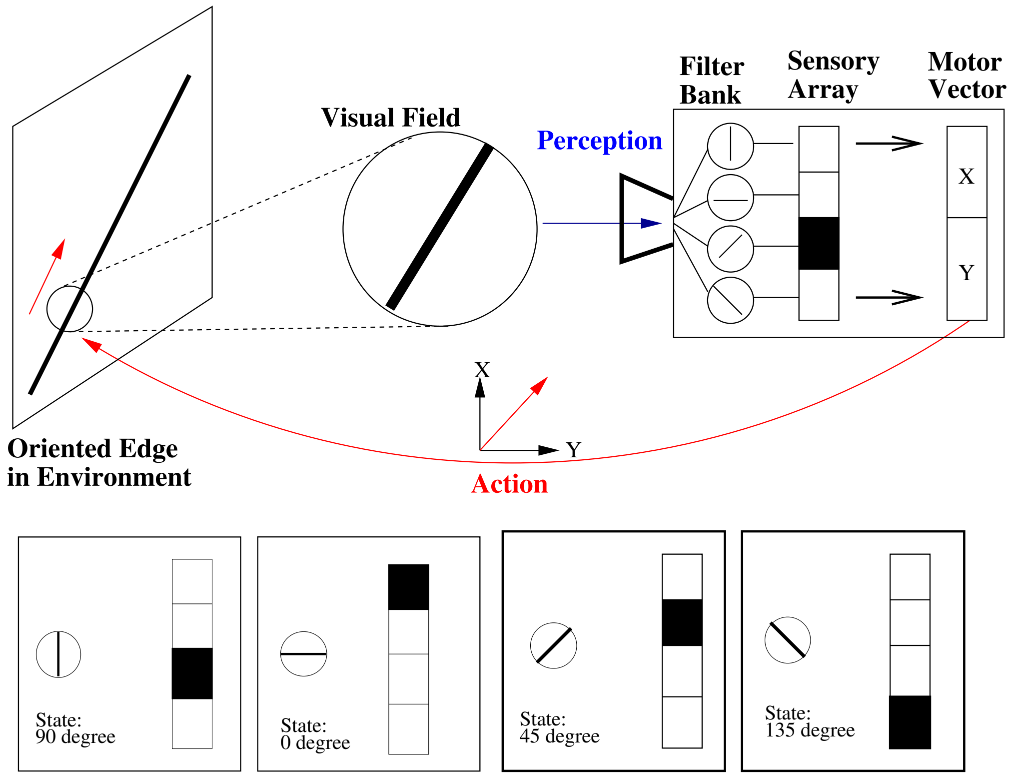

Fig. 12. Spatial patterns in a sensorimotor agent. The four basic spatial patterns which can be observed in a sensorimotor agent are shown. Each spatial pattern arises in response to a unique oriented edge as input. The four patterns represent the 0°, 45°, 90° and 135° sensory states of the agent.

in the form of an octagon (see Fig. 8), its tracing of the octagon generated four spatial patterns which continuously repeat in a cyclic fashion (see Fig. 13) over the duration of observation by the agent. Each individual spatial pattern was generated in response to a single side of the octagon, whereas the spatio-temporal pattern as a whole was generated by the octagon over a period of time.

In the previous experiment we have seen that the agent generates MVSs in response to complex shape inputs. In this experiment the agent generates spatio-temporal sensory patterns for complex shapes. Moreover, the MVS generated by the agent reflected the property of a class of shapes presented as input. In this case the

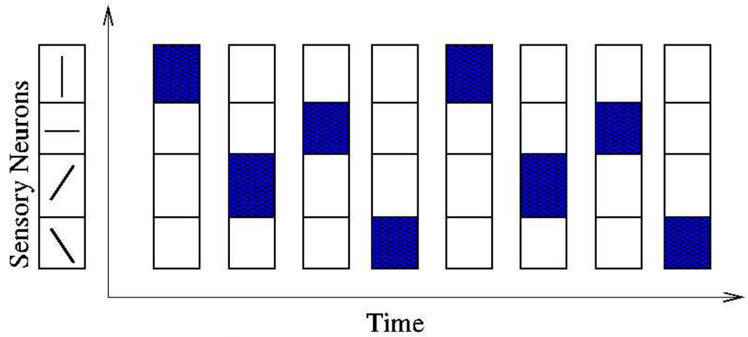

Fig. 13. Spatio-temporal pattern for an octagon. An illustration of a spatio-temporal pattern for an octagonal input is shown. There are four spatial patterns generated for an octagon. Each octagon gives rise to a sequence of eight spatial patterns (twice each pattern). These patterns are distributed over time as the agent traces the octagon with its visual field.

spatio-temporal patterns are in direct response to input shapes. Hence, we hypothesize that spatio-temporal patterns and MVSs are somehow related to each other.

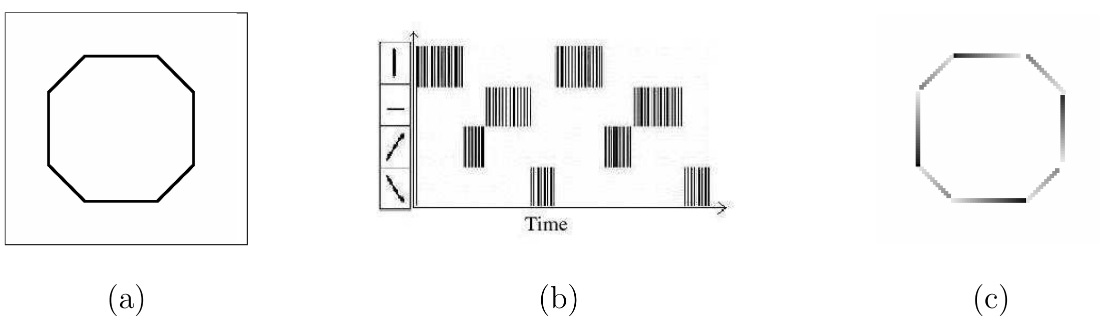

Fig. 14. MVS generation for an octagon. (a) The agent is presented with an octagon as a visual input (b) The octagon input generates a spatio-temporal pattern in the sensory array of the agent. (c) The agent uses SIDA to generate a Motion Vector Sequence (MVS) for the octagon.

Fig. 14 shows the spatio-temporal pattern and the corresponding MVS generated by the agent for the same input, i.e., octagon. This relationship between MVS and spatio-temporal pattern is used by the agent to acquire the meaning of spatio-temporal patterns in terms of MVS. This approach is a logical extension of how the agent acquired the meaning of a simple spatial sensory pattern in terms of motion generated by the agent during its training (refer section A). The advantage of establishing such a relationship between MVS and spatio-temporal pattern is that the usability of the MVS to generalize shapes can be exploited to understand the meaning of wide range of spatio-temporal patterns and perform simple tasks such as recognition of shapes.

Pattern recognition in the following experiments is defined as the ability of the agent to identify and classify various input shapes into a category of predefined class of shapes. For example, if the agent is presented with circles of various radii at various positions in the input scene, the agent should be able to identify all of them as circles. In this experiment the agent was introduced to the task of pattern recognition and pilot experiments are carried out. The agent was used to classify the shapes belonging to two categories, i.e., squares and octagons.

The first task of the experiment was to build into the agent a database of category of shapes, by defining how to store and label the information it gathers using SIDA. Here, we describe how the agent was trained to build a database for the octagon. Initially the agent is presented with an image of an octagon (Fig. 14(a)) along with the label identifying it as an 'octagon'. When this visual input is provided, a spatio-temporal pattern is generated in the sensory array of the agent as seen in Fig. 14(b), from which it generates a MVS for the octagon (Fig. 14(c)). Recall from the previous section that the MVS can be used by the agent to generally reflect the category of shapes using normalization. Using the technique, the agent stored the MVS for category of octagons in this experiment. For the octagon presented the size of the motion vectors forming the MVS is given in the Table III. In the table each column

Table III. Normalization of MVS

Direction	NE	E	SE	S	SW	W	NW	N
Length(Input Octagon)	28	71	28	71	28	71	30	74
Length(Normalized Octagon)	50	126	50	126	50	126	50	126

is a motion vector. Each row is a Motion Vector Sequence (MVS). Top row shows the MVS for input octagon and the bottom row shows the normalized MVS for the same input.

While normalizing the vectors the agent allows for a 10% difference between lengths of the motion vectors while deciding whether they should be considered equal or not. So in this example the vector NE28 differed from N71 by 43 $>$ 10% of NE28. Hence, they are considered to be of different sizes. On the other hand, N71 and E74 are equal because they differ by $<$ 10% of N71.

This adjustment is based on the behavior of a natural agent. A natural agent, upon observation can only approximately identify the sides of a polygon as equal or unequal. It cannot distinguish between lengths of sides of a polygon if they differ by minute amounts. In this experiment this allowance has been empirically chosen to be 10%.

The vectors of the MVS were normalized to a size of 50 (can be of any value, if it is consistently used), i.e, the smallest motion vector of the MVS is of length 50 and all other vectors are proportional to it. This normalized vector was now labeled as an 'octagon' and stored in the agent. The same process was carried out for the square MVS and the normalized MVS stored in the agent. The agent now, has a database of category of shapes: squares and octagons. Now the task of the agent is to identify the inputs presented to it as squares or octagons irrespective of the sizes and positions in the visual scene.

For the task of recognition, the agent is presented with three isomorphic shapes as shown in Fig. 15. For these inputs, the agent generates spatio-temporal patterns in the sensory array as shown in the Fig. 16.

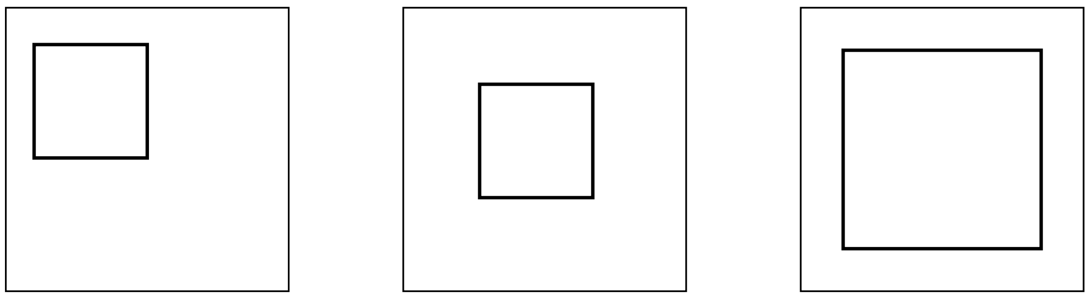

(a) (b) (c)

Fig. 15. Inputs for recognition.

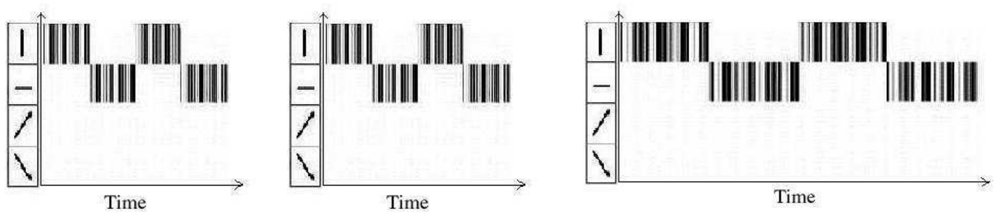

(a) (b) (c)

Fig. 16. Spatio-temporal patterns for the three inputs shown in Fig. 15.

These inputs gave rise to MVS as shown in Fig. 17. The sensorimotor agent then normalized them to a standard size as described earlier. The normalized MVS generated in this fashion was compared with with normalized labeled vectors. For all the three inputs shown here, the agent was able to find a matching normalized MVS in its database. Hence the agent identified all three inputs as squares.

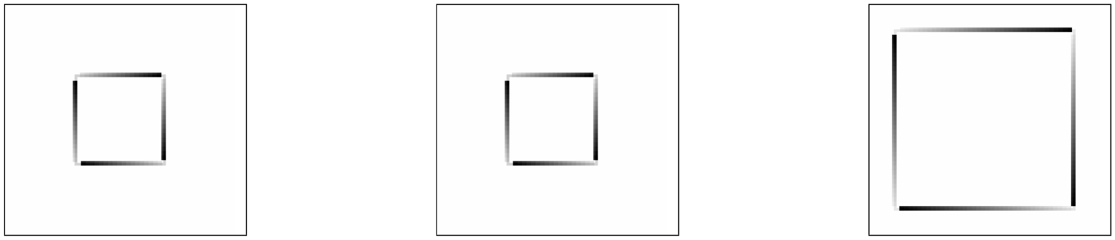

(a) (b) (c)

Fig. 17. MVS for the three inputs shown in Fig. 15.

This task of shape recognition demonstrates that the agent can interpret the spatio-temporal patterns presented to the agent using MVS. Moreover, it also proves that the learning of association between generated action and simple sensory patterns for oriented edges can be used by the agent to understand spatio-temporal patterns generated in response to complex shapes.

D. Summary

In this section a set of three experiments have been carried out. The first experiment demonstrated that the sensorimotor agent can learn the association between a sensory pattern and the action it generates, using sensory invariance as a criterion. This criterion forms the basis of the SIDA framework. From the results it is evident that the agent can anticipate the next step it has to take for the type of input provided to it. This ability to anticipate implies that the agent has been able gain knowledge of properties of the input and has been able to use them to generate appropriate action steps. The rest of the experiments seek to extend the ability of the agent beyond spatial patterns to spatio-temporal patterns. In the second experiment, it is argued that the agent can generate motion vector sequences (MVS) for complex shapes. These MVSs reflect the properties of the input shapes. The use of MVS also makes the comparison of shapes presented easier. The MVSs can then be used for the generalization of input properties for a set of shapes of varying sizes and positions. The third experiment introduces spatio-temporal patterns into the SIDA framework. The sensorimotor agent learned to use MVS for understanding the meaning of the spatio-temporal patterns in its sensory array. The agent used the generalization concept to normalize motion vectors and build a database of shapes. These normalized MVS were then used by the agent to perform complex tasks such as recognition of complex shapes.

In the following chapter, the SIDA framework and its basic concept, i.e., sensory invariance, are discussed in light of the experimental results and compared with the work of other researchers in this field.

CHAPTER V

DISCUSSION AND FUTURE WORK

This chapter discusses the main contribution of the SIDA approach with respect to related work by other researchers using sensorimotor approaches (section 5.A), followed by discussion on issues related to the SIDA framework (section 5.B), and future directions of this work (section 5.C).

A. Main Contribution

The main contribution of this thesis is that it demonstrated that a sensorimotor agent can acquire the meaning of neuronal spikes through its own actions. From the experiments it has been demonstrated that the sensorimotor agent is able to interpret the spike patterns in terms of actions it generates to reflect the input stimulus properties. This is similar to how the brain understands the objects in the world, i.e., by the features and properties of the object.

The sensorimotor agent presented in this thesis is similar to the robot described by Cohen and Eal in their work on natural semantics [9]. Natural semantics, as explained in chapter 2, refers to the ability of a system, an agent in our case, to acquire and maintain meaning for itself. In their work, Cohen and Bea report a semantically capable robot which ‘learns concepts and meanings by interacting with its environment’ [9]. The sensorimotor agent described in this thesis is also able to learn about the stimulus by generating behavior which influences the way the agent perceives the stimulus. However the two models differ in terms of scope of natural semantics. Cohen and Beal’s robot interprets the sensory patterns resulting from objects in the world, based on the role played by the object in the robot world. Now the role played by the object can vary depending on the action performed by the robot. Hence the robot is not actually trying to understand the sensory activity in terms of the properties of the object but rather, in terms of the role of the object in the robot's world. On the other hand the SIDA agent is trying to learn the meaning of the sensory activity by generating action which could reflect the object properties without taking into consideration the role played by the object in the agent's experience of the world. Because of this difference, Cohen and Beal's robot can be considered to be at a more abstract level, because it seems to be aware of the roles it is playing. SIDA agent makes no such assumption and assumes that the agent does not have any concept of the roles the stimuli can play in its world.

The mechanics of the sensorimotor agent are however much similar to the learning agent used by Pierce and Kuipers [23, 24]. In both cases, the agent tries to explore an unknown world and use a 'generate and test' approach to learn about the features in the world. However, they differ from each other in that the sensorimotor agent uses the sensory invariance criterion to generate action, where as the learning agent in [23, 24] randomly generates action and uses the resulting experience as a feedback to gain knowledge about its actions and consequences. Moreover their method does not assume that the agent is aware of its action capabilities which is an implicit assumption in SIDA. Their learning agent thus begins by trying to gain knowledge of its own actions before trying to interpret the sensory patterns in terms of its own actions. On the other hand however, their agent can perform complex computational analysis (linear regression, statistical analysis, controlling the robot etc.,) which means their learning agent makes some other assumptions about the capabilities of the agent. The sensorimotor agent in our case is however, based on very simple model based on the invariance criterion, which does not require complex mathematical knowledge or mechanics built-in into it.

The most significant contribution of this thesis is the *sensory invariance criterion* for the agent to learn the meaning of spikes. This criterion makes for a simple yet powerful tool for learning of meaning in sensorimotor agents. Moreover, the invariance used in this thesis offers a different perspective to invariance from the conventional one. Conventionally, invariance is something that is sought in the environment, e.g., invariant feature detection in vision in which the detecting system seeks features which are supposed to be invariant spatially or temporally. However in the SIDA framework, invariance is *enforced* by the system 'internally' in the agent through a well choreographed action. Although the methodology of using internal neural sensory invariance is different from external invariance in sensory features, they can be regarded as related to each other. For example, consider an object approaching an observer. As the object approaches, a neuron fires in the observer, and the same neuron will fire if the object is stationary and the observer moves towards the object. Thus, the meaning of neuron firing can be understood in terms of action that would turn on the neuron reliably [8]. Hence even in the case the object is moving towards the observer and the observer is stationary, since the same neuron is firing, the brain interprets the spike activity as equivalent to moving forward towards the object [12].

Lastly, the criterion of sensory invariance provides a solution to the question, "what quantity should a perception-action cycle cycle system maximize, as a feed-forward channel might maximize its channel capacity?" (Bell [6]). This question originates from Shannon's information theory [13], which seeks to explain what function is maximized in the information channel for effective information transfer. By considering the reward criterion used in training the agent, i.e., sensory invariance as the function to be maximized during the learning process, it has been shown that the learning of meaning of spike patterns can be achieved effectively. Thus, hypothetically, sensory invariance can be seen as the quantity which the perception-action cycle system must maximize, and as a result acquire understanding of the environment.

B. Issues with the New Framework

This thesis lays the foundation for a new framework, the SIDA framework. As with any new framework, there are a few issues which need to be separately addressed for the better understanding of the framework. Some of the issues are addressed in this section.

1. Why Not Use a Direct Analysis?

In this thesis the agent used reinforcement learning to recover the features of the stimulus. However, with a sufficiently large two dimensional array of sensors it is possible to recover several or all properties of the visual environment within the agent using unsupervised learning. The reason why several such sensors were not used is, even with several sensors the direct knowledge about the stimulus is not available beyond the first stage of sensory processing. What we set out to seek was to understand how the properties are recovered at a later stage in the brain. In the scenario of using several sensors, it is only the initial stages which would have access to properties of stimulus as they are in the visual world. At each of the later stages, the information received is further encoded until it reaches the final stage, the sensory array where it is interpreted. So for the purposes of the understanding how the spikes are interpreted in the brain, it is not feasible to analyze the properties at the initial stages.

2. How Can We Have Direct Knowledge of Our Own Movements When We Cannot Have That for Perception?

The SIDA framework *assumes* that the sensorimotor agent *already* possesses knowledge of its own actions, based on which the agent is able to associate meaning to spikes. In a realistic situation, this assumption may or may not be true. Then the agent should be able to learn about the actions it can perform. There is enough research work in the area which suggests that sensorimotor agents can learn about their own actions. For example, Philipona et al. [22] identify two sets of sensors, extroceptive and proprioceptive, of which agents can have complete control over proprioceptive sensors. Proprioceptive sensors lets the agent learn about the position, location, orientation, and movement of body and its parts. Hence, it might be possible for action, and closely tied proprioceptive sensors, to provide direct knowledge to the agent about its own motor abilities.

C. Future Work

In the experiments performed using the SIDA agent, the stimuli presented were always related to visual world. Can it be used for other sensory modalities like audition, somatic sense, olfaction, etc.? The approach used here is general enough that it can be potentially extended to all senses which are related to action in some manner. For example, somatic sense, which like vision depends to a considerable extent on the actions involved, can use SIDA to follow the surface contours of an object by maintaining sensory invariance. However, SIDA cannot be used in domains which do not have much correlation between action and perceived sensory state e.g., olfaction, audition, etc.

Another direction for future work would be to use MVS across sensory modalities.

Recall that MVS was generated in response to spatio-temporal patterns in the sensory array of the agent. Such MVSs are independent of the modality giving rise to them. Hence they can be used by other modalities as well. For example, consider the shape of an octagon is stored as a MVS. This MVS is built from the visual input. Now, given that this MVS is a motor concept and the stimulus features are in the form of a sequence of motion vectors, the MVS can be relayed outside the visual system of the brain and into the tactile sensory system. When the tactile system is provided with such a set of MVSs, the subject can recognize shapes which are octagonal in shapes through touch. Moreover MVS generated through any sensory modality can be used to directly generate motion. If the MVS is provided to generate motion in limbs, the shape reflected by the MVS can be generated through motion of limbs. This logical extension is based on how a child learns to draw different shapes. For example, when the child sees the shape of an octagon on the blackboard, the visual sensory activity generates a MVS in his brain. That the child is able to draw the same shape on paper could be because he is generating action according to the MVS. This simple observation can thus be used a basis to explore how the MVS from various modalities can be used to generate desired action, like in this example, the MVS from visual system being used to generate the octagon on paper using the hand.

In the previous section it has been pointed out that the SIDA framework assumes that the agent is aware of its actions. However, as part of future work it is proposed that SIDA be explored so that the sensorimotor agent can learn about its own actions. This attempt will be made based on the work of Philipona et al. [22] and Pierce and Kuipers [24], [23].

Lastly, the simulated neurons in the agent were tuned to only four orientations. In later experiments, the number of simulated neurons can be increased to be more in proportion to the number of orientation tuned neurons in brain. Moreover, the inputs to the sensorimotor agent so far involved complex 2D geometrical shapes. In the future the inputs can be objects from the real world. It would be interesting to observe the behavior of SIDA agent in response to such inputs. This could also be the next step in building realistic agents which can learn natural semantics.

CHAPTER VI

CONCLUSION

This thesis brings out the problem that neural decoding methods that require direct knowledge of the stimulus cannot provide satisfactory explanation to how the brain, which does not have any direct knowledge of the world, can understand the meaning of neural spikes. This thesis provides a solution to this problem through a sensorimotor approach. A new framework called Sensory Invariance Driven Action (SIDA) based on a sensory invariance criterion was proposed that can be used by a sensorimotor agent to acquire the meaning of its internal state (i.e., spikes) in terms of its own actions. Several computational experimental results are presented which show that the SIDA framework can be used to easily and effectively associate meaning to spike patterns. The success of SIDA in the experiment raises the hope that we can be built autonomous agents that can learn natural semantics based on their own actions.

REFERENCES

[1] E. D. Adrian, “The impulses produced by sensory nerve endings,” *Journal of Physiology (London)*, vol. 61, pp. 49–72, 1926.

[2] M. A. Arbib, “Language evolution: The mirror system hypothesis,” in *The Handbook of Brain Theory and Neural Networks*, 2nd ed., M. A. Arbib, Ed. Cambridge, MA: MIT Press, 2003, pp. 606–611.

[3] P. Bach y Rita, *Brain Mechanisms in Sensory Substitution.* New York, NY: Academic Press, 1972.

[4] ——, “Tactile vision substitution: Past and future,” *International Journal of Neuroscience*, vol. 19, pp. 29–36, 1983.

[5] H. B. Barlow, “Summation and inhibition in frog’s retina,” *Journal of Physiology*, vol. 119, pp. 58–68, 1953.

[6] A. J. Bell, “Levels and loops: The future of artificial intelligence and neuroscience,” *Philosophical Transactions of the Royal Society of London*, vol. 354, pp. 2013–2020, 1999. [Online]. Available: http://www.cnl.salk.edu/˜tony/ptrsl.pdf

[7] S. R. Cajal, *Histologie du systéme nervéux de l’homme et des vertebrés,* translated by N.Swanson and L.W.Swanson. New York: Oxford University Press, 1995.

[8] Y. Choe and S. K. Bhamidipati, “Learning the meaning of neural spikes through sensory-invariance driven action,” Department of Computer Science, Texas A&M University, Tech. Rep. 2003-8-3, 2003. [Online]. Available: http://faculty.cs.tamu.edu/choe/ftp/publications/choe.tr03.8-3.pdf

[9] P. R. Cohen and C. R. Beal, “Natural semantics for a mobile robot,” Department of Computer Science, University of Massachusettes, Amherst, Tech. Rep. 00-59, 2000. [Online]. Available: http://www-eksl.cs.umass.edu/papers/cohen-ECCS99.pdf

[10] H. L. Dreyfus, *What Computers Can't Do: The Limits of Artificial Intelligence*, revised ed. New York, NY: Harper & Row, 1979.

[11] A. P. Georgopolous, A. Schwartz, and R. E. Kettner, “Neuronal population coding of movement direction,” *Science*, vol. 233, pp. 1416–1419, 1986.

[12] J. J. Gibson, *The Perception of the Visual World*. Boston: Houghton Mifflin, 1950.

[13] P. Hacker, “Languages, minds and brain,” in *Mindwaves: Thoughts on Intelligence, Identity, and Consciousness*, C. Blakemore and S. Greenfield, Eds. Oxford, UK: Basil Blackwell, 1987, ch. 31, pp. 485–505.

[14] H. K. Hartline and F. Ratiliff, “Inhibitory interaction in the retina of Limulus,” in *Physiology of Photoreceptor Organs*, M. Fuortes, Ed. New York, NY: Springer-Verlag, 1972, ch. 10, pp. 381–447.

[15] A. L. Hodgkin and A. F. Huxley, “A quantitative description of membrane current and its application to conduction and excitation in nerve,” *Journal of Physiology*, vol. 117, pp. 500–544, 1952.

[16] D. H. Hubel and T. N. Wiesel, “Receptive fields, binocular interaction and functional architecture in the cat's visual cortex,” *Journal of Physiology (London)*, vol. 160, pp. 106–154, 1962.

[17] S. Hurley, “Perception and action: Alternative views,” *Synthese*, vol. 129, pp. 3–40, 2001. [Online]. Available: http://www.warwick.ac.uk/staff/S.L.Hurley/papers/pa.pdf

[18] T. M. Mitchell, *Machine Learning*. McGraw-Hill, 1997.

[19] R. Möller, “Perception through anticipation: An approach to behavior-based perception,” in *Proceedings of New Trends in Cognitive Science*, 1997, pp. 184–190. [Online]. Available: http://www.psy.mpg.de/CR/PEOPLE/MORA/documents/ntcs97.fm.ps.gz

[20] M. W. Oram, P. Földiák, D. I. Perrett, and F. Sengpiel, “The ‘ideal homunculus’: Decoding neural population signals,” *Trends in Neuroscience*, vol. 21, pp. 259–265, 1998.

[21] J. K. O’Regan and A. Noë, “A sensorimotor account of vision and visual consciousness,” *Behavioral and Brain Sciences*, vol. 24, no. 5, pp. 883–917, 2001. [Online]. Available: http://www.bbsonline.org/Preprints/ORegan/

[22] D. Philipona, J. K. O’Regan, and J.-P. Nadal, “Is there something out there? Inferring space from sensorimotor dependencies,” *Neural Computation*, vol. 15, pp. 2029–2050, 2003. [Online]. Available: http://nivea.psycho.univ-paris5.fr/Philipona/space.pdf

[23] D. M. Pierce and B. J. Kuipers, “Map learning with uninterpreted sensors and effectors,” *Artificial Intelligence*, vol. 92, pp. 162–227, 1997.

[24] D. M. Pierce, “Map learning with uninterpreted sensors and effectors,” Ph.D. dissertation, Department of Computer Sciences, The University of Texas at Austin, 1995.

[25] F. Rieke, D. Warland, R. de Ruter van Steveninck, and W. Bialek, *Spikes: Exploring the Neural Code.* Cambridge, MA: MIT Press, 1997.

[26] G. Rizzolatti, L. Fogassi, and V. Gallese, "Neurophysiological mechanisms underlying the understanding and imitation of action," *Nature Reviews Neuroscience*, vol. 2, pp. 661–670, 2001.

[27] S. S. Scott, "Population vectors and motor cortex: neural coding or epiphenomenon?" *Nature Neuroscience,*, vol. 3, pp. 307–308, April, 2000.

VITA

Sarvani Kumar Bhamidipati was born in India on the 21st of August, 1979. After completing his schooling in the Kendriya Vidhyalaya Sanagathan, India, he went on to attain his Bachelor of Engineering degree in Mechanical Engineering at Gujarat University, Ahmedabad, India, in June 2000. He received his Master of Science degree in Computer Science at Texas A&M University, College Station, USA, in May 2004.

Permanent Address:
E-4, Swapnashilp Apts,
Near ONGC Colony,
Sabarmati,
Ahmedabad 382424,
India.

The typist for this thesis was the author.