

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

UMI

A Bell & Howell Information Company
300 North Zeeb Road, Ann Arbor MI 48106-1346 USA
313/761-4700 800/521-0600

**BEHAVIORAL CHOICE IN AN ELECTRONIC
NERVOUS SYSTEM**

by

Nef Straub

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Computer Science and Engineering)
in the University of Michigan
1996

Doctoral Committee:

Professor Keki B. Irani, Chair
Associate Professor Edmund Durfee
Professor John Holland
Associate Professor John Laird

UMI Number: 9624740

UMI Microform 9624740
Copyright 1996, by UMI Company. All rights reserved.

**This microform edition is protected against unauthorized
copying under Title 17, United States Code.**

UMI
300 North Zeeb Road
Ann Arbor, MI 48103

**In memory of my mother,
who always knew I could; and with love to my best friend,
who continues to be part of my growth.**

ACKNOWLEDGEMENTS

The completion of this work is due, in part, to the influences of the following delightful people, whom I sincerely thank: Carolyn Finzel, Diane Onan, Usama Fayyad, Jane Havens, Victoria Howell, Keki Irani, Deb Kyle, W. Paul Lang, Linda Lang, Mark Latessa, Darol Straub, Lois Thompson and H. Ray Ziegler, Jr. Thanks for the encouragement!

Among those contributing to the start of this work are Bill Bunkelman, Rodney Brooks, James Grna, Robert Franklin, Jr., Stephen Grossberg, John Holland, Keki Irani, Darol Straub, Neff Walker and all the other lecturers who brought with their voices the joy of learning.

This research was supported in part by a National Science Foundation Fellowship for Women in Engineering.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	vi
LIST OF TABLES	xi
LIST OF APPENDICES	xii
CHAPTER	
I. INTRODUCTION AND SCOPE	1
Behavioral Choice and Artificial Intelligence	2
SENSI - Approach and Summary of Findings	4
Why an Electronic Nervous System?	5
Outline of Contents	8
II. REVIEW OF RELATED LITERATURE	10
Prototypes, Past and Present	11
Learning in Animals	15
Maps and Motivations	18
Development and Learning in the Nervous System	20
III. SYSTEM COMPONENTS AND VALUES	28
Motor Mechanisms	28
Sensory Systems	29
Environment	30
Drives, Reflexes, and Innate Values	31
IV. NETWORK IMPLEMENTATION OF INNATE RESPONSES	33
Neuromimes	33
Sensorimotor Circuitry	35
Recoil, Startle, and Frustrated Motion	41
Drive Implementations	46
V. ANALYSIS OF THE INNATE NETWORK	53

Neuromimes	53
Sensorimotor Circuitry	55
Mutual Inhibition – Directional Neuromimes	59
Recoil and Startle Reflexes	63
Drive Networks and Interactions	64
Mutual Inhibition Revisited – Drive Networks	71
Overall Innate Behavioral Analysis	78
Summary and Utility of Results	80
VI. NETWORK IMPLEMENTATION OF CLASSICAL CONDITIONING	84
Hardware Plasticity	84
Sampling Cells	85
Learning Network Basics	86
The Arousal Role of Drives	88
Other Drive-Related Considerations	90
Application of Learning Networks to SENSI	91
Consequences of Generalized UCS Arousal	93
Learning Refinement: Sensory Focus	94
Learning Refinement: Synapse Specifics	94
Consequences of Synapse Placement	96
Learning Refinement: Simple Reinforcement	97
VII. ANALYSIS OF THE LEARNING NETWORK	99
Lasting Fear	99
Learning and Behavioral Choice	100
Learning Network Activity	101
Behavioral Analysis	109
Plasticity and Permanence	110
Summary and Utility of Results	112
VIII. CONCLUSIONS AND FUTURE DIRECTIONS	116
Review of Research Goals	116
Accomplishments and Conclusions	117
Areas for Continued Research	123
APPENDICES	128
BIBLIOGRAPHY	144

LIST OF FIGURES

Figure

1.	Motoneuromime connections for the left motor.	36
2.	Directional control connections for the two motors.	37
3.	Influence of halt interneuron on motor interneurons.	38
4.	Diagram of the motion sensor; perspective view from bottom.	39
5.	The design and mechanics of the pain and touch sensors.	40
6.	Neuromime connections of the recoil reflex.	43
7.	The network implementation of the startle reflex involving touch from the direction opposite to motion, and sudden variations in light intensity.	44
8.	The frustrated motion response network. The extra power provided in these cases can help SENSI if it is bogged down. Note that this response also occurs if the motion sensor fails.	45
9.	The neuromime connections of the boredom drive.	48
10.	The neuromime connections of the hunger drive.	50
11.	The fear drive neuromime network. The convergence of outputs from the pain neuromimes is discussed in the text.	51
12.	Activity of a general neuromime as it fatigues.	53
13.	Initial responses of the motoneuromimes upon activation.	55
14.	Motoneuromimes after fatigue has set in for FF and FR begins to fatigue.	56
15.	The effect of activation of the halt interneuromime in the common interneuromimes, and the resulting deactivation of the motoneuromimes.	56
16.	The motion sensor's output.	57
17.	Touch and pain signals from a whisker distorted through its range.	57
18.	The responses of the light, greater, and sudden change in light senses associated with the two photocell 'eyes'. A bright light was quickly	

	passed by both lenses in this experiment.	58
19.	The response of a light sensing neuromime as a light is quickly turned from dim to bright.	59
20.	All directional neuromimes are receiving comparable excitation. The dominant forward neuromimes are beginning to fatigue.	61
21.	This real-time continuation of figure 20 shows the fatigued forward neuromimes and inhibited, but unfatigued backward neuromimes competing for dominance.	62
22.	The continuation of figures 20 and 21 illustrates the backward neuromimes gaining complete dominance over the forward neuromimes. The competitive toggling is completed.	62
23.	Excitation of the startle reflex neuromime by means of a touch on the robot opposite the primary direction of motion.	65
24.	Interruption of the excitatory input to the boredom drive, showing the inhibition of the boredom drive neuromimes and the response of the fear drive network to pain.	66
25.	The end of the pain signal and the re-activation of the boredom drive. The subliminal fear (FD activity) continues for several seconds. Times shown are actual times from the start of the experiment, shown in figure 24.	67
26.	Fatigued hunger (but no light) while the boredom drive is active. Without light, the hunger drive network cannot become active in the drive competition.	68
27.	The stimulus of light is added to the experiment that began with figure 26. The hunger drive becomes dominant over the boredom drive. There is insufficient light for feeding, and the boredom network remains partially active.	69
28.	SENSI begins to charge its batteries, inhibiting the boredom drive.	71
29.	The hunger drive network and the fear drive network, as directly involved in their competition for dominance.	72
30.	Pain excites the fear network while the hunger network is dominant.	74
31.	This real-time continuation of figure 30 shows the fear network inhibiting the hunger network. Neither drive network is dominant. The reduction of HD's activity after loss of positive feedback is evident.	75

32.	This continuation of figure 31 illustrates the fear network becoming dominant. HD is quite fatigued, and HDC remains inhibited.	76
33.	The introduction of light to a hungry SENSI which has felt pain long enough for the fear drive network to be fatigued.	77
34.	After 400 milliseconds, HD had fatigued without feedback, but FD is fatigued despite feedback. Fear is still the dominant drive.	78
35.	The hunger drive network gains control over the fatigued fear network. The effects of feedback, and lack of it, are evident.	79
36.	The initial state of the network showing the UCS and its response, the CS, and the missing connections where learning circuitry is to be inserted.	86
37.	A first attempt at the learning mechanism connections. The shaded cell is the sampling cell; plastic synapses are shown as filled circles.	87
38.	A revised learning network whereby the CS becomes conditioned both to the desired response and the arousal (labeled 'A') initially provided by the UCS.	87
39.	The learning network for the case of a stimulus conditioned to an innate reflex response. The conditioned response cannot occur in the absence of drive activity (which may have other sources of activation in addition to the UCS shown).	89
40.	The learning network for a stimulus conditioned to an innate, drive-dependent response (note the drive competition neuromime, as described in Chapter IV). Here also, the drive must be active for the conditioned response to occur.	89
41.	The additional inhibitory connections required to maintain the competition for motor control which is the basis of behavioral choice.	91
42.	The learning network showing the involved components of SENSI's network for two arbitrary whiskers. Analogous wiring is required for the remaining whiskers.	92
43.	The diagram representing our synapse-on-a-synapse structure.	96
44.	A method for implementing active, selective, positive reinforcement on the motodirectional neuromimes. See text for an explanation.	97
45.	Touch and fear combined before SENSI has learned to respond. Although the plastic synapse shows some output, it is not enough (together with FD) to fire FDR. The second sampling cell (which has	

	synapses to the motor response) cannot fire without FDR active.	102
46.	Pain on whisker 1 results in fear, sampling of FDR, and sampling of the motor reflex response.	103
47.	Pre-existing fear, combined with touch, elicits learned activity on FDR. The motor response is excited by plastic synapses emanating from S2 _{TFL} .	104
48.	An illustration of retraining the motor synapses of the front, right learning network. (a) synapse output after four learning trials using whisker 1; (b) synapse outputs after four relearning trials using whisker 10; (c) synapse outputs after three additional relearning trials with whisker 10.	105
49.	A sample of reduced learning and/or response due to inhibition of S2 _{TFL} by HDR. FDC is also being inhibited, but it is less fatigued. Fear is dominant.	107
50.	A sample of competition between hunger and fear, where hunger is regaining dominance. Learning and/or the learned response is inhibited.	108
51.	Schematic of the general neuromime, showing resistor and capacitor labels and their default values. Power connections and pin numbers (for the timer and operational amplifier ICs) are not shown. A description of the circuit is given in the text.	129
52.	Schematic of motor control circuitry. All diodes are of type 1N914; all unlabeled transistors are type 2N3906.	133
53.	Motion sensor schematic. Refer to figure 4 for actual sensor diagram.	134
54.	Schematic of the pain/touch neuromimes, with a table indicating the proper component values for each case.	135
55.	The light sensing circuitry. An explanation of function is in the text.	136
56.	SENSI's feeding circuit. A batteries charging signal occurs when the current flow from the solar cell array to the batteries is strong enough.	137
57.	The batteries low (hunger) and batteries high (sated) circuitry.	138
58.	SENSI's power bus configuration.	139
59.	A preliminary synapse diagram, indicating function. An explanation of terms and mnemonics is in the text.	140
60.	The plastic synapse. An explanation of the circuit is in the text.	141

61. **SENSI's connection overview. See table 3 for an explanatory list of neuromime labels.**

143

LIST OF TABLES

Table

1.	Connections from pain neuromimes, implementing the recoil response.	42
2.	Component values and characteristics of several neuromime types.	130
3.	Neuromime labels and their meanings.	142

LIST OF APPENDICES

Appendix

A.	Neuromime Model	128
B.	Motor Control Circuitry	132
C.	Sensory Circuitry	134
D.	Synaptic Plasticity	140
E.	Overall Nervous System	142

CHAPTER I

INTRODUCTION AND SCOPE

One of the first things the neophyte researcher must learn is . . . not to pick the most important problem for research, but the most important problem for which the tools are available and a method he can handle, with the expectation that when the first step has been taken he will be better able to see how to take the next.

Donald Olding Hebb

How do complex interconnections of relatively simple components result in the emergence of intelligent, or even coherent behavior? While examples of such interconnected structures exist within and about us, their very complexity keeps us struggling to understand their function (and dysfunction). This research explores the nature of neuronal connections which result in behavioral choice: How, when more than one behavior is possible or desirable, does a network of interconnected neuron-like components initiate a subset of behaviors at the expense of other choices?

The ability to choose behaviors conducive to survival in a complex, dynamic, and potentially hostile environment is an important facet of natural (animate) intelligence. A useful framework for researching this aspect of intelligent behavior is complex enough to incorporate its developmental context, yet simple enough to be subject to analysis. This analysis permits evaluation of the criticality of the components involved in the formation and expression of each behavior. The issue of behavioral choice is complex: available behaviors can depend on sensory capabilities, motor facilities, quick-action reflexes, internal expression of needs, methods of associating actions with reactions, and processes of storing and retrieving associations for future reference. Components involved in the dynamic expression of behavior include physical capabilities, interactions among distinct, possibly incompatible motivations, expectancies derived from stored experiential associations, and current sensory input. It is believed that intelligence does not develop in a vacuum; the

how's of intelligence cannot be properly understood without taking into account the *what's* and *why's* — what is learned, and why it is important. If we understood the dependencies of intelligent behavior on the motivations that drive it and on the capabilities which affect its development and expression, then we might better understand the processes by which experience is accrued and adaptively utilized.

This chapter discusses the fundamental nature of choice with regard to intelligence, describes where this research fits with respect to other work in behavior mediation strategies, and provides the rationale for the approach. An outline of the remaining contents is given.

Behavioral Choice and Artificial Intelligence

The study of artificial intelligence, or the attempt to create mechanisms of intelligence in synthetic life-forms, is a complex set of subfields interacting in untold ways. This section illustrates how choice is a major issue in many artificial intelligence (AI) paradigms to date, and reviews a number of general choice mechanisms which have evolved for autonomous systems.

Many formalized AI techniques revolve around the notions of search and problem reduction. It is revealing to consider that effective search means *making good choices* about where and how to look, and that efficient problem reduction involves *making good choices* about problem division. The process of making choices in these situations is usually referred to as a *control strategy*. For instance, whatever the control strategy chosen for search — be it forward, backward, heuristic, and-or tree, means-end analysis — each strategy is attempting to optimize the way in which choices are made such that search is either minimized or exhaustive.

One goal of the application of AI to robotics is the development of robust, autonomous, task-oriented robots. The choice mechanisms used in attempts at this goal are varied. Some researchers (Fikes & Nilsson, 1971; Sussman & McDermott, 1972) have taken a global, planning approach, using traditional AI search procedures within a problem domain to develop a sequence of steps from start to goal states. Others (Walter, 1953; Braitenberg, 1984) have described more local, action-based schemes. Global planning generally involves a software loop which compares the current state against a set of conditions, evaluates the comparison, and applies an action rule based on the evaluation. This has been referred to as a “sense-think-act” paradigm (Pfeifer & Verschure, 1992). The more reactive schemes omit the “think” step in what might be called sense-act systems. Neither of these approaches

has achieved the robustness of behavior AI researchers strive for. The sense-think-act paradigm suffers primarily from a lack of responsiveness, the sense-act scheme, from a lack of directedness.

Autonomous robots ideally have both the reactivity to deal with a dynamic environment and the directedness to work toward non-immediate goals. Hierarchical control strategies, where lower levels provide reactivity and higher levels pursue goals, have received much attention as a potential solution to achieving both. Examples of systems utilizing hierarchical decision making are ABSTRIPS (Sacerdoti, 1974), NOAH (Sacerdoti, 1977), and MOLGEN (Stefik, 1981). Distributed hierarchical control paradigms, containing parallel processes within levels which compete for control of the levels' outputs, have also been a subject of research. Examples of this approach include blackboard architectures (Nii, 1986a; 1986b) and, to some extent, subsumption architectures (Brooks, 1986).

A related method of handling the responsiveness/directedness dilemma is that of finer-grained control schemes such as parallel distributed processing (Rumelhart et al., 1986), classifiers (Booker, 1982) and neural networks (Beer, 1990). These systems utilize numerous, simple decision makers, which compete with each other for the privilege of being active. Active processes then converge on an appropriate action (the convergence mechanism is the choice mechanism, here). While hierarchies of control may be identified in this type of system, they are often not clearly defined.

Finer-grained control schemes have an associated complexity that soon has the researcher looking for ways for the system to "tune" itself. (Kaelbling, 1992) describes a software algorithm which uses a statistics-based reinforcement strategy for adaptation. (Booker, 1982) used genetic algorithms (Holland et al., 1986) to provide the means for adaptability in his insect simulation. A particular example of learning in a hardware implementation of neuromimes is (Walter, 1953); the development of learning algorithms for use in neural networks is an ongoing field of research (Anderson & Rosenfield, 1988; Anderson et al., 1990).

(Pfeifer & Verschure, 1992), in an excellent summary of the need for adaptation in general, propose a strategy for autonomous robot designers which they call distributed adaptive control (DAC). DAC involves designing robots such that adaptive behavior emerges from basic reflexes and drives by means of associations built up within a neural network. Their paradigm fits so well with this research that I feel it appropriate to quote the DAC development steps (pp. 22,23):

1. Define the physical setup of the robot with its sensors and effectors and specify the types of environments in which the robot is to function. This is the framework determining the interaction potential of the robot.
2. Define the "value scheme". The value scheme comprises a set of basic sense-act reflexes (direct connections of sensory patterns to motor actions), a set of elementary drives, as well as the specific properties of the sensors and effectors. The value scheme is pre-defined. It provides the constraints on the self-organization.
3. Define the network architecture and the learning mechanisms. The assumptions about the value scheme are translated into a network architecture. Sensory input is projected into a set of corresponding network layers. A minimal set of assumptions to achieve adaptive behavior is (a) that the sense-act reflexes be implemented by linking the respective sensory layers to the motor system, and (b) that there be an associative mechanism to enable the transfer of these reflexes to more sophisticated sensors.
4. Let the robot interact with its environment. Analyze the robot's behavior by informal and systematic observation, and by formal analyses. Relate its behavior to the dynamics of the network.

The choice of network architecture is left to the designer using the DAC strategy. Many options exist, both in the form of the neuron model and in the structure of the interconnections.

SENSI - Approach and Summary of Findings

The purpose of this investigation is to explore mechanisms of behavioral choice within a simple network of neuromimes. Following the DAC design strategy, the research is accomplished via a mobile robot. Our robot, SENSI, is a tracked vehicle incorporating light, motion, and contact sensors, and motor actuators. The value scheme is defined by the interactions of recoil and startle reflexes, and the hunger, boredom and fear drives. The robot is controlled by a heterogeneous neural network, much like (Beer 1990), but the network is organized in part using principles developed by Grossberg (Grossberg, 1988), which outline network architectures for drive, attention, and adaptation interactions. Our neuromimes incorporate input summation, inhibition, threshold, and fatigue characteristics.

The network is affected by external sensory events and internal needs, and it mediates among available motor actions. Under appropriate conditions, it can adapt

itself, transferring reflex responses to sensory events other than the original reflex-generating stimuli. The characteristics of the network's components and connectivities, which lead to choice and learned choice, are the foci of our work.

Issues of choice occur within the innate behaviors (e.g. directional choice given conflicting reflexes), among the drives, and as an integral ingredient of effective learning. This is not a simulation. This research was conducted with a mobile robot for which sensory and motor apparatus are noisy and imperfect. The choice mechanism is an electronic nervous system, and each of its handwired neuromime components is unique in its combination of imperfect resistor, capacitor, and diode characteristics. It is within this context of imperfect reality that robust behavior mediation strategies are sought.

What emerges from our work is the need for certain neuromime characteristics such as spike-and-decay output signatures and fatigue. We have found a form of dendritic tree to prove useful for normalization of inputs to individual neuromimes. We also determined a necessity for certain interconnection structures such as mutual inhibition and positive feedback. Our incorporation of learning leads to several insights regarding connections which realize ongoing plasticity versus those which result in permanent associations. Plastic synapses which include their own learned strength during sampling form permanently effective connections. Those which exclude their learned strengths are subject to dynamic decay with use. Active reinforcement is required to maintain useful learning in the latter case.

Why an Electronic Nervous System?

There have been recent studies of behavioral choice within neural networks (Maes 1989; Beer 1990; Grossberg 1988). These studies have primarily been conducted in theoretical or simulation contexts. It is the belief of this author (and several other researchers (i.e., Brooks, 1986; Pfeifer and Verschure, 1992; Patel & Schnepf, 1992)) that behavior is too environment-dependent to investigate thoroughly without direct environmental interactions. While simulations have much to offer in terms of exploring theoretical issues, an interactive robot, which must deal with real objects using real (imperfect) sensors and actuators, provides a more valid indication of a system's robustness *in the environment it was designed for*. Such a system is also the focus of attention for achieving groundedness and situatedness: "If a system is 'situated' it can get the information it needs to act directly from the particular situation it is in" (Pfeifer & Verschure, 1992).

Perhaps a hybrid system, with real sensors and actuators, and a simulated (software) neural network, as implemented by Pfeifer and Verschure, is an optimum approach to this issue. There are several reasons why this approach was not taken here. These reasons and related discussions are enumerated here:

(1) The first and foremost reason was a desire to use real-time, parallel processing at the neuron level.

What I am arguing against, here, is the use of a state space approach to neural networks, where a software loop examines the status of each neuron and alters it according to the inputs that impinged upon it a loop-cycle ago. With the increasing processor speeds available, and the relative slowness of the nervous system's signal processing, it is likely that these loops can be frequent enough, even with a large number of modelled neurons, to adequately simulate the dynamics of each neuron. There are advantages to such a simulation, for example (a) the activities of any individual neuron can be saved, examined, and analyzed; (b) any neuron's parameters or connections may be simply and effectively altered; (c) learning algorithms, even growth algorithms, can be easily implemented and tested; and (d) because of its location within a separate processing system, knowledge gleaned by a learning simulation can be immediately 'transferred' to any other compatible system just by copying the data structure of the network.

So, why hardware? A strong advantage of the nervous system over artificial controllers is the stability of its behavior even given the loss of a number of neuron components. This facet of nervous system models has no equal in a software simulation; the functioning of the simulated network is dependent on the functioning of the underlying processor. Also, everything biological has its quirks and foibles, i.e. no two neurons are likely to be exactly equivalent or entirely predictable in their behaviors. But they are in simulation. One wonders how much the inherent diversity and fuzziness of individual neurons' behaviors contributes to the robustness of the system as a whole. (While this is an interesting point it won't be answered by this research, either!)

Finally, and this is a matter of subjective evaluation, there is always a bias associated with the tools used to create a system. The preference here is to deal with hardware, electrical artifacts rather than software, programming language and data structure artifacts. The source of this preference is the nature of the model. What we are attempting to model is the complex system of unique, dynamically interacting components, the interactions of which define the behaviors of animate,

autonomous life. By creating the model in electronics, where each component is unique and the interactions are by nature (not algorithmic design) dynamic, the entire flavor of the research is more realistic.

(2) Another reason for developing a hardware nervous system was to explore alternative computer architectures from a neural network perspective. Content-addressable memory, real-time priority-based scheduling, evolution of knowledge structures and dependency relationships, and intelligent interactive computing are just a few of the benefits that the neural model offers. First, however, the neural model must be examined. The necessary and sufficient features of the nervous system need to be identified.

While there are many roads to this goal, the vision here is first the development of a functional, real-time hardware system with adaptation. Explorations of necessary and sufficient neuron features come from a modifiable, testable neuron model. With this model established, the door is opened for the miniaturization and mass-manufacturing that characterize today's computer hardware. With mass production, further experiments in connectivity constraints are realizable. From that point, innate systems could be manufactured; different systems could be tested in varying conditions to establish fitness, and so forth – until the dream of accomplishing the successes of animate life in hardware becomes the reality.

(3) Finally, there was a strong inclination to build an architecture in which nonlinear behavior was evident, and chaotic behavior probable. As the field of chaos theory has expanded, the application of its analyses has extended to numerous biological processes. (Gleick, 1987) notes:

Many other scientists began to apply the formalisms of chaos to research in artificial intelligence. The dynamics of systems wandering between basins of attraction, for example, appealed to those looking for a way to model symbols and memories... With or without chaos, serious cognitive scientists can no longer model the mind as a static structure.

Grossberg, in his lectures at the Wang Institute, has stressed that the interactions occurring in the nervous system which lead to learning are nonlinear, nonlocal, and nonstationary. While the potential chaotic behavior of our network is not analyzed here, it is also not constrained out of the research. Our use of hardware and its accompanying non-linear behavior is an interesting subject for further work.

Further inspiration to work along these lines comes from a view of neural network diversity well characterized by (Gleick, 1987) in his discussion of the chaotic structure of snowflakes:

Sensitive dependence on initial conditions serves not to destroy but to create. As a growing snowflake falls to earth, typically floating in the wind for an hour or more, the choices made by the branching tips at any instant depend sensitively on such things as the temperature, the humidity, and the presence of impurities in the atmosphere. The six tips of a single snowflake, spreading within a millimeter space, feel the same temperatures, and because the laws of growth are purely deterministic, they maintain a near perfect symmetry. But the nature of turbulent air is such that any pair of snowflakes will experience very different paths. The final flake records the history of all the changing weather conditions it has experienced, and the combinations may as well be infinite.

Outline of Contents

This project necessarily involves a host of subproblems including engineering, robotics, neuromime design, sensor and motor selection and implementation, mathematical analysis, and so forth. Since the focus is on electronic nervous system connections for behavioral mediation, the body of this treatise is devoted to issues at that level of abstraction, namely, the structure of neuromimes and their connections which result in innate biases, individual behaviors, and learned associations. The other concerns are summarized where appropriate and detailed in the appendices.

Chapter II contains a review of relevant literature concerning the nervous system in general, innate associations or reflexes, and theoretical mechanisms of learning new associations and behavioral responses. SENSI's components and environment are discussed in Chapter III. Chapter IV is concerned with the neuron model, sensory and motor control features, reflex connections and drive implementations. In effect, Chapter IV describes the implementation of the innate structures outlined in Chapter III. An analysis of individual neuromime behaviors, and of the causal behaviors of selected networks of neuromimes is given in Chapter V. This chapter concludes with a summary of results from the innate network, and their applicability to other research.

Chapter VI discusses plasticity and experience-based adaptation, in other words, where and how learning occurs in our system. An analysis of the functionality of the learning system follows in Chapter VII, along with a summary of conclusions and their relevance. Chapter VIII concludes with a discussion of results,

needed improvements and possible enhancements. The appendices are devoted to the general neuromime schematic and analysis, motor control circuitry, sensory interface systems, plasticity schematics and analysis, and the overall nervous system layout.

CHAPTER II

REVIEW OF RELATED LITERATURE

But if the performance of a model is to be demonstrably a fair imitation of cerebral activity, the conditions of stimulation and behavior must equally be comparable with those of the brain. Not in looks, but in action, the model must resemble an animal. Therefore it must have these or some measure of these attributes: exploration, curiosity, free-will in the sense of unpredictability, goal-seeking, learning, forgetting, association of ideas, form recognition, and the elements of social accommodation. Such is life.

W. Grey Walter

Recognizing that behavioral choice is the observational correlate of what is often termed intelligence, this chapter seeks to summarize aspects of research regarding intelligence and learning in animate systems. This is such a rich area of literature that no attempt has been made at completeness; rather, influential works in selected related fields are reviewed. A top-down tour is taken through considerations involved in modeling intelligence. The first section briefly describes several efforts at creation of devices which interact with, and learn from, their environments.

A categorization of animal learning is presented in the second section. By providing observational descriptions of increasingly complex learning abilities, we illuminate varying degrees of intelligent behavior and thereby provide a guideline for analysis and evaluation of modeling results.

The third section briefly reviews theoretical issues pertaining to the organization of stored experiential knowledge. The purpose of this is to indicate possible roles of perception and motivation in the determination of which stimuli should be attended, and how associations might be arranged in memory to facilitate recall, prediction, generalization, and the like.

A sequence of related theories concerned with neural organization, development, and the coding of associations is examined in the fourth section. It is intended

to illustrate how the mechanisms of the third section might realistically occur in the nervous system, and to review research upon which many neural modeling efforts are based.

Prototypes, Past and Present

Models of living systems which sense and interact with their environments provide immediate feedback to the designer with respect to their coping skills. Some researchers, such as (Brooks, 1991), feel this mastery of successful environmental interaction is a major portion of intelligence itself:

I argue that [perception and motor skills] are the hard problems solved by intelligent systems, and further that the shape of solutions to these problems constrains greatly the correct solutions of the small pieces of intelligence which remain.

This approach is shared by (Kaplan, 1973), who hypothesizes that “the structure underlying the spatial map of the world... is not different from the structure that underlies *all* cognitive processes.” The interactive being must be able to learn the features of its environment which affect it, store abstractions of these features along with their positive or negative entailments, be able to recognize them from varying and obscure inputs, and act appropriately to ensure survival. Much of what is grouped under intelligence — classification, abstraction, recognition, inference, and planning — is contained in the effective realization of sensory and motor capabilities.

Thus our first interest is to examine models which have been built to interact with their environments. Several electro-mechanical models are examined here. Each is simply, yet functionally built. They are designed to interact in specific ways with respect to the environment within which they are placed. While these may seem to be mere engineering feats, so much of learning appears to be based on the mobility which these robots achieve that it is felt these attempts begin to capture significant roots of animal intelligence.

The Tortoise. In 1953, W. Grey Walter wrote about a robot, “*Machina speculatrix*, inevitable name of the species for the discerning, though ‘tortoise’ to the profane,” which he had created from two tubes, two relays, two condensers, two

small electric motors, two batteries, and various supporting hardware. This simple electro-mechanical system wandered, avoided slopes and obstacles, looked for and moved toward light of moderate intensity, and avoided very bright lights.

Walter's learning research centered around the implementation of conditioned reflexes. The tortoise was enhanced to respond to sound, and various associations were investigated between light, sound and touch. Since light and touch were preset to evoke distinct behavioral responses, learned associations were readily observable, e.g. learning 'sound-means-light' would be demonstrated by an attraction of the mechanism to a whistle; that of 'sound-means-touch' would result in evasive behavior after any significant noise.

Investigations of conditioned reflexes led Walter to break this form of learning down, first into two processes, then into seven distinct steps. The processes are those of selection and construction: selection of the subset of sensory input deserving of attention, and construction of the memory of the sensory events and their association. In other terms these are the processes of discovering salient patterns, of forming a means for recognizing the initial pattern when it next occurs and for predicting, or anticipating, the associated pattern or event. With his analysis, Walter constructed a circuit which embodied the seven steps. It was taught to associate sound with light, sound with touch, light with sound, and touch with sound. The resultant behavior was in accord with Walter's predictions and also provided a few insights into the nature of memories and neuroses. It seemed to mirror, quite adequately, some of the behaviors exhibited by the animals after which it was patterned.

It is surprising, therefore, that this learning mechanism wasn't extended to more basic, internal associations. In particular, had Walter provided the tortoise with an internal stimulus indicating low power level, perhaps the mechanism would have learned to associate light with power, and taken a step toward self-sufficiency. While external stimuli form a basis for learned behavior in animals, prediction, planning, and so forth seem to require increasing amounts of internally generated signals.

Toto. Brooks and fellow researchers at MIT (Brooks, 1986; Mataric & Brooks, 1990) have developed several varieties of mobile robots in the past few years. They are built in layers of activities, or behavior-producing systems, each of which is a stimulus/response interaction with the world. These activities, implemented as augmented finite state machines, interact via suppression of input and

inhibition of output. Otherwise stated, a higher level may subsume the effect of a more primitive level; hence, the overall paradigm is called a subsumption architecture. Because the behavioral systems are added incrementally, one activity's response can provide a stimulus for a different activity.

The robots, called Creatures, have operated with some interesting behaviors. Varieties have been built which avoid light, which follow heat, and which move toward sound. All are provided with locomotion and obstacle avoidance abilities. Most of these Creatures do not learn, but behave in a predetermined, reactive manner. An exception of interest is the work of Mataric and Brooks (Mataric & Brooks, 1990) on Toto, a Creature which builds a map of its environment. What makes Toto different is an architecture which permits the environment to permanently affect some of the Creature's activities, i.e., during traversal of some area, the features of the areas are internally recorded for future reference, and thereafter affect the Creature's perception of its world.

Toto's architecture consists of three primary layers: collision-free wandering, landmark detection, and map learning/path planning. Collision-free wandering is the behavior emerging from the four navigational rules of stroll, avoid, align, and correct. The visible attribute of this behavior is the tracing of environmental boundaries. Via the use of a compass, sonar proximity sensors, and the obstacle-tracing behavior, the landmark detection layer extracts environmental features from the Creature's motion. As Toto wanders around the boundaries of its office environment, hypotheses are entertained as to the type of landmark to which it is adjacent. Consistent sensor readings adjust confidence levels in each hypothesis until a threshold is attained and the associated landmark is acknowledged (i.e., a match is chosen).

Once a landmark has been identified, it is broadcast to the map building layer. The map is an active graph of nodes, with each node corresponding to a landmark. The active node always reflects Toto's position within its environmental map. Since each node is an activity, capable of receiving and transmitting messages, path planning is simply accomplished by a message broadcast from the goal node. In effect, the first such message to reach the active node is retraceable to the goal and presents the most direct route to it.

(Mataric, 1990b) has pointed out several similarities between Toto's navigational system and that of a laboratory rat. The rat seems to have a distributed spatial representation based on landmarks. Many researchers have investigated the function of *place cells*, neurons in the rat hippocampus which fire with maximum

frequency when the rat is in a certain absolute position and orientation, called a *place field*. Each of Toto's graph nodes can be thought of as representing a place-cell, with its place-field being the associated landmark. Experiments also indicate that, in the rat, proximal landmarks are often recalled simultaneously, motion changes affect place-cell activity, and goal locations are externally selected based on animal motivations.

To better understand the limitations of Toto's navigational skills, it is helpful to consider the simplifications, with respect to a rat, under which it operates. Foremost, it need not avoid threats or find sustenance. This means it does not have to perceive or remember environmental utility. Any biologically sound cognitive map must contain semantic, or utility information to be useful to the animal. "The capacity to recognize some familiar landmark... helps very little if one does not have any other information connected to it" (Kaplan & Kaplan, 1982). Second, no provision has been made for altering the map structure given environmental changes; this is a monumental simplification given the rat's capability to deal with complex and *dynamic* environments. The rat hippocampus contains neurons similar to place cells but which are responsive to unrealized expectations about the current location (Nadel, 1980). Finally, the office environment in which Toto operates is fixed enough in structure that the Creature was built with a static and very limited set of landmark types which it could recognize.

Lola. (Pfeifer & Vershure, 1992), using their DAC (distributed adaptive control) paradigm, have developed a control strategy which demonstrates adaptability in the robot Lola. DAC, as described earlier, involves specification of the robot's sensorimotor capabilities, environment, basic values, and learning mechanisms. The purpose of designing in this manner is the robot-centeredness of the process, i.e., basic parameters chosen from the robot's point of view are augmented by learning strategies such that the robot's experience leads to adaptive behaviors.

The design of Pfeifer and Vershure was first implemented in simulation. The robot was equipped with sensors for collision, target, and range detection, and was capable of locomotion. The environment specified was a walled enclosure with various obstacles and an optional target. There were two reflexes: collision leading to retract and turn, and target detection resulting in motion toward the target. The single implemented drive was for forward motion.

Their architecture was a simulated, layered neural network. There were sensory, intermediate, and motor layers. The intermediate layer contained approach

and avoidance areas, and were connected such that avoidance was dominant. Weights in the network were adjusted in a Hebbian fashion (coincident firing of connected neurons resulted in an incremental strengthening of the connection), with active decay (while a neuron was firing, its connection weights would incrementally decay unless reinforced by post-synaptic firing).

In simulation, the robot initially collided with the wall and obstacles. Through learned associations, the avoidance mechanism resulted in less collision and relied more on the range finder. Thus the mechanism *learned via conditioning*, to avoid collision. After such behavior was learned, the robot could apply it in novel ways. For example, the robot was capable of backing out of a cul-de-sac without collision. Experiments were also performed with a target present in the environment. If the target was placed behind a wall, the robot demonstrated edge-following behaviors (learned through the interaction of the attraction and avoidance mechanisms).

The trial of this mechanism on the robot Lola differed in sensor types utilized (e.g., infrared was used) but was quite successful. The authors note that Lola without learning tended to get stuck in corners, but with learning this cyclic trap resolved itself. They also discussed the apparent “anticipation” of collision exhibited by the robot. While the network itself contained no strategy for short-term memory, the robot behaved in a manner which indicated its existence. The emergence of this behavior from the simple network illustrates both the potential power of this approach and the intrinsic fallibility of the attempts to design networks based upon desired, observable behaviors. In the words of (Braitenberg, 1984), “Interest arises... when we look at these machines or vehicles as if they were animals in a natural environment. We will be tempted, then, to use psychological language in describing their behavior. And yet we know very well that there is nothing in these vehicles that we have not put in ourselves.”

Learning in Animals

Many researchers (e.g. Hebb, 1949; Thorpe, 1956; Holland, 1975; Charlesworth, 1976; Booker, 1982; Nelson, 1989; Brooks, 1991) feel that intelligence and learning are functionally described and demonstrated by an organism's environmental interactions. (Kandel, 1979) defines learning as “the ability to modify behavior in response to experience.” This section explores a categorization of learned behaviors according to complexity, presents examples of observed occur-

rences of these behaviors in various species of animals, and provides a guide for evaluation of models designed to emulate animate behaviors.

The distinctions between types of learning are often subtle, and in this categorization are defined along observational lines. (Thorpe, 1956) notes:

One of the greatest difficulties encountered in the separation and classification of different types of learning is the fact that we can hardly ever find in nature an example of a learned response which can be regarded as belonging to one type and one type only.

Thorpe hypothesized several levels of learning, increasing in their complexity. The simplest type seems to be that of *habituation*, defined as learning not to respond to stimuli which lack significance to the well-being of the animal. Occurrence of habituation is realized by the tendency for an animal to quit responding to repeated, similar stimuli which otherwise portend no relevant consequences. This type of learning appears to be universal in animals, from the protozoa through the birds and mammals. Its universality can be understood by considering the impracticality of animals possessing a specific response to any specific stimulus. Rather, they have an instinctive self-preservation response to anything new or predator-like, and learn with experience which stimuli do not seem to represent an impending threat to life.

The next simplest type of learning is a process of association, as in *classical conditioning*. It often occurs in conjunction with habituation, as objects which fail to satisfy some need become less and less attended, while those with positive aspects become associated with reduction of the internal motivation. This is the form of learning studied by Pavlov in the conditioned salivation experiments and by Walter in his tortoise model. (Kandel, 1979) describes *sensitization* as "the prolonged enhancement of an animal's pre-existing response to a stimulus as the result of the presentation of a second stimulus that is noxious. This would seem to be a form of negatively reinforced conditioning.

Simple conditioning is difficult to distinguish observationally from the next two forms of learning, called *instrumental conditioning* and *trial-and-error learning*. It is perhaps best to use the words of Thorpe, who makes the distinction:

...we see that classical conditioning is the establishment of an association between a normal reward or reinforcement and a new external or exteroceptive stimulus, which is initially indifferent in the sense that it does not innately release any specific responses and so does not originally have any 'meaning' for the animal. Instrumental condition-

ing, on the other hand, is the establishment of an association between a voluntary motor act as part of appetitive behavior primarily as perceived by the animal's proprioceptive organs, and the normal reward or reinforcement which again involves a consummatory act. Thus both combine to make up the full normal process of trial-and-error learning; the first by conditioning the external situation to the innate releasing mechanism, the second by conditioning a voluntary motor act of the appetitive behaviour to the innate motor mechanism of the consummatory act.

An example of a combination of habituation and instrumental conditioning which is often confused with trial-and-error learning is found in chickens. For a chick, the pecking action is innate, or instinctive. However, what to peck at is not. Through finding seeds and grain pleasant to ingest, and leaves and pebbles less satisfactory, the chick learns to associate the proper substances with reduction of hunger. It is the lack of simple conditioning which keeps this example from being an instance of learning by trial and error.

According to Thorpe, trial-and-error learning is the most important learning process involved in adaptation of voluntary responses to environmentally posed situations. He feels that any adaptive behavior, no matter how insightful or intelligent it may appear, involves trial-and-error learning to some extent. In fact, it is felt that play in the higher animals involves a good deal of this form of learning, whereby the animals become practiced at the motor actions required for survival.

Play enters into the next form of learning as well—*latent learning*. There are three main characteristics of latent learning: it is learning devoid of specific motivation; its result is a training later applicable to motivational needs; and it enables an animal to selectively attend to a previously learned whole, without trial-and-error, in relation to current needs. This behavior was discovered in the rat, which tends to explore a new territory without any particular goal. Then, having learned the characteristics of the terrain, it can easily relocate resources when necessitated by hunger, thirst, or fear.

Fully half of Thorpe's treatise is devoted to a review of experimental evidence of learning in the various species. These examples include habituation in protozoa, trial-and-error learning in earthworms, maze learning in turtles, play in birds, and latent learning in birds and fish.

An interesting aspect of the categorization of learning presented is the dominance of association as a learning paradigm. Habituation is an association of non-importance with a specific stimulus. Conditioning is an association of stimulus with

effect, and of action with effect. Trial-and-error learning, as defined, then appears to be a meta-association. Finally, latent learning associates various stimuli with potential future uses. How are these associations formed and represented? What mechanisms help an animal discriminate between salient and meaningless stimuli? The next sections address hypotheses which have been formed with respect to these issues.

Maps and Motivations

Foremost in theories of attention and memory is the process of perception. "What is learned is in terms of what is perceived; what is not perceived can hardly be remembered" (Hebb, 1949). (Kaplan & Kaplan, 1982) contend that perception and cognition are intimately related processes. Surely, most animals are aware of their environment only to the extent that their sense organs permit. For the animals we are trying to model, "it seems probable that the perception of stability or variation in phenomena is important in defining and delimiting all concepts and schemata" (Vernon, 1966).

"Perception is not simply an automatic response to simple sense-data...but is, on the contrary, an active organizing process" (Thorpe, 1956), and what animals organize is the plethora of incoming signals to the sense organs. According to (Uhr, 1966):

The perceptual mechanisms of living organisms have developed around wavebands of energy that are commonly emitted by objects in our physical world... The purpose of perception is to reduce the signals that the mechanism senses...and to judge whether [the reduction] belongs to any of a class of signs that are of interest to the organism because they suggest actions that it should take. The judgment that some part of the flow of experience belongs to such a class is 'pattern recognition.'

Each animal species has evolved to be attentive to certain salient stimuli or patterns (Young, 1966; Booker, 1982); often these patterns are related to motivations or drives. The behaviors elicited by the presence of these stimuli are the reflexive and instinctive behaviors. Patterns, definable as any sequence of events in time or any set of objects in space, distinguishable from or comparable with another sequence or set (Walter, 1953), become the building blocks of memories and associations. Innate responses to fixed patterns are the product of evolution; learned responses to the

patterns of a dynamic environment are the product of experience.

In what form are these patterns and associations stored? "The problems of recognition, spatial understanding, dealing with sensor noise, partial models, etc. ...are all tied up with the representation of the world used by an intelligent system" (Brooks, 1991). One proposal which has gained prominence in the research community is that of world representation within a *cognitive map*. A cognitive map is the representation of relative locations and attributes of phenomena encountered in an organism's everyday spatial environment (Downs & Stea, 1973). (Kaplan, 1973) postulates four major types of knowledge stored within such a map:

1. Where one is;
2. What is likely to happen next;
3. Whether it will be good or bad;
4. Some possible courses of action.

The first supposes that spatial information is stored; the second presupposes an association between the current situation and some other pattern(s). The third knowledge type involves an attribute or evaluation associated with the expected situation. The fourth suggests an association with some successful, appropriate behaviors with which to meet the anticipated event(s). These four types of knowledge are also referred to as recognition, prediction, evaluation, and action information (Kaplan & Kaplan, 1982).

Represented locations and events are likely to contain common features (subpatterns) causing overlaps among them. Thus "sequences of experiences are stored as overlapping patterns, yielding a whole network of associations" (Kaplan & Kaplan, 1982). This network of associations is felt to be both an economical means of experiential storage and a basis of new sequence generation due to existing cross-over among patterns held in common by different experiences.

Researchers stress the necessity for evaluative codes to be associated with these stored patterns (Nadel, 1980; Downs & Stea, 1973). The cognitive map is assumed to include codes relating to basic biological motivations and to pleasure and pain. In terms of storage, these codes may be represented in the map in a fashion similar to that of sensory properties (Downs & Stea, 1973). More importantly, the existence of such motivation may play a crucial role in an organism's determinations of which incoming signal patterns to attend to, thereby shaping the contents of the cognitive map (Routtenberg, 1979).

In summary, learning seems to be a process of forming associations within a representation of environmental features. The building blocks of the representation

are patterns; they may be temporal, spatial, or motivational (evaluative) patterns. What patterns are stored is much dependent on the evolved sensory capabilities of the organism as on the complexity of the environment. What an organism is capable of learning appears to be dependent both on sensorimotor capabilities and associational capacities of its nervous system (Thorpe, 1956). How such a representation might be embodied within the nervous system is considered next.

Development and Learning in the Nervous System

The human nervous system contains billions of neurons, which are cells with input dendrites and an output axon. An axon generally branches extensively to make contacts (synapses) with hundreds to tens of thousands of dendrites. Connections between neurons are inhibitory or excitatory. Each neuron has a varying excitation threshold which, if exceeded via overall dendritic inputs, will result in the neuron emitting an action potential (firing). When a neuron fires, no amount of excitatory input can make it fire again within an interval called the absolute refractory period. This period is equivalent to the neuron having, briefly, an infinite threshold. The threshold then gradually lowers (relative refractory period). A strong stimulation fires a cell more frequently by being able to exceed its threshold earlier in the relative refractory period. In this manner, intensity translates to frequency in the central nervous system (Hebb, 1966). Neurons are also thought to prime if inactive or inhibited for an extended period of time; this may result in spontaneous firing.

The neuron described above is actually an idealized model. In fact there are many different types of neurons in the nervous system. Some have no obvious axon; some dendrites form synapses on other dendrites; an axon terminal may itself receive a synapse, causing a 'local' effect (Crick & Asanuma, 1986). The activity outlined above occurs by means of both electrical and chemical signals. A neuron's 'spike' is often a release of neurotransmitter, of which a variety have been identified. Neurotransmitters are released in quanta from vesicles at synaptic endings. The number of quanta released at a given synapse is related to the strength of the connection or the amount of influence the presynaptic neuron has upon the postsynaptic neuron (Kandel, 1979).

Not only are there various neurotransmitters released at synapses, there are different chemical receptors on the receiving dendrites. The receptor type seems to

determine the function (inhibitory or excitatory) of the synapse for a given transmitter (Kandel, 1979). Further complicating the interactions, there have been identified a number of neuropeptides which, when released at a synapse, appear to modulate synaptic function rather than activate it. The effect of these peptides is more diffuse and longer lasting than that of the transmitters (Crick & Asanuma, 1986). The cellular picture of the neuronal system is thus one of many parallel interconnected cells communicating among themselves by electrical, temporal, and chemical signals. This system is responsible for the behavior, memories, and learning abilities of animate life.

Different areas of the brain seem to serve different functions. In mammals, the cerebellum controls balance and movement, and the spinal cord supplies the brain with sensory input, directs muscle control, and directly controls simple reflexes (Coward, 1990). The hippocampus seems to be the center of spatial learning and memory (Olten et al., 1979). The cerebral cortex "seems highly specialized to represent and process detailed and complex sensory and sensorimotor information...[and] appears essential for language and complex spatial memory in humans. It is certainly necessary for cognitive functions" (Thomson et al., 1983). Thus a functional view of the brain has emerged which shows specialization of regions to certain behaviors, with appropriate interchange of information among regions.

Lashley. The manner in which this complex mechanism stores memories and accommodates learning has been a topic of interest to researchers and philosophers for centuries. This review covers only a few of the more prominent theories of this century, in particular, those which present a system approach to learning through neural-element changes. K. S. Lashley (Beach et al., 1960), known for disabusing researchers of the concept of local memory traces, concluded after years of study and experimentation that:

The cortex must be regarded as a great network of reverberatory circuits, constantly active. A new stimulus, reaching such a system, does not excite an isolated reflex path but must produce widespread changes in the pattern of excitation throughout a whole system of already interacting neurons.

In the same essay, he presented his theory of cerebral organization and function. Lashley proposed that the cerebral network was composed of a large number of 'trace systems', where each system consisted of numerous traces of habits

or memories. The systems were not anatomically separate; a neuron might be a member of many systems. Traces within a system were more connected with each other than with any other system. An external stimulus could activate a particular set of traces. This tonic activity would make remaining traces of the system available for recall and would inhibit other entire systems—an explanation of attention. A system which had recently been fully activated was assumed to retain a high level of sub-threshold activity, temporarily influencing other systems. Such a dominant system was believed to contain memories.

Hebb. While not directly concerned with development and learning, Lashley's theory and experimental work paved the way for investigation of ways by which the nervous system might form the distributed memories he hypothesized. One of the most influential results has been D. O. Hebb's 1949 essay, "The Organization of Behavior." In this book Hebb presented his theory of the formation of cell assemblies and phase sequences, and advocated learning via the alteration of neural connection strengths.

Hebb's basic premise was that a growth process accompanying synaptic activity makes that synapse more readily traversed. More precisely,

When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased. . . . The general idea is an old one, that any two cells or systems of cells that are repeatedly active at the same time will tend to become 'associated so that activity in one facilitates activity in the other.

In Hebb's model, some neurons are innately sensitive to various features of sensory input. As multi-featured and sequenced stimuli are repeatedly experienced by an organism, the temporal qualities of neuronal responses result in selective connection enhancement. The cell assemblies thus organized reflect an association of stimulus elements; each assembly is in effect a pattern recognizer, with neurons corresponding to distinct features shared among assemblies. Cell assemblies resulting from repeated experience of objects and temporal events form the system's memories. Since connections are strong, assemblies with neurons in common may excite one another. A sequence of such activations, called a phase sequence, was used to explain anticipation, prediction, and 'trains of thought.'

One the whole, Hebb's system bears close resemblance to the tonic organiza-

tion proposed by Lashley. Traces correspond to cell assemblies; both share distributedness across collections of neurons, and share neurons among collections. Systems, or related sets of traces, correspond to phase sequences, and both are used to describe a mechanism of attention. What Hebb has added is a plausible theory of trace or cell assembly formation which takes into account the functional distributivity of the nervous system demonstrated by Lashley.

Edelman. The cell assembly theory did not address the issues of evolutionary effects and phylogenetic development. "Hebb imagined that the intrinsic original connectivity structure of the network was of secondary importance and might for theoretical purposes be considered as random or homogeneous" (MacGregor, 1987). This assumption has been countered by (Edelman, 1978) whose theory of neural Darwinism excludes the possibility of randomly connected networks:

A cell group is considered to be a collection of contiguous neurons whose intrinsic connectivity is defined by events in ontogeny and development. The connections within a group are not random but are definite. Each such group may have divergent or convergent extrinsic connections to and from other such groups; these connections are also neuroanatomically defined and nonrandom.

Edelman proposes that prenatal development results in cell groups composed of hundreds to tens of thousands of neurons which are more interconnected than connected to other groups. This development takes place by means of dynamic processes of cell adhesion which control cell motion and dendrite and axon extension (Edelman, 1983). The overall resulting organization is called the primary repertoire, and is a collection of sets of groups with diverse intrinsic connectivities but similar extrinsic connectivities. Each set is a degenerate recognition system, meaning elements of the set are connected differently internally, but respond to input in a functionally similar manner.

The second phase of development occurs postnatally via a process of selection. Those subgroups that respond to input at any given time undergo an alteration of some synapses, functionally changing the intrinsic or extrinsic connectivity. The effect is an alteration of the probability that these subgroups will become active again when presented with a similar input pattern. A sufficient repetition of input selects, in this manner, various subgroups over their neighbors, and produces a new organization called a secondary repertoire. Edelman summarizes the result:

Thus a secondary repertoire is a collection of different higher-order neuronal groups whose internal or external synaptic function has been altered by selection and commitment during experience. Moreover, repetition of input need not be confined to external signals, but may include reentrant inputs from the brain itself.

Much like Hebb's cell assemblies, selection of more than one group in response to a pattern provides for associative memory and recall. Edelman advances 'phasic reentry'—cyclic repetition of a sequence of neuronal events—to account for consciousness and reaction to novelty; this concept bears a resemblance to Hebbian phase sequences. But cell assemblies and phase sequences are entirely based upon experience. The selection aspect of the group-degenerate model helps explain general competency and common innate behaviors across a species while allowing for diversity and individuality within the species.

Coward. An issue not yet addressed in the suppositions reviewed is that of the inclusion of novel information into an otherwise developed representation of environment and experience. (Coward, 1990) presents a system-level hypothesis of brain function which attempts to explain this and several other issues such as declarative and procedural memory encoding, system effects resulting from pleasure and pain, and the functionality of sleep. Coward calls this model a *pattern extraction hierarchy*.

He begins with the following definition:

A pattern is a set of components, each of which is itself a pattern. A component can indicate positively or negatively for the presence of the pattern, and the degree of indication can vary between components. A pattern repeats when the weighted sum of identified components exceeds a threshold. This threshold can be variable.

While not explicitly stated, it must be the case that there exist minimal patterns, and these are assumed to be the signals from the environment to which individual sensory neurons respond.

Like Edelman, Coward recognizes the existence of genetic mechanisms which control initial neuronal interconnection. In primitive species this control results in identical nervous systems. Similar to the selection concept of Edelman, Coward proposes that the initial connectivity of more sophisticated animals is reduced by

early experience to those connections required by the individual animal.

In the pattern extraction hierarchy, each neuron is a pattern recognizer. The pattern components are given by the inputs to the neuron and the recognition threshold is that of the neuron as modulated by local arousal level. Close to sensory input, a neuron's firing is interpretable as simple pattern recognition. Further into the hierarchy the firing of a neuron may also be viewed as an action recommendation. When it influences a muscle, a neuron's activity represents a choice of action.

A neuron's arousal level determines how much of a pattern need be present for it to fire, and its firing rate can be interpreted as the relative importance of the pattern it recognizes. Since neuron activity can result in a release of neurohormones which serve to arouse a functional region of the system, self-regulation of attention is a fundamental aspect of the model. A similar self-regulation mechanism is proposed to explain the effects of pleasure and pain upon the system. In this case, the release of hormones indicative of pleasure or pain is thought to affect the connection strengths of neurons recently active at the action-selection level of the hierarchy.

In the pattern extraction hierarchy model, alteration of connection strengths is assumed to account for *procedural* learning. *Declarative* learning within the cortex is explained by an entirely different process:

When a region of the cortex receives large amounts of input, but no output, the state can be visualized as cascades reaching partially through the region, but not to the output level. This happens when an initially strong cascade (i.e. one with much sensory input) peters out and is eventually blocked by lack of appropriate pattern nodes. This condition leads to the postulate of a population of previously unused neurons that are equipped with rich connectivity drawn from the same sources as their neighbors. In a blocking situation, the model assumes that the arousal level of a set of the previously unused neurons rises. Some of these fire for the first time and become permanently imprinted with the pattern that triggered them. This firing of neurons fills gaps in the template and allows the cascade to reach the output state. These neurons become the engram of declarative memory, distributed over many brain regions. The existence of this type of connection sensitivity imprinting is a critical physiological prediction of the model.

It is proposed that active sleep is an internal review of recent experience which results in the appropriate preprogramming of unused cortical neurons for regions whose sources have become depleted.

Grossberg. Stephen Grossberg and his colleagues have amassed an impressive set of theoretical work on neural mechanisms (see, for example, Grossberg, 1982; 1987a; 1987b), which attempt to provide *formal network architectures* to account for such behavioral issues as cognitive codes, conditioning, attention, reinforcement, memory, vision, and motor coordination. Grossberg's theories approach these mechanisms by way of the necessity of real-time world interactions.

The dynamics of interactions are important enough to elucidate further. Beginning with a finite capacity for attention, it follows that neural elements which receive external stimulation compete among themselves for dominance. The dominant network thus represents what is currently attended to in limited short-term memory (STM). The competition is based upon the relative significance of each signal as biased by the current status of internal needs and the innate or learned knowledge of the signal's consequences. So attention is dynamically directed and strongly dependent on the agent's history. But the agent's history (future) is likewise dependent upon those signals which have been (are being) attended. For those stimuli previously attended have led to the current knowledge of consequences, the process being learning and the result being long-term memory (LTM). These are the very memories serving, at present, to modulate attention.

It is this complex, dynamic, and inseparable mess of dependencies which ultimately results in the current choice of action. And it is this mess of dependencies that Grossberg has focused upon, using mathematical properties of dynamic systems to both explain and predict neurological and psychological data. The real-time approach has yielded several network principles. One is the unit of LTM, which previous theories indicate as being the synapse.

Grossberg focuses on *neurons*, called *sampling cells*. In his work, the unit of LTM is not a single synapse between two cells, but rather the *entire set* of synapses from the sampling cell to its destination cells. What is learned is a reflection of the average activities of destination cells *when the sampling cell was active*. This last emphasis is the presumption of the stimulus sampling operation, namely, unless the sampling cell is itself active (thereby the synapses are actively relaying signals), no changes in synapse efficacy (for that sampling cell) occur. When the sampling cell *is* active, the synapses' efficacies, or LTM traces, adjust according to the pattern of activities present on the post-synaptic cells. In this way the sampling cell may learn, over repeated exposure to a statistical pattern of cell activities, to reproduce that pattern (when active) on the post-synaptic cells.

Another tenet advanced is the isolation of the learning cells from arbitrary

firing. Grossberg proposes that sampling cells must receive two distinct inputs in order to fire. These inputs are (1) a cue or stimulus input, and (2) an arousal or drive input. This isolation has practical consequences besides restraining learning to instances when appropriate patterns exist to learn. Because a sampling cell needs both cue and arousal inputs to fire, and since, after conditioning, it should be able to produce the conditioned response given the cue, a second learning must occur: learning the arousal signal. This can be thought of as associating a drive with the cue (recall considerations of this sort in the discussion of cognitive maps). This key issue will be covered in more detail in chapter VI.

Along the lines of attention, Grossberg outlines a network in which drive representations and sensory cues interact. Competition must occur so that motivationally incompatible behaviors are not simultaneously elicited by the network. The competition takes place *after* drives and sensory inputs combine. This allows those behaviors most appropriate for the given environment to take precedence over extremely high drives which the environment cannot (at that time) satisfy.

There exists experimental support for many aspects of the preceding theories; Grossberg's work alone is riddled with references to neurological evidence supporting his theories. (Thomson, et al., 1983) present a review of experimental results which indicate neuroanatomical plasticity as a fundamental process of memory and learning. Other work implies that plasticity may be highly dependent on temporal aspects of synaptic inputs (Segev & Parnas, 1983) which in turn seem to be influenced by the geometrical structure of dendritic trees (Horwitz, 1983). The structures of these trees have been found to vary over the life of an animal, usually most rapidly in early development (Thomson, et al., 1983). Evidence exists indicating more highly interconnected and developed cortex regions in animals raised in an enriched environment with respect to the cortical regions of those raised in a relatively stimulus-free area (see, e.g., Rosenzweig, et al., 1972). There have been studies which demonstrate genetic control of initial neuronal interconnections (Bunge, et al., 1978; Kandel, 1979; Edelman, 1983). (Routtenberg, 1979) has provided an analysis of the apparent influence of neural pain and pleasure systems on memory and learning.

CHAPTER III

SYSTEM COMPONENTS AND VALUES

Of each particular thing, ask: "What is it in itself, in its own constitution? What is its substance and material? And what is its causal nature? And what is it doing in the world? And how long will it abide?"

Marcus Aurelius

Our goal is to use the knowledge and principles summarized in the previous chapter to guide our development of an environmentally interactive autonomous robot capable of flexible and adaptive behaviors. The basic building blocks we used to obtain such behaviors are described in this chapter. SENSI's physical components determine what the robot can react to as well as what responses are possible. This interface also defines the environmental niche for which SENSI is best suited, so a brief summary of an ideal environment is given. Finally, SENSI's innate drives and reflex behaviors, which form the foundation for all of its actions, are discussed.

Motor Mechanisms

While environmental interactions are critical to SENSI's development, it itself doesn't have a great effect upon its environment. It is capable of moving through it, making sounds and lights, and it can power up from adequately intense light sources. So, SENSI is a system which learns, not to change the environment to its liking, but to acquire a fit with its surroundings adequate for survival.

After several attempts at chassis design, the base of a 16:1 scale model King Tiger tank, measuring approximately 16 by 9 inches, was chosen for the motor element. The model was modified to replace the original single motor and clutch mechanism with dual, high torque motors, each of which drives a track. Speed and direction are independent for each motor; this provides reasonable flexibility and sharp turnability. The motors and circuitry can be operated on 16 C-cell batteries. Details of the electrical interface to the chassis are given in Appendix B. An oval

plastic plate, affixed to the tank base, forms the platform on which the remainder of the robot hardware rests.

A voltage-sensing circuit applied to the batteries provides visual indication of low and full power, by means of red and green LEDs, respectively. A speaker is used to audibly denote (via different frequencies) activation of nociceptors (pain sensors), touch, and feeding. These signals would be much more significant, ethologically speaking, among a herd of robots, where they could communicate warnings or indicate an individual's status to the rest of the group. For our purposes, they serve as a diagnostic aid — they make visible or audible those internal conditions which may affect behavior in otherwise inexplicable ways.

The robot can maintain adequate running power if it has enough exposure to bright lights. A 12 by 6 inch photocell panel provides current to the rechargeable batteries.

Sensory Systems

While we have attempted to keep SENSI sensor rich, we have been partially thwarted due to space, hardware and complexity. A sensor rich agent has more information about its environment and can therefore perform more in terms of recognition, differentiation, classification, and reaction. This is thought to be accomplished by the activation of various neurons in response to simple, distinguishable features — the more features discernible, the finer the perceptual grain. We have not been able to incorporate this level of complexity and have simplified the perceptual components in at least two significant ways. First, there is no redundancy of neurons regarding a given feature. Either the neuromime works and the feature is noted, or it doesn't, and that stimulus can no longer be a basis for response. Second, we have restricted 'features' pretty much to 'modalities'. While some directional information is provided by a receptor's location, primarily a stimulus is either present or not. Thus, recognizable 'patterns' are extremely simple in this implementation.

Compared with animals and their wide ranges of sensory capabilities, the sensing modalities implemented are also quite limited. Various animals utilize mechanical (touch, pressure, position, tension, acceleration, vibration, and sound), chemical (smell, taste), visible light (eyes, skin, pineal body), thermal, infrared, and electrical receptors (Waterman, 1989). SENSI has only mechanical (touch and motion) and visible light sensors. Also limiting, with respect to animals, is the

passive nature of SENSI's receptors. It cannot independently focus any of them; rather, it must reposition its body to examine the world from a different perspective.

For the purpose of discussion, SENSI's receptors are partitioned into two classes: proprioceptors, which sense and indicate the robot's internal status, and exteroceptors, which sense and indicate environmental conditions. The schematic of each sensory system's electrical interface can be found in appendix C. Each sensor is described in further detail in the next chapter, and physical diagrams are provided in cases of hand-manufactured devices.

Proprioceptive Sensors. There are two types of proprioceptive sensors implemented: power sensors and a motion sensor. The power sensors indicate low power (need to charge), full power (charged), and increasing power (charging). These senses are used to determine when a search for sustenance is a priority by activating the refueling drive, and to verify that refueling occurs. The motion sensor is used to verify that SENSI is moving when it wants to be (a good situation), or is stuck or restrained (a bad situation), resulting in a need for either more applied power, screaming panic, or both. The motion sensor is capable of indicating whether SENSI is moving predominantly backward or forward.

Exteroceptive Sensors. The robot is equipped with 14 whiskers, most of which are capable of signaling two degrees of touch. They are used to indicate contact and heavy contact (recognized as pain). Nociceptors signal invasion of body space, and their activation is to be avoided by the robot.

SENSI has two light sensors, situated on the front of its chassis. They are to the left and right of center, and point predominantly forward. Each sensor produces three outputs; one indicating that a large, sudden change in light intensity has occurred, one proportional to the intensity of light received, and one proportional to the received light's relative intensity with respect to the other sensor's input. The latter output is available to aid in directing motion toward brighter (feeding) areas.

Environment

SENSI's ideal environment is an enclosed area with a firm, relatively smooth surface. There are no steep drops or slopes. Because the robot is tracked, with high torque motors, it should be capable of traversing rougher areas, or rugs, but the weight of the electronics will cause high current drain on such surfaces. Obstacles

may exist, but should be (1) at least four inches high, so the whiskers contact them, (2) of smooth texture, to avoid snagging, (3) of sufficient extent that contact will be made with at least one whisker, and (4) in contact with the floor (so SENSI isn't decapitated). Obstacles need not be stationary. Sources of light should be located high enough that the robot can sense them, and be intense enough for solar-cell charging. As the robot has no sensory mechanism regarding cold or warmth, the climate should be temperate. Given the electrical nature of the device, standing water or high humidity is not a good idea.

Drives, Reflexes, and Innate Values

SENSI has drives representing boredom, fear, and hunger. It has recoil, startle, and 'frustration' reflexes, where the recoil reflex is an abrupt motion away from the area of discomfort, the startle reflex is a freeze and wait response, and frustration results (ideally) in an increased motor response. The hardwired responses to stimuli can be thought of as providing SENSI with a set of innate values; they dictate what SENSI pays attention to in its world, and what actions it can take.

The boredom drive, active unless suppressed by strong stimuli or the activity of other drives, causes SENSI to move slowly forward. Boredom will control the motor response so long as none of the other drives becomes dominantly excited.

The refueling (hunger) drive activates when the batteries reaches a predetermined level of discharge. Its activity increases in intensity as battery voltage drops. The effect of this drive becoming predominant is an innate tendency to move in the direction of brightest light, and a predisposition to stop in adequately intense light for battery charging.

The fear drive is excited by the presence of pain, and all other stimuli which excite a reflex. It does not innately cause a motor response; the reflex behaviors control the motors in such situations. Fear exhibits internal hysteresis, however, and the presence of subliminal fear can affect other behaviors.

With respect to reflexes, recoil is brought on by excitation of one or more pain sensors (built into the whiskers). Startle is associated with unexpected sensory input, such as touch on a side opposite motion, or a sudden fluctuation in perceived light. Frustration is excited by the presence of motion requests without corresponding motor activity detected by the motion sensor.

The drives and reflexes provide the foundation for everything SENSI does (and is capable of learning). It can attend to touch, pain, light, sudden variation in

light, motion, and battery voltage level. It is capable of moving forward, turning in either direction, and moving backward. Given the necessary connections and experience, it could learn to associate any of the perceivable stimuli (regularly occurring in the context of some drive) with a suitable combination of motor responses.

Given that its innate responses are motion forward, halt, light tracking, and recoil, these are the behaviors most simply associated with other stimuli (more 'sophisticated' sensory input, in the words of Pfeifer & Verschure). The drives under which these stimuli can acquire meaning are limited to the innately present fear, hunger, and boredom. SENSI's adaptation capabilities are described later, in chapter VI. The next chapter describes how the innate drives and reflexes are implemented in our neural network.

CHAPTER IV

NETWORK IMPLEMENTATION OF INNATE RESPONSES

In constructing an induction machine, we, the architects of the machine, must decide *a priori* what constitutes its 'world'; what things are to be taken as similar or equal; and what *kind* of 'laws' we wish the machine to be able to 'discover' in its 'world'. In other words, we must build into the machine a framework determining what is relevant or interesting in its world: the machine will have its 'inborn' selection principles.

Karl R. Popper

We call our neuromime network an electronic nervous system because of its control function in determining all of SENSI's actions, and because its design and implementation have been guided by knowledge of the animate nervous system. "Although we cannot hope, at present, to duplicate the intricacies of real neural networks, neurobiology can provide us with the inspiration required to translate the concepts of artificial neural networks into working silicon" (Murray et al., 1991). The neuromimes in SENSI's nervous system were developed using biological neurons as models. Concepts of network organization have been gleaned from theories of neurological function. This chapter provides a description of our general neuromime, sensor and motor circuits, reflex connections, and drive implementations.

Neuromimes

The features implemented in the general neuromime were taken from knowledge of its biological counterpart. It is a 555 timer based circuit, augmented with operational amplifiers, capacitors, resistors and diodes. A schematic and a more complete description of the circuit may be found in Appendix A.

Neuromime inputs may be excitatory or inhibitory. The amount of influence each input has on the neuromime can be easily varied by changing a resistor. In the general case (given in the appendix), there are two inhibitory inputs, each capable of

canceling an excitatory input, and one totally inhibitory input. There is a provision for self-excitation (positive feedback). Inputs are summed and compared to a threshold which itself is adjustable. The threshold varies dynamically with neuromime activity via fatigue; this negative feedback causes the threshold to rise as the neuromime remains active. Both the rate of fatigue and the rate of recover from fatigue can be varied.

The output of each neuromime is a pulse train. Pulse width and frequency (i.e., the duty cycle) of the output are dependent on the intensity of the input sum relative to the threshold. (Murray, et al., 1991) have presented pros and cons of pulse-stream neural networks. They discuss pulse amplitude, pulse width, pulse code, and pulse frequency modulation. The fact that our output pulses vary in both width and frequency is a consequence of the manner in which the 555 timer IC is being used, but it is a robust and current-efficient method of pulse-stream encoding. Note, however, that with this form of signal, pulses must overlap in time to combine as input.

The general neuromime was electrically balanced in such a way as to require two or more excitatory inputs (in the absence of inhibition) to fire. When it receives two such inputs, it responds with a train of V^+ pulses at a frequency of about 20 Hz and a 50% duty cycle. All else kept equal, fatigue sets in after approximately 1 second, and reduces the output to a 17 Hz, 8% duty cycle. The output limit at input saturation for the general neuromime (with 3 excitatory inputs and no inhibition or fatigue) is 70% at 40 Hz.

In the course of testing our neuromime, it became clear that certain, well-defined uses did not require the full generality of the initial circuit. Simpler neuromimes have the additional benefit of requiring less space — a practical concern for SENSL. The number of inputs for a given neuromime can be reduced if, for example, it is a sensory receptor, with only a single source of activation. Fatigue need not be present for all neuromimes, although this would certainly be more biologically correct. Other neuromimes do not require the inhibitory input circuitry. These specialized neuromimes are used in the sensorimotor interface. The specializations are discussed as they come up, a summary of component and response variations is given in Appendix A, and simplified sensory neuromimes are described in Appendix C.

Sensorimotor Circuitry

Motoneurons. The motor driving neurons are modeled after mammalian motoneurons in a specific manner. Studies of motoneurons (Shepherd, 1990) have revealed two primary classes, called type S (slow) and type F (fast). Type F motoneurons are subdivided into types FR and FF, meaning resistant to fatigue and fatigable, respectively. In the nervous system, these motoneurons become active on demand, in the order S, FR, FF. This occurs due to the relative thresholds of the types. The motor effect is a smooth motion beginning with the slow (low force) neurons, incorporating recruitment of the FR (high force, slow fatigue) neurons, and finally, the FF 'burst of power' if required. The motoneurons deactivate in the reverse order. One can imagine how such a scheme lends itself to smooth efficient muscle actions. Since type S is more resistant to fatigue than type F, the movements associated with type S are those which require minimal force of possibly long duration (such as stance). Type F neurons have analogous movements (consider the jumps of a cat, for example).

The robot is powered via two motors which drive the treads of a model King Tiger tank. This is hardly analogous to muscle movements of limbs. However, the principles of the type F and S neurons were felt to be applicable. In a very loose interpretation of the motoneuron scheme, three motoneuromimes per motor were implemented: one each of types S, FR, and FF. As in the motoneurons, fatigue rates effectively increase $F_S < F_{FR} < F_{FF}$ and output effects (speed) increase $E_S < E_{FR} < E_{FF}$. Instead of the thresholds increasing $T_S < T_{FR} < T_{FF}$, they are equal, and the input resistances are successively higher for the respective neuromimes.

The basic idea is to forward requests for motion to a motor's S, FR, and FF neuromimes, and to funnel all outputs from these neuromimes to the motor itself. Due to the response nature of the network, the motoneuromimes will all become active simultaneously, but fatigue will reduce their effects in the order FF, FR, S. In this manner, the intensity and duration of requests for motion modulate motor speed. Figure 1 depicts the form of the motoneuromime connections. The neuromimes and connections shown are for the left motor. The use of the interneurons is divided between reflex commands and all other motion requests. While not biologically correct (Burke, 1990), this separation of function was useful to us for testing and robustness purposes. The network for the right motor is identical.

Interneuromimes were included in the motor circuit because of the possibility of needing the following benefits (Burke, 1990):

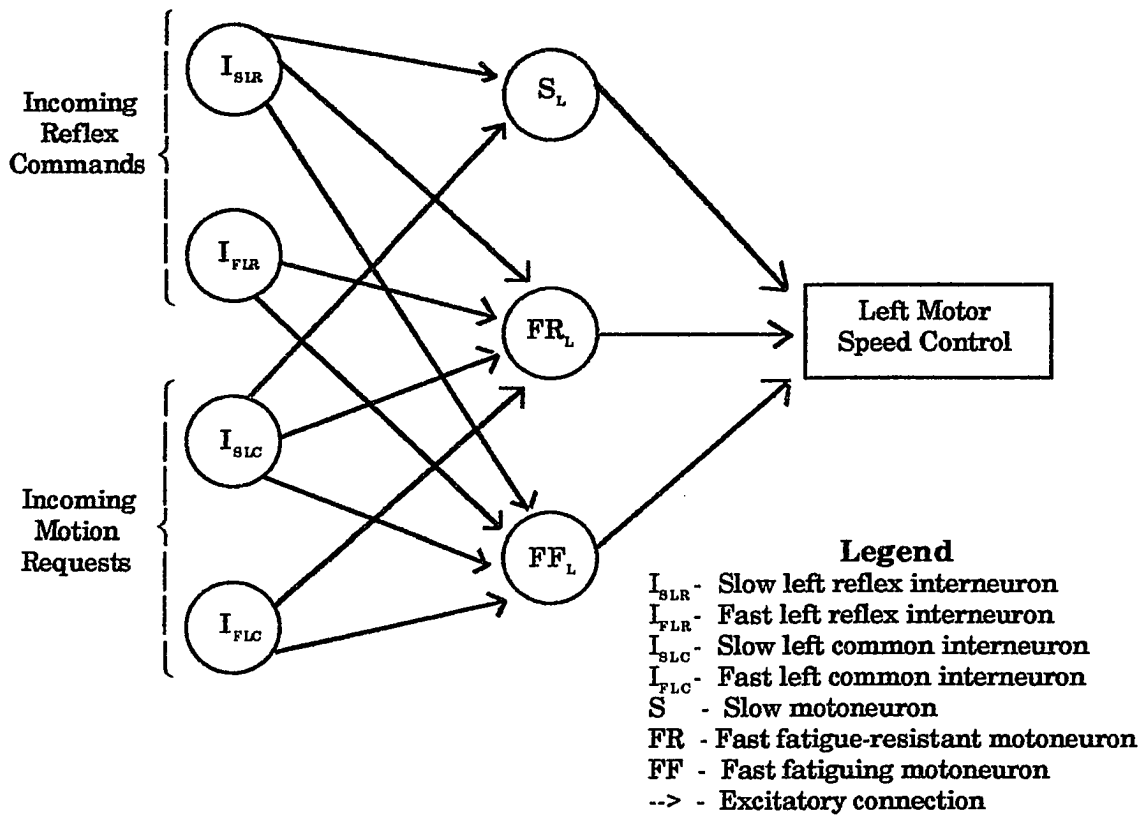


Figure 1. Motoneurone connections for the left motor.

...when interneurons are interposed between an afferent system and motoneurons (a "polysynaptic" pathway), two significant advantages accrue. First, the sign of the effect at motoneurons can be changed. So far as is known, all primary afferents excite the neurons to which they project, but an interposed interneuron can produce an inhibitory synaptic effect. Second, transmission in a polysynaptic reflex pathway can vary from zero to considerable amplification, by virtue of excitatory and inhibitory effects converging onto the interposed interneurons. Multisynaptic circuits can thus function as logical elements (in effect, digital gates) as well as analog signal amplifiers.

The connections between the interneurones and the motoneurones, as depicted in figure 1, follow a pattern of synaptic input organization as proposed in (Burke, 1990). In this scheme, one interneuron (in our case I_{Sxx}), provides input to all three types of motoneuron. The efficacy of the synaptic input decreases from S to FR to FF. A separate interneuron (I_{Fxx}) provides input to the same set of motoneurons, but only to the FR and FF types. Burke discusses how such a scheme can produce great flexibility in the use of a given motor pool.

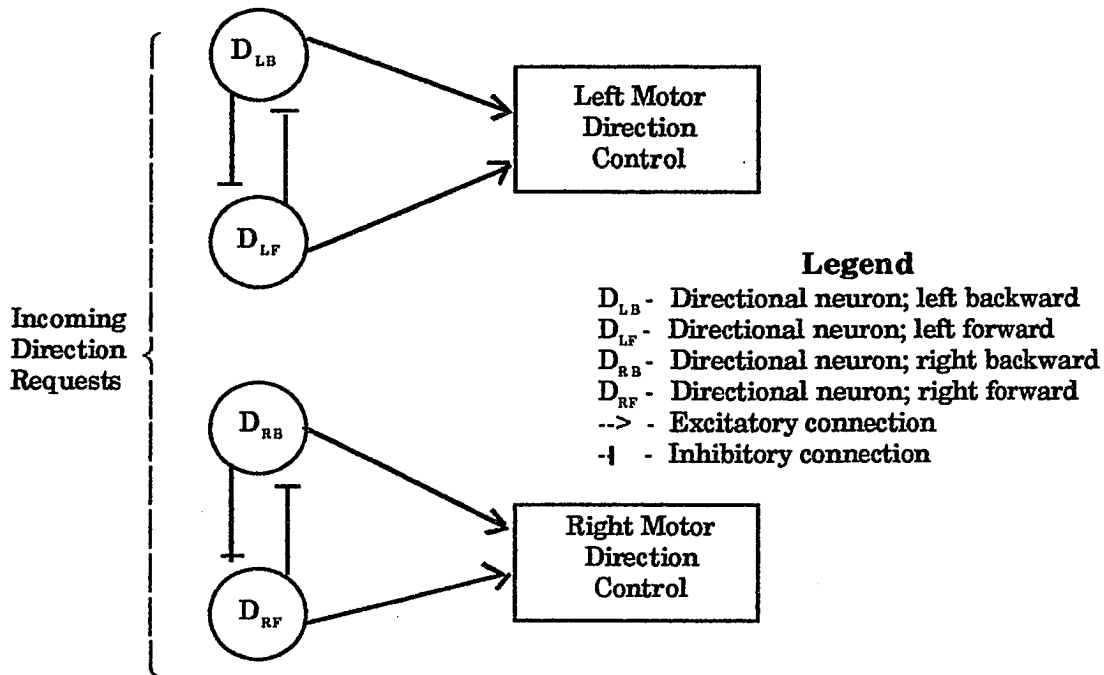


Figure 2. Directional control connections for the two motors.

In the synaptic pattern described by Burke, the FR and FF type motoneurons can be excited by the second interneuron alone, and this provides for possible rapid, ballistic movements. For our purposes, this isn't necessary, so we have used the second interneuron to enforce the recruitment order of the motoneurons. By setting the input resistances of the FR and FF motoneurons to require at least two active inputs to overcome the threshold, we have ensured that the fast motoneuromimes only fire when the slow ones are already activated.

Motor direction is determined by a pair of neuromimes per motor, one for forward and one for backward. These neuromimes are interconnected via mutual inhibition, as shown in Figure 2. They connect to a signal stretching circuit so that the activity of a directional neuromime maintains the motor's directionality until altered by the activity of the opposing directional neuromime. This deviation from ideal neuromime use was necessitated by the substitution of one neuromime for many in our network. While it is possible (and necessary in our case) to simplify the network in this manner, adjustments must be made to compensate for certain effects of a neuron group such as persistent firing, which one neuromime is not functionally capable of providing. In biological systems, persistent firing probably occurs by a cycling of activation within a group of neurons, according to fatigue, with the over-

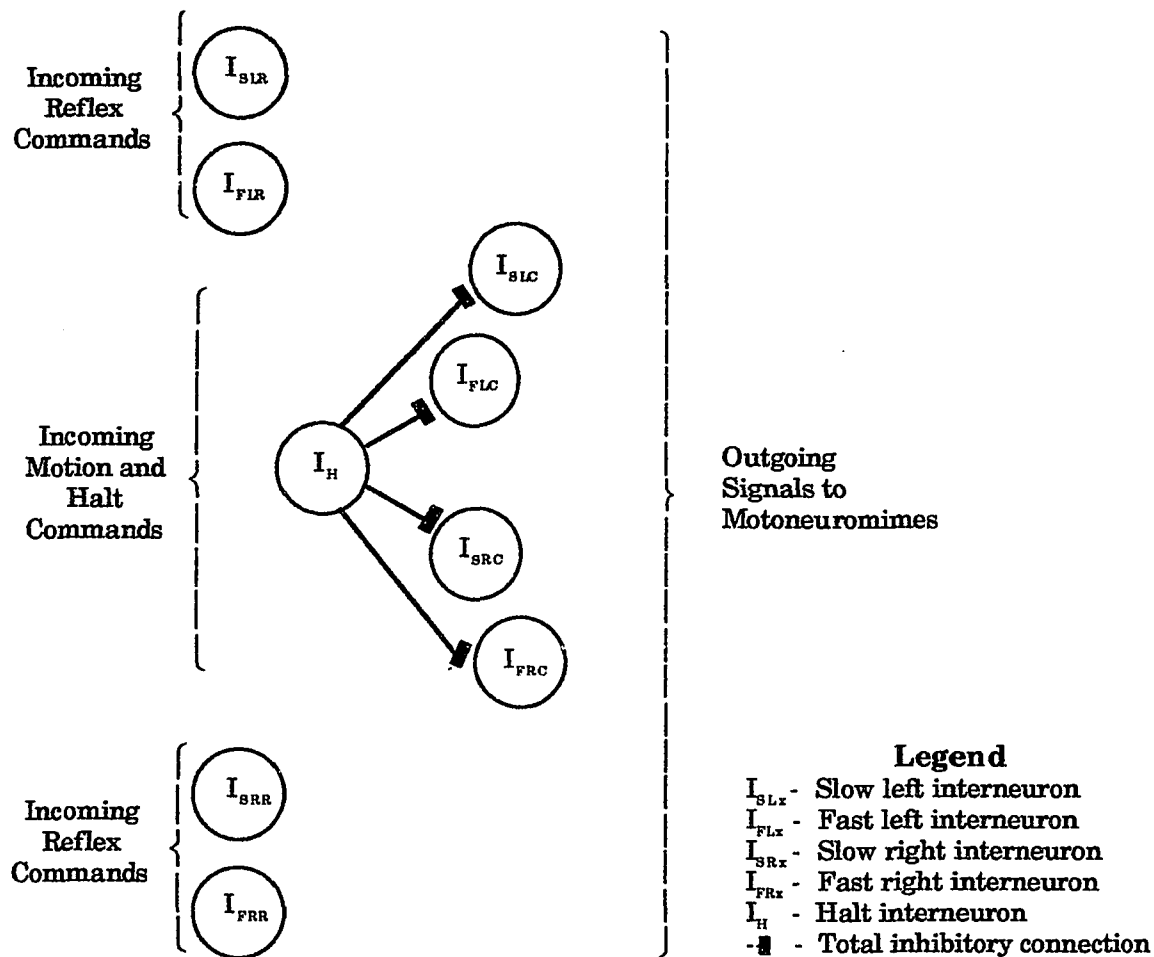


Figure 3. Influence of halt interneuron on motor interneurons.

lapping outputs of the group presenting a 'continuous' response.

An interesting consequence of the directional neuromimes' behaviors is indecision. If both forward and backward are highly excited for a given motor, neither neuromime will be able to immediately 'win' the competition by inhibition of the other, and the behavioral result is that the motor cannot proceed in either direction. Strong and conflicting motion requests can therefore lead to temporary immobility. If this condition persists, a toggling effect will occur, which is the result of each neuromime maintaining control until it fatigues.

The last of the motoneuronomimes is the halt interneuronomime which, when firing, strongly inhibits four of the eight interneuronomimes for motor control, as shown in figure 3. The neuromimes which are totally inhibited are those which receive inputs from all sources other than the pain reflex. Thus, the robot can stop if

it wants to, but will still respond to pain (the reflexes bypass the halt command). The halt neuromime is very simple, having only excitatory inputs. Only one input is required for it to fire (no threshold), and it does not fatigue in the usual sense. It exhibits hysteresis in that a single excitatory input results in an enduring (though decaying) output signal.

Figures 1 and 2 show black boxes for the interface between the neuromimes and SENSI's motor system. For a schematic of the hardware interface, and a brief description of function, please refer to Appendix B.




Figure 4. Diagram of the motion sensor; perspective view from bottom.

Motion Sensing. SENSI needed some method of ascertaining if it was moving when it intended to be. For this purpose, we developed a proprioceptive motion sense. In animals, this sense is a function of the inner ear; in SENSI it is a function of the tail. Figure 4 illustrates the mechanical design of the sensor. It is basically a wheel which rotates as SENSI moves over terrain. This rotation causes metallic strips to pass a photodetector. Each time this occurs, photosensor activity excites a neuron indicating motion. The sensor provides rough speed information by virtue of the evenly spaced strips, and provides directional information due to the separation of the plate to which the sensor is attached. Since the sensor can rotate a full 360 degrees, its contact with the base plate rotates to touch either the front or back half. Which half it touches determines which motion neuromime fires.

The signal output by the motion sensor is a single pulse per strip, or eight pulses per sensor rotation. The pulse duration is about a tenth of a second. A schematic and description of the interface circuitry is given in Appendix C.

Touch and Pain Sensors. SENSI has fourteen whiskers, eight of which can signal two degrees of contact. Because SENSI needed to be more sensitive to touch on the front, six whiskers exist where there were originally two. Of these, four are shorter, and signal pain; two are long, and signal touch. The touch-only whis-




Figure 5. The design and mechanics of the pain and touch sensors.

kers are mounted in place of the original front two dual-purpose whiskers, and each is flanked by a pair of pain-only whiskers. This arrangement results in a higher touch awareness prior to the indication of pain. The whiskers are arranged symmetrically around the oval base plate of the robot. As depicted in Figure 5, each dual-purpose whisker consists of a long piece of wire (in our case, G-string guitar wire) mounted in a base and centered within two rings. The rings differ in their diameter and their distance from the base, and provide the contacts which complete the whisker circuit when the wire is bent to touch them. Because it requires a greater distortion of the wire to touch it to the larger ring, this circuit signals pain. The circuit including the smaller ring signals touch.

(Burrows and Laurent, 1989) state that in the locust, “moving a hair evokes a burst of spikes in its only afferent neurone...” Because the majority of our hairs have dual functions, it was necessary for us to implement two afferent neuromimes per whisker, one each for pain and touch. (The expanded set of front whiskers share the four neuromimes that were implemented for the original two, dual purpose whiskers.) Both neuromime types are quite simple in that they require only one input (from the associated whisker), which is excitatory, so all other inputs need not be implemented. It was also decided not to implement fatigue for whisker neuromimes, as the interneuromimes to which they necessarily connect could be responsible for loss of sensitivity if desired. (As it turns out, sensory fatigue is quite necessary. More on this in Chapters VI and VII.)

The 'burst of spikes' referred to by Burrows and Laurent indicated that the whiskers should connect to neuromimes which exhibited hysteresis, which allows

the response of movement along or away from an object to continue briefly after contact is lost. This idea was substantiated by (Beer 1990), who found it helpful to incorporate hysteresis into the tactile sensors of his simulation. We have implemented a burst of output pulses in our pain and touch whiskers. The pain neuro-mimes output a 1 second burst of 95% duty cycle, 50 Hz pulses. The touch neuro-mimes have the lesser response of a 0.5 second burst of 67% duty cycle, 50 Hz pulses. Whisker electronics are described in Appendix C.

Light Sensors. SENSI's two light sensors are mounted on a small stand affixed to the front of the base plate. It's 'eyes' are round lens photo transistors (encased in a pair of 'Teddy Ruxbin' eyeballs to decrease sensitivity) which, in the presence of light, produce a voltage proportional to the light's intensity. The use of two such lenses permits the comparison of light intensities, allowing SENSI to direct its motion toward or away from light sources. The comparison is built into our light sensor circuitry, which outputs five distinct signals: one signal per lens proportional to light intensity, and one signal per lens which is active if the light is brighter on that lens' side than the other, and a final signal occurring if either lens experiences an abrupt change in light intensity input. A schematic of the circuitry and a diagram of the correlation of output to rough light exposure is given in Appendix C.

Power-related Sensors. Maintaining adequate power to function is perhaps SENSI's primary goal as an autonomous robot. To assist in meeting this goal, it has several ways of judging its current power status. There are three types of power signals made available to the rest of the nervous system: batteries charged, batteries discharged, and batteries charging. The first two signals control the hunger drive, and are discussed further in the section on drive implementations. The last, batteries charging signal, serves the purpose of indicating to SENSI when it is in intense enough light to feed. Schematics for the circuitry comprising the power-related senses can be found in Appendix C.

Recoil, Startle, and Frustrated Motion

Recoil, startle and frustrated motion are SENSI's only implemented reflexes. Recoil is a response to the activation of one or more pain sensors, and its effect is to move away from the source of irritation. Startle is actually a freeze response, occurring after a sudden change in perceived light, or an unexpected touch. Frustrated

motion is the internal raising of motoneuromime activity in response to a continued lack of motion even though the motoneuromimes are active.

The neuromimes and connections which constitute the recoil reflex are drawn in Figure 6, but only to illustrate the connectivity complexity for a simple reflex. The connectivity information is duplicated in Table 1, in a more readable form! Note that there is a symmetry to the pain responses, following the symmetry of the whisker positions. Recoil is a sharp movement away from the area of heavy contact, merged with a preference for forward movement. Pain contact on the rear two sensors (whiskers 6 and 7) results in maximum motion forward; on the center, side sensors (whiskers 4 and 9), the effect is forward motion with a turn away from the side of contact. Incidence of pain on the other rear whiskers (numbered 5 and 8) causes SENSI to turn away from the side of contact. Finally, pain on the front four sensors (whiskers 1, 2, 3, and 10) results in a combination of backing up and turning away.

The recoil responses were designed with forward motion in mind, thus movements attempt to re-align SENSI such that it can continue to make forward progress. In initial tests, we found that it might take several painful contacts for SENSI to finally bypass an object, and we considered altering the recoil response to eliminate this. On further thought, however, we decided that elimination of this feature would be an interesting test of SENSI learning to use touch information to avoid the pain signals.

The startle reflex is a much simpler reaction to sudden variations in sensed light, or to unexpected touch. Unexpected touch means any sensed touch on the half

Speed (Slow, Fast)	Pain whisker neuromimes									
	P ₇	P ₈	P ₉	P ₁₀	P ₁	P ₂	P ₃	P ₄	P ₅	P ₆
Left	F	F	F	S		F	F	S		F
Right	F		S	F	F		S	F	F	F
Direction (Fwd,Bwd)										
Left	F	F	F	F		B	B	B		F
Right	F		B	B	B		F	F	F	F

Note: The pain whisker positions and neuromime legend are the same as in Figure 6.

Table 1. Connections from pain neuromimes, implementing the recoil response.

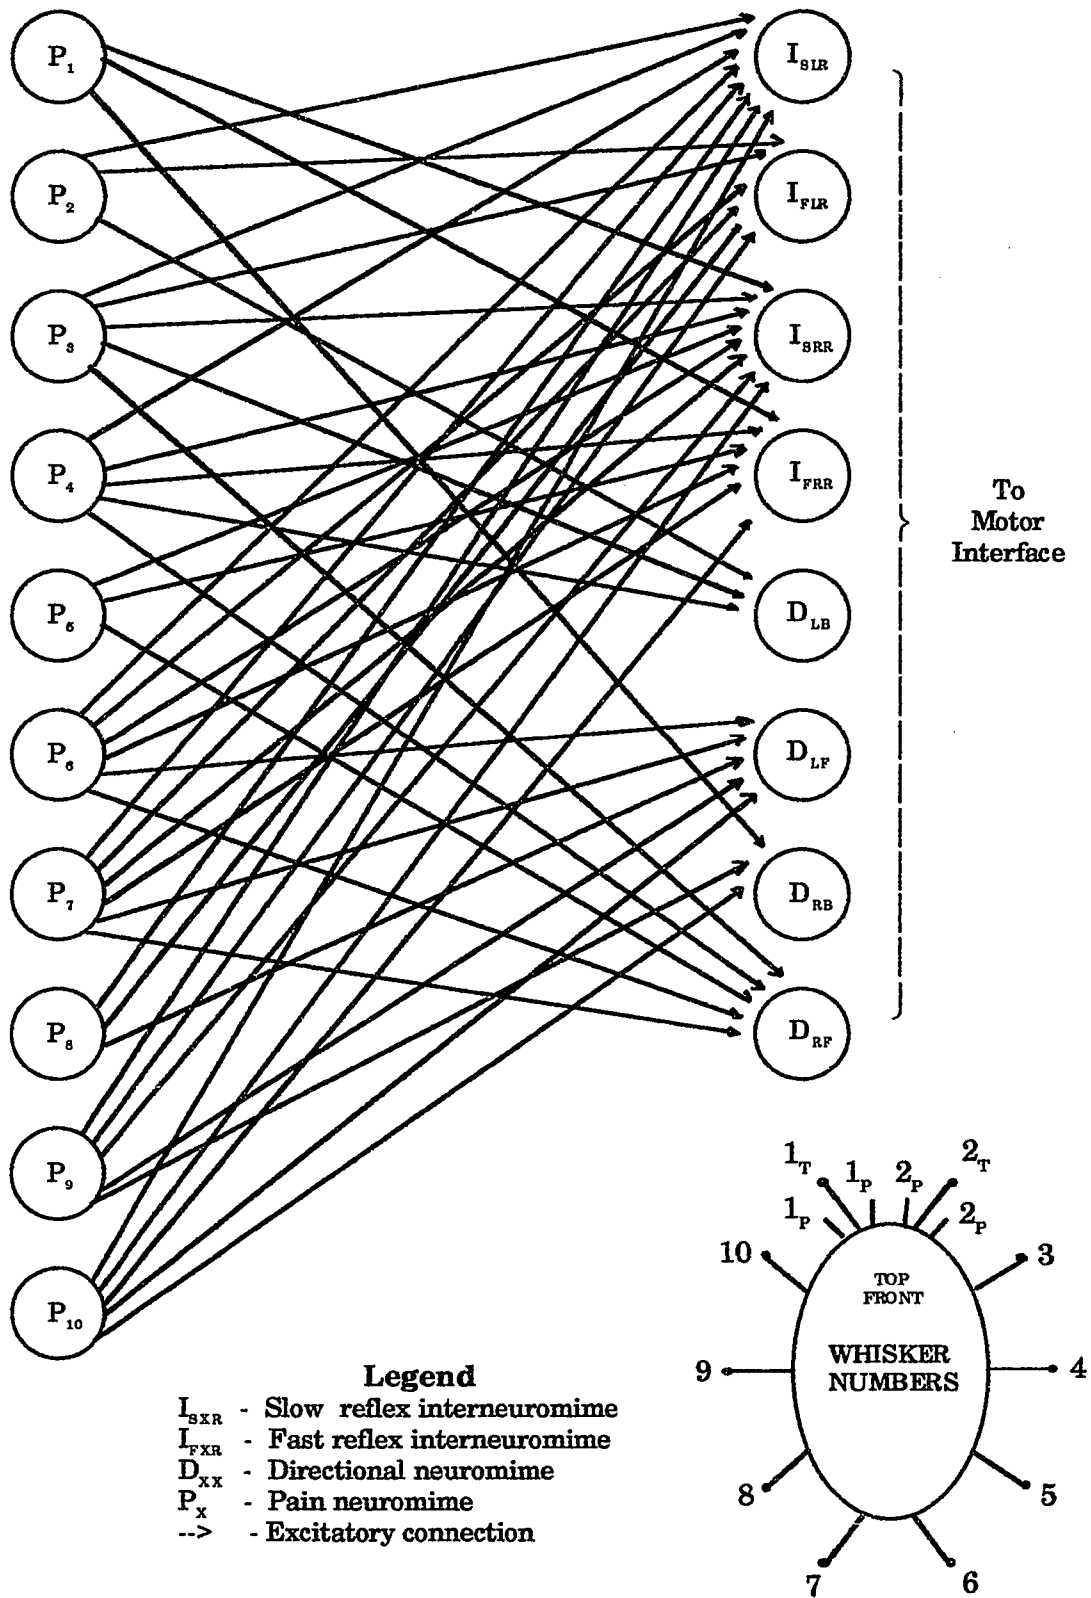
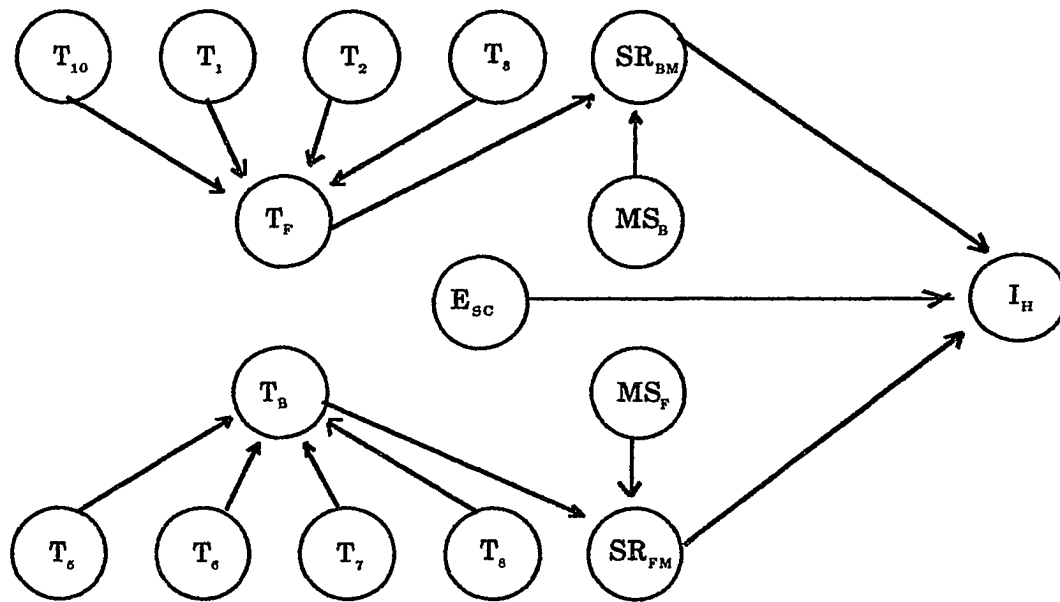


Figure 6. Neuromime connections of the recoil reflex.

of the robot opposite the perceived direction of motion. Startle merely causes SENSI to freeze temporarily. Since SENSI has a halt interneurmime, with a function exactly corresponding to the startle reflex, all that is required is to excite this halt neuromime with the onset of the appropriate sensory signals.

Figure 7 illustrates the startle reflex connections. Touch from the front portion of the robot is combined with a signal of motion backwards (in SR_{BM} , the startle reflex – reverse motion neuromime) to excite the halt interneuron. A similar set of connections causes a halt if touch to the rear of the robot combines with motion forward (the SR_{FM} neuromime). We were originally going to implement the startle reflex for touch given no motion, but this is a potentially cyclic trap (i.e., SENSI is not moving and touching, so it doesn't move and is touching, etc.) The neuromime E_{sc} directly excites the startle reflex when activated by a sudden change of light intensity received by either phototransistor.



Legend

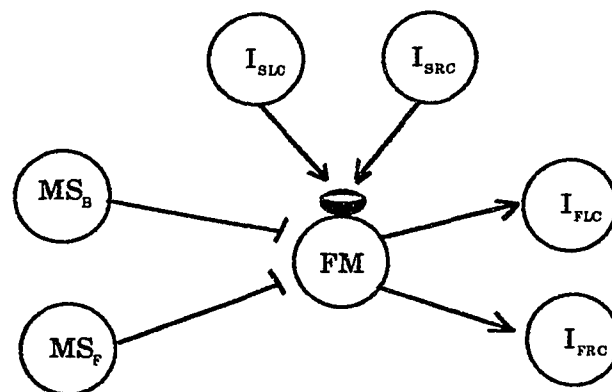
- | | |
|-------------------------------|--|
| MS_F - Forward motion sense | MS_B - Backward motion sense |
| T_f - Forward touch present | T_B - Backward touch present |
| I_H - Halt interneurmime | SR_{BM} - Startle Reflex – backward motion |
| T_x - Touch neuromime | SR_{FM} - Startle Reflex – forward motion |
| -> - Excitatory connection | E_{sc} - Eye neuromime; sudden change |

Figure 7. The network implementation of the startle reflex involving touch from the direction opposite to motion, and sudden variations in light intensity.

The last innate response to consider is the frustrated motion reaction. In the event that one or both slow motor interneuromimes have been firing, yet there is no sensation of motion, both fast interneuromimes are excited. This is a built in attempt to ensure SENSI moves when it wants to. The neuromime connections for this response are given in Figure 8.

If only one slow motoneuromime was had been excited initially (either for the left or right motor), the FM neuromime's activity will result in only that motor receiving greater excitation. This effect has nothing to do with the FM network however, it is due to the motoneuromime network discussed earlier. Recall that the fast motor interneuromime cannot fire unless the slow interneuromime is already active.

Figure 8 includes a new neuromime symbol which indicates a delay. In our neuron model, a delay is created by requiring excitatory input to build up to a suprathreshold level. We can control the rate of buildup and the rate of decay by our choice of components, thereby affecting how long an excitatory situation must exist for action to be taken. We didn't implement this at first, but with testing it became obvious that frustration should inherently build up to a response over time, not cause an instantaneous reaction, as do the other reflexes.



Legend

FM - Frustrated Motion neuromime	MS _F - Forward Motion Sense neuromime
I _{SLC} - Slow Common interneuromime	MS _R - Reverse Motion Sense neuromime
-> - Excitatory connection	I _{FxC} - Fast Common interneuromime
-- - Inhibitory connection	⬤ - Excitatory Buildup Well

Figure 8. The frustrated motion response network. The extra power provided in these cases can help SENSI if it is bogged down. Note that this response also occurs if the motion sensor fails or jams.

This innate response is a very simple reaction to what may be a more major drive in animate systems. Frustration is, however, likely much more complex, tied into expectations, for instance. In fact, it is probably a learned response. We chose to build this form of proprioceptive frustration primarily because of the help it affords SENSI in the navigation of lesser obstacles.

Drive Implementations

In the common usage of the term, SENSI has three drives: boredom, fear, and hunger. In a neural network sense, calling each of these a drive implies that there exists a locus in the network for each of them, influenced by internal and external sensory factors, and causing an effect, when excited, on behavior. Because boredom and fear are hard to identify directly with both an internal and external stimulus, it is not clear if they really should be referred to as drives. They are treated as drives here because their placement in the competitive drive network directly or indirectly influences and/or motivates behavior.

The connection format used for all three drives involves a drive *representation* neuromime, as described in (Grossberg, 1987a). In this type of network, a drive neuromime reflects the *internal* state of the agent's need, and the drive representation neuromime's activity is dependent upon both the drive neuromime's activity and relevant *external* sensory cue activities. Ideally, for example, the maximum probability of feeding behavior occurs in the presence of both low batteries and high light (hunger and food). Other combinations of drive and sensory input, for example, very bright light and a fairly high battery charge, or very low batteries and little or no light, may also excite feeding, but it is not as likely. This is true particularly if other drives are actively competing with the hunger drive.

Boredom. The boredom drive serves to motivate SENSI to move around in its world in the absence of other stimuli. The network includes boredom to capture a fraction of what, in animate systems, is referred to as exploration. In those animate systems, this exploration results in much greater benefit, such as learning the locations of resources, than in this research network. In our system, the only real advantage seems to be that it keeps SENSI moving – the only behavior it has, and therefore necessary for the learning of behavioral associations (this is an interesting supposition as to why it may have evolved in animate systems, too).

The network for the boredom drive is shown in Figure 9. The boredom neuro-

mime, B, receives tonic excitation and is active unless inhibited by significant stimuli. In SENSI's case, pain and feeding are considered to be significant. B serves as the internal measure of boredom, and its level of excitation affects BD, the boredom drive neuromime.

BDA is an arousal neuromime which is also tonically active. If BD fires, the drive representation neuromime BDR will become active based on the input overlap of BD and BDA. If the boredom competition neuromime BDC is not inhibited by the representation neuromimes (FDR, HDR) of the other drives, the resulting behavior is slow forward motion.

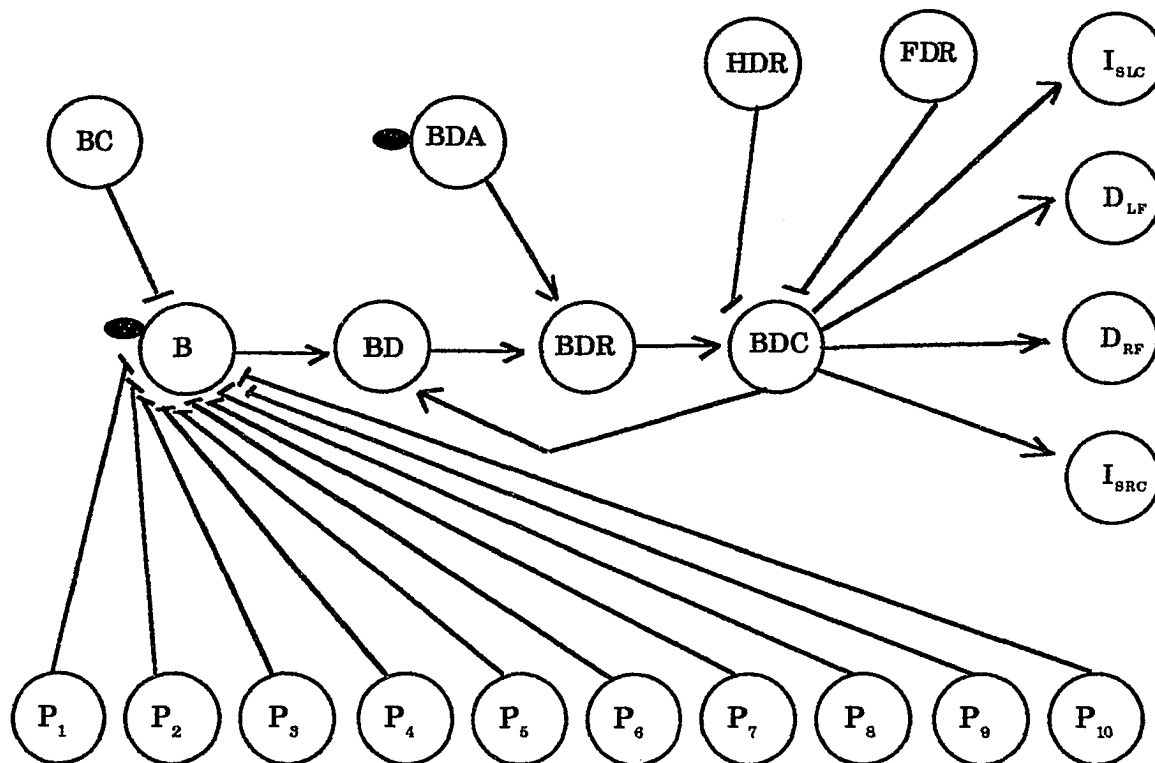
The positive feedback implemented from BDC to BD serves to maintain drive activity and drive-motivated behavior in cases of brief interruptions of input to BD. Other drives can provide sufficient inhibition at BDC to break this cycle. Also, fatigue of the drive network neuromimes limits the extent to which the feedback cycle alone can maintain drive activity.

Note again that pain and batteries charging will inhibit the boredom neuromime. This inhibition is somewhat redundant, as in either case a drive (fear or hunger) will be concurrently inhibiting the boredom competition neuromime, and thereby the effect of the boredom network. The interaction of boredom with the other drives is philosophically motivated. Sensory input of relevance reduces or eliminates boredom. The fearful or hungry SENSI is not bored. We also could not justify the inhibition of fear or hunger effects by the boredom drive, which is why the competition shown in figure 9 is one-sided.

Hunger. The hunger drive is the most straightforward in SENSI's network with respect to the model the drives are based upon. The low power sense is connected so as to excite the drive neuromime HD, and the full power sense such to inhibit it. Light (the external stimulus) and low power (the internal state) must be simultaneously present for the hunger drive representation neuromime HDR to fire.

As with the boredom drive network, activation of the representation neuromime enters the drive into competition with the other drives, and, if it wins, the competition neuromime maintains activity via positive feedback. In the case of hunger, the positive feedback allows SENSI to continue to feed (charge) even if the batteries low neuromime BL is no longer active. When the batteries high neuromime BH becomes active, meaning SENSI is 'sated', its inhibition of HD breaks the feedback cycle.

An interesting question regarding hunger was what effect it should have on



Legend

- | | |
|--|--|
| P_x - Pain neuromime | B - Tonic Boredom neuromime |
| BC - Batteries Charging neuromime | BD - Boredom Drive neuromime |
| I_{src} - Slow common interneuromime | BDA - Tonic Boredom Drive Arousal neuromime |
| D_{RF} - Directional neuromime | BDR - Boredom Drive Representation neuromime |
| - - Inhibitory connection | BDC - Boredom Drive Competition neuromime |
| --> - Excitatory connection | HDR - Hunger Drive Representation neuromime |
| ● - Tonic Excitatory connection | FDR - Fear Drive Representation neuromime |

Figure 9. The neuromime connections of the boredom drive.

SENSI. We decided to have the feeding instinct built into the robot. This is not too unreasonable, as even cats and dogs are born with the heat sensors to locate the source of food and the reflex for nursing. If the hunger drive competition neuromime HDC is active, **SENSI** will use its light sensors to move toward the most intense available source for power. When it locates an adequate light source, the AL neuro-mime becomes active, and **SENSI** will remain stationary to feed.

The neural mechanism of the hunger drive is illustrated in figure 10. If **SENSI** is hungry, but cannot sense light at all, the drive representation neuromime will not be able to fire. If light is present, **SENSI** will move toward the sensed light,

adjusting its direction (turning toward the brighter side) so as to equalize the intensity signals of its eyes. This is accomplished by the cross-connection of the relative light intensity sensors, LG_R and LG_L , to the forward turning neuromimes FT_R and FT_L , which cause turning by accelerating one track with respect to the other. Once SENSI has lined itself up on a course to the light source, it will continue forward toward it until the neuromime BC is activated. BC fires when the batteries are actually charging. Its contribution toward exciting the halt interneuromime provides the stationary feeding behavior.

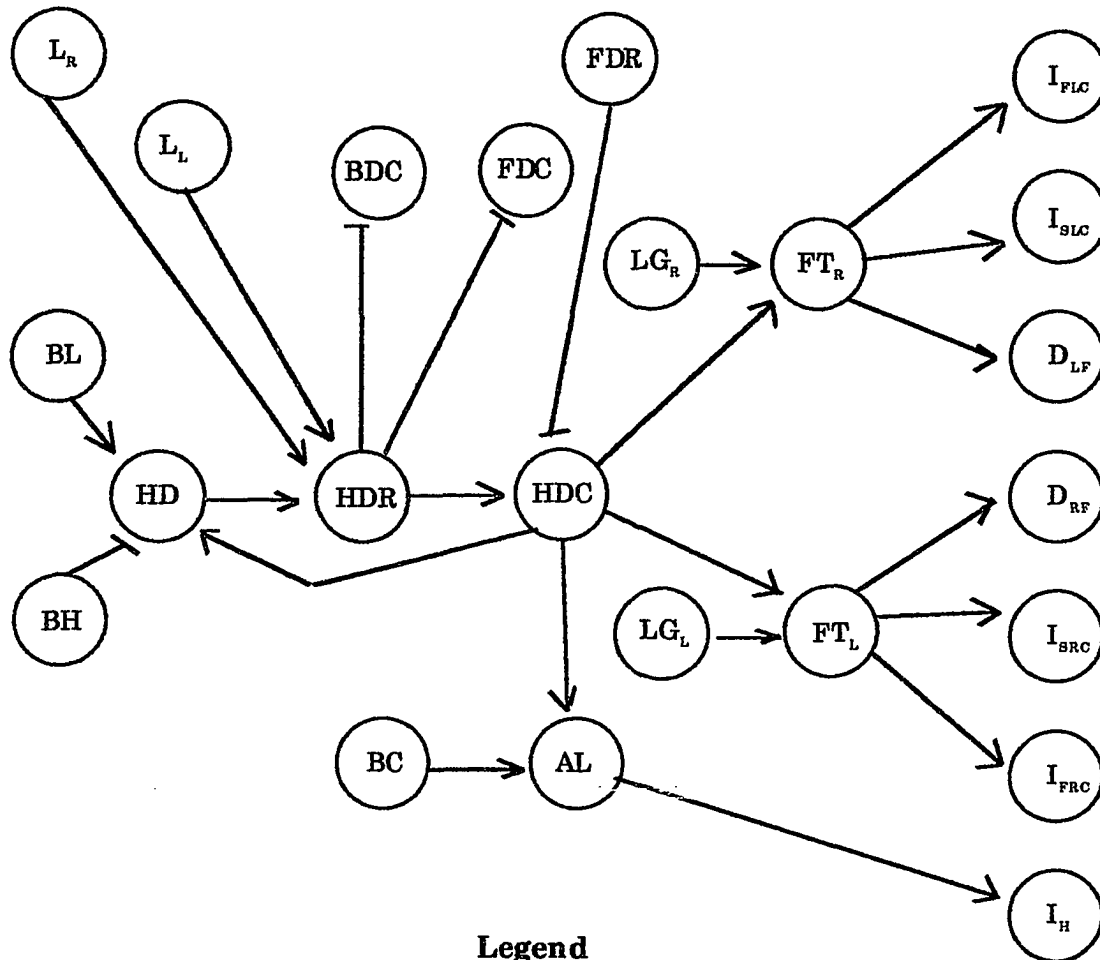
Note that there are two distinct circumstances in which SENSI will stop its feeding behavior. If it becomes sated, HD will be inhibited and HDR will not be able to fire. Without HDR, HDC cannot fire, and HDC must be active for AL to fire. So, even if adequate light exists to charge, SENSI will discontinue feeding when 'full'.

The other reason it will stop feeding is due to fear. If the fear drive is activated, it competes, via FDR, to gain control of SENSI's immediate behavior. Strong signals from the fear drive network will inhibit HDC, which, again, means AL cannot fire.

Fear. The fear drive network, shown in figure 11, differs from those of boredom and hunger in that there is no specific activity associated with the drive being active. In other words, activation of the fear network itself does not result in motor activity. The associated activity is wired as an innate response to pain and startle. Pain is an unconditioned stimulus, with the associated (reflex) unconditioned response. With respect to fear, an occurrence of pain provides both the cue and arousal required to make the drive active (this possibility is suggested by Grossberg).

It can be seen in figure 11 that even though the fear drive competition neuromime's output does not result in any overt action by SENSI, it still participates in competition with the other drive networks. Excitation of fear can dominate both hunger and boredom. The innate reflex action completed, it is possible for SENSI to remain 'paralyzed' by fear.

Also of note is the new construct introduced into the connection between the pain neuromimes and the fear drive representation neuromime. The lines coming together as they do indicate the presence of a single excitatory synapse to the representation neuromime. It is helpful to think of the outputs from the pain neuromimes as forming a dendritic tree. In essence, this type of connection limits the total amount of excitation the pain neuromimes can provide to the destination



Legend

- | | |
|--|---|
| BL - Batteries Low neuromime | HD - Hunger Drive neuromime |
| BC - Batteries Charging neuromime | HDR - Hunger Drive Representation neuromime |
| BH - Batteries High neuromime | HDC - Hunger Drive Competition neuromime |
| AL - Adequate Light neuromime | BDC - Boredom Drive Competition neuromime |
| L _r - Light Right neuromime | FDR - Fear Drive Representation neuromime |
| L _l - Light Left neuromime | FDC - Fear Drive Competition neuromime |
| D _{xx} - Directional neuromime | LG _r - Light Greater Right neuromime |
| FT _r - Forward Turn Right neuromime | LG _l - Light Greater Left neuromime |
| FT _l - Forward Turn Left neuromime | I _{FXC} - Fast common interneuromime |
| --> - Excitatory connection | I _{SXC} - Slow common interneuromime |
| - - Inhibitory connection | I _H - Halt interneuromime |

Figure 10. The neuromime connections of the hunger drive.

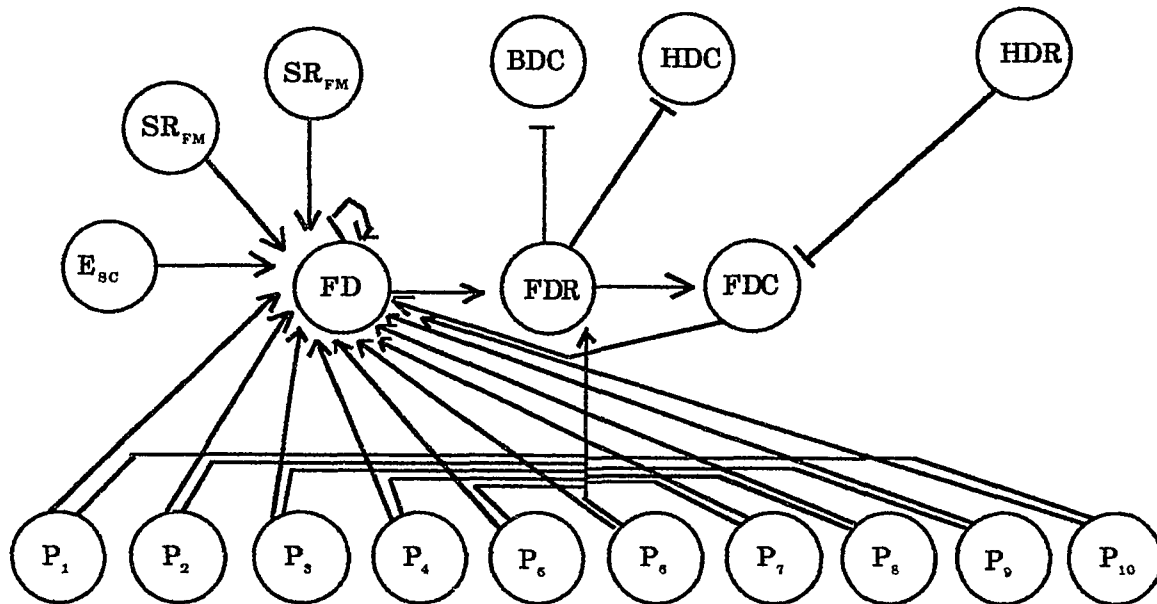
neuromime. While not as important given the special nature of fear, some form of normalization can be crucial in regulating absolute sensory effect.

Recall that the drive representation neuromime should fire only in the case of input from both the drive neuromime and from a relevant sensory cue. However,

the representation neuromime cannot determine the source of its incoming excitation. It is therefore critical that the sum of all possible sensory cue input NOT be sufficient to fire the representation neuromime in the absence of input from the drive. Yet, each sensory input must be sufficient alone to fire the representation neuromime if the drive neuromime is active. Hence the need for sensory input normalization.

It should be noted that this issue surfaced earlier – in the hunger drive network of figure 10. In that instance, with only two light inputs, we were able to balance the circuitry in such a way as to meet the drive representation criteria. With more realistic (and vast) sources of sensory input, however, the approach used in the hunger network rapidly fails.

The last novel item in figure 11 is the addition of self-excitation to the fear drive neuromime FD. This addition ensures that FD remains active for a brief



- | | |
|--|--|
| P_x - Condensed Pain neuromime | Legend |
| SR_{xM} - Startle Reflex neuromime | FDR - Fear Drive Representation neuromime |
| FD - Fear Drive neuromime | FDC - Fear Drive Competition neuromime |
| E_{sc} - Eye Sudden Change neuromime | HDR - Hunger Drive Representation neuromime |
| - - Inhibitory connection | HDC - Hunger Drive Competition neuromime |
| --> - Excitatory connection | BDR - Boredom Drive Representation neuromime |
| | BDC - Boredom Drive Competition neuromime |

Figure 11. The fear drive neuromime network. The convergence of outputs from the pain neuromimes is discussed in the text.

period after the cessation of other excitatory signals. This subliminal fear has no effect on the innate responses of SENSI's network, but will become crucial in regulating its learned responses. The use of subliminal fear in learning is discussed in chapters VI and VII.

CHAPTER V

ANALYSIS OF THE INNATE NETWORK

...analysis is much more difficult than invention in the sense in which, generally, induction takes more time to perform than deduction... A psychological consequence of this is the following: when we analyze a mechanism, we tend to overestimate its complexity.

Valentino Braitenberg

Acting together, the networks depicted in the previous chapter provide the basis for all of SENSI's innate behaviors. This chapter is concerned primarily with analyzing the functionality of the electronic nervous system as a whole. To aid in this analysis, reproductions of graphs from a multi-channel sampler/recorder are utilized. These give a visual picture of the simultaneous activities of several neuromimes (up to eight). Together with the knowledge of neuromime interconnections, these graphs allow us to visualize the activity of subnetworks under controlled input conditions.

Neuromimes

As in the previous chapter, we begin with a look at an individual neuromime's behavior. Figure 12 shows the activity of a general neuromime, given continuous, suprathreshold excitation. Note the square wave characteristic of our model. As time passes, fatigue sets in, and the neuromime exhibits the reduced, fatigued output. Here you can see how the duty cycle and frequency change as a




Figure 12. Activity of a general neuromime as it fatigues.

result of the fatigue circuitry. This characteristic, while still permitting the neuromime to remain active, results in its having less potential effect on the networks it influences.

In the diagrams to follow, neuromime characteristics will vary from that illustrated in figure 12. This demonstrates the heterogenous nature of the network – not all neuromimes are created equal. Component variations result in changes in duty cycle, period, fatigue rate, fatigue recover rate, fatigue duty cycle and period, and so forth. Also, as is seen in figure 12, some anomalies occur in neuromime output. Transitions from normal to fatigued output are not always immediate and precise, but may occur with a few 'spurious' pulses. We have not found that these unexpected pulses give rise to any serious behavioral results.

In practice, the neuromime design deserves improvement. It requires numerous components including several capacitors, which would be difficult to implement in integrated circuit format. Such a format is desirable due to the reduced space and weight. One improvement considered was the use of a bipolar power supply, which also would permit the design of a field-effect transistor-based neuromime. This change was of too high a magnitude to pursue at this time.

The behavior of the neuromimes is also lacking in some respects. The biggest problem to date has been the implementation of competition, or mutual inhibition, given this model. As discussed earlier, the neuromime design's output of pulses, combined with the sparseness of neuromimes, results in opportunistic firing. If a neuromime inhibits another, the latter is inhibited only at the times when the former neuromime's output is high (during a pulse). This is a very poor model of the chemical depletion associated with biological neuromime inhibition. The presence of more inhibiting neuromimes, with overlapping outputs, would also help achieve consistent inhibition, but redundancy of neuromime function is impractical in our limited network.

In fact, mutual inhibition was not attainable given our initial model. A compromise made in cases of mutual inhibition (in fact, for all inhibition) to meet this drawback was to use an additional capacitor to stretch the inhibition inputs of neuromimes. The shape of the inhibition input, with this addition, is a superior model of chemical replenishment. It allows the effects of inhibition to last long enough to 'cover' the interlude between normal pulses. Naturally, when the inhibiting neuromime fatigues, the inhibited neuromime can still opportunistically fire, which gives it a chance to attain control of related behaviors.

The addition of a capacitor was also utilized to enhance positive feedback in

the drive networks. Since these neuromimes fatigue, and overlap of excitatory pulses is required to fire a drive neuromime, it was found necessary to stretch the feedback input. This increases the likelihood the inputs will coincide adequately for maintaining the drive activity.

The lesson really learned here is that the shape of biological neuron outputs is important! Had we included this characteristic in our neuromime (output stretch and decay as provided by an additional capacitor), it would not have been necessary to alter the forms of individual inputs for inhibition or feedback.

Sensorimotor Circuitry

Motoneurons. Figures 13 and 14 illustrate the activities of the S, FR, and FF motoneuromimes when the motor system has been excited. The variations in neuromime characteristics are evident. Note the S neuromime's stable firing rate throughout the sampling period (the S neuromimes fatigue quickly; shown is the steady output after that point). Figure 13 shows the activities as activation begins. The starting pulses are somewhat longer than regular pulses, but the full fatigued state is not evident in the FR and FF neuromimes until later in the activation. These fatigued states, shown in figure 14, again have a lesser influence on the motor interface. The FR neuromimes, beginning to fatigue in figure 14, reach full




Figure 13. Initial responses of the motoneuromimes upon activation.




Figure 14. Motoneurones after fatigue has set in for FF and FR begins to fatigue.




Figure 15. The effect of activation of the halt interneurone in the common interneurones, and the resulting deactivation of the motoneurones.

fatigued firing in about 4.2 seconds; FF neuromimes fatigue in about 1 second. Behaviorally, the initial excitation of the motor neurons is evident as a burst of motion, which then subsides to a steady crawl with fatigue.

Figure 15 shows the effect of the halt interneuromime on the motor network. As can be seen, once the halt interneuromime is activated, all motoneuromime activity due to the common interneuromimes ceases. Previous connection diagrams show the halt interneuromime totally inhibiting the common interneuromimes (which drive the motoneuromimes). Of course, any activity via the *reflex* interneuromimes is not affected by the active halt state.

Of note in figure 15 is the immediacy of response in the end of a chain of connections. The inhibition of the motoneurons follows the excitation of the halt interneuromime very closely. Delays are practically non-existent in the time scale of this network's responses.




Figure 16. The motion sensor's output.

Motion Sensing. The output pulse train from the motion sensor is given in Figure 16. As claimed, the output pulse is about 1/10 of a second, or 100 ms.




Figure 17. Touch and pain signals from a whisker distorted through its range.

Touch and Pain Sensors. Figure 17 shows the outputs from the touch and pain sensors of a single whisker, as it is distorted from the touch range into the pain range. The ongoing touch signal even after the commencement of the pain signal is due to the whisker design (see Figure 5).

Light Sensors. The light sensor circuitry has varying responses that serve to sense light intensity, abrupt changes in light, and variation in light across the two photocells. The next two diagrams illustrate these features. Figure 18 shows the responses of all five sense neuromimes as a bright light is quickly passed across the two sensors. The figure shows the excitation of the 'sudden change' neuromime both at light onset and offset.

This experiment illustrated in figure 18 demonstrated a need to make the light sensors more responsive to gradations in light intensity. The result of this alteration is portrayed in figure 19. In this second experiment, a light was brought




Figure 18. The responses of the light, greater, and sudden change in light senses associated with the two photocell 'eyes'. A bright light was quickly passed by both lenses in this experiment.




Figure 19. The response of a light sensing neuromime as a light is quickly turned from dim to bright.

from dim to bright in front of the sensor. The figure illustrates the full range of output for sensed light. The 'light greater' neuromimes (not shown in Figure 19) were also tracked in this case. We learned that if the light is at maximum intensity, both of these side-oriented sensors will fire. As this sense is used to direct the robot forward toward bright light when it needs to charge, this anomaly is not a cause for immediate concern.

Mutual Inhibition - Directional Neuromimes

The issue of mutual inhibition, or competition, has been mentioned in the context of the directional neuromimes. Although we worked without properly operating mutual inhibition for quite some time, the lack of it eventually became critical, and prompted the following analysis. Given that we wanted only one direction, forward or backward, active at any given time, how could we implement this and still guarantee immediate reflex responses?

Our naïve starting condition was to set up reflex directional inputs to be capable of overcoming any amount of inhibition. The idea was that *whatever* SENSI was doing, a recoil reflex would always be able to re-direct it. Since a fatigued neuromime cannot necessarily fire, we removed fatigue from the directional neuromime as well.

The result? Under inevitable cases of conflict (pain on more than one whisker can result in opposing directional signals), both directional neuromimes for a given motor would receive fire-despite-anything input signals. Both would fire maximally – and the motor would lock in the immobile state. Because neither could fatigue, if SENSI didn't move for some other reason (which generally didn't happen due to the over-riding control of the reflex network) it would remain frozen in place, screaming, until rescued or until its batteries became depleted.

It is obvious now that mutual inhibition's purpose is competition between/

among neuromimes. In such a context, fire-despite-anything inputs have no place. Fire-despite-anything may as well be a hardwired switch that turns a neuromime on under certain conditions, *regardless* of its history of activation, or of competing signals. This is *not* competition, but a fixed behavior. Therefore, the first conclusion is: competition requires that a neuromime or group of neuromimes be able to inhibit the opposing neuromime(s) under comparable excitation conditions.

We have previously seen that the square shape of neuromime pulses does not lead to good inhibition properties for small numbers of neuromimes (see the section on neuromimes in this chapter). In the directional case, there are just two neuromimes competing with each other. Thus, to meet the first condition of competition, it was necessary to alter the shape of the pulses. This could have been done either at the neuromime's output or at the inhibition input. We chose the latter due to its lack of effect on the rest of the network (changing the shape of an output pulse affects all neuromimes to which the altered one connects).

To accomplish satisfactory inhibition, it is necessary for the effects of an inhibitory pulse to span the gap between pulses. The concluded constraint on the shape is: during normal (unfatigued) firing of the inhibiting neuromime, the effect of inhibition on the inhibited neuromime should be strong enough at all times to mitigate (comparable) incoming excitation to the extent that the net excitation does not exceed the threshold of the inhibited neuromime (it can't fire).

Comparable, in the above constraint, implies equivalent excitation is received by both the inhibiting and the inhibited neuromimes. The issue of whether excitation levels are equivalent further defines the expected results of mutual inhibition. Suppose both competing neuromimes receive equivalent excitatory input, and one somehow acquires the edge, namely, it fires and inhibits the other. What then? In animate systems, networks get tired, and other, less tired networks can gain the advantage. This leads to the third conclusion: given persistent, comparable, excitatory inputs, the activity of the competing neuromimes should toggle, or alternate.

To attain this behavior, we reach the fourth conclusion: fatigue is a critical component of competition. Without fatigue of the active neuromime(s) in a mutual inhibition setting, toggling will never occur. The inclusion of fatigue is no new insight (see discussion of Hebb's work, for example), but its criticality was emphasized to us in this analysis.

Here is the scenario of the competing directional neuromimes with fatigue, inhibition stretching, and persistent, comparable inputs. Initially, both may fire, but one quickly gains the advantage and inhibits the other. As the 'winner' contin-

ues to fire, it fatigues, effectively raising its threshold. It begins to fire less often and shortly, the other neuromime gets an opportunity to become active. When it does so, it is not fatigued, and its full firing strength inhibits the previously active one. This result provides the desired toggling behavior at the neuromime level.

A sample of actively competing directional neuromimes' outputs is displayed in figures 20 through 22. The figures were obtained by sampling the directional neuromime activity while both a front and a back whisker were bent down to elicit




Figure 20. All directional neuromimes are receiving comparable excitation. The dominant forward neuromimes are beginning to fatigue.

pain signals (in particular, whiskers 1 and 7). In figure 20, the forward neuromimes are beginning to fatigue. Figure 21 shows the backward neuromimes, which are not fatigued, opportunistically beginning to fire as the forward neuromimes produce less inhibition. Figure 22 shows the inhibition of the forward neuromimes by the now dominant backward neuromimes. These three figures capture the real-time toggling behavior of equally excited, competing neuromimes.

At the behavioral level, the result is realistic and potentially life sustaining! It allows the animate being to switch behaviors under equivalent excitations, and thus permits varying attention, and possibly the generation of random actions. For SENSI, it is the difference between being stuck and unable to act at all, and being

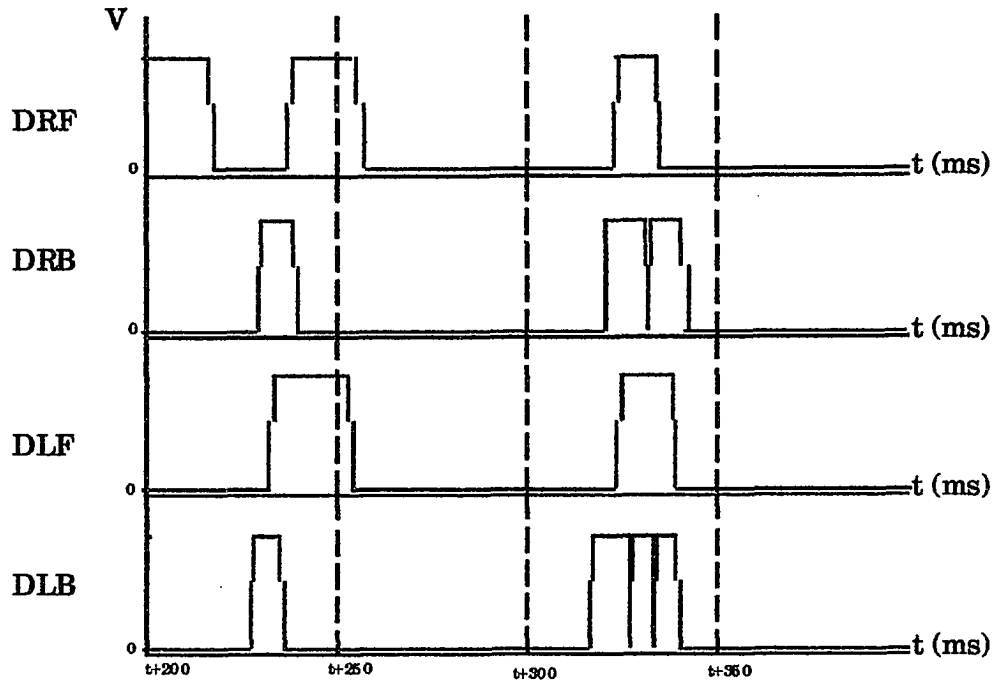


Figure 21. This real-time continuation of figure 20 shows the fatigued forward neuromimes and inhibited, but unfatigued backward neuromimes competing for dominance.




Figure 22. The continuation of figures 20 and 21 illustrates the backward neuro-mimes gaining complete dominance over the forward neuromimes. The competitive toggling is completed.

able to alternate attempts to move when the current activity doesn't extricate it from whatever tight spot it is in.

Also interesting, behaviorally, is the effect of the rate of fatigue and the rate of recovery on the competition. (These rates are individually determined by the values of resistors and capacitors; refer to Appendix A.) If the rate of recovery is substantially faster than the rate of fatigue, the toggling occurs at practically uniform time spacing. But suppose the neuromimes fatigue faster than they recover. In this case, recovery is not necessarily complete for the neuromime that wins the dynamic competition. Because it is not fully recovered, it fatigues more quickly. The other neuromime(s) takes over, but it has had even less time to recover. The behavior result of this increase in toggle rate is reminiscent of frustration in biological systems. SENSI tries hard, alternately, to go in each direction. Failing at both, it tries each more quickly. If it continues to fail, it begins to resemble a quivering wreck!

In floor tests, this condition doesn't usually occur, however, and this is due to the individuality of the directional neuromimes. Each of the four directional neuromimes is unique. This means that competition is not necessarily even. Thus if one pain reflex of driving both motors backward competes with another pain reflex of driving both motors forward, the resulting behavior is *not* necessarily that of SENSI moving forward and backward! Each motor's directional neuromimes are independently vying for control, and the resulting winners may have SENSI moving in *any* of its possible directions at any given time. From our perspective, this is terrific, because it is just this type of unpredictable, apparently random motion that might get SENSI out of the spot that is causing the competition in the first place!

Finally, this discussion began with the concern that reflex inputs to the directional neuromimes should be able to override any currently active directional inputs. Within our new framework this is still the case. Recall that fatigue and inhibition both have the effect of raising the threshold of a neuromime. A reflex input is strong, and has the potential to exceed this higher threshold. Should both directions be excited by reflexes, the comparable input scenario remains valid.

Recoil and Startle Reflexes

The recoil and startle reflexes, as defined in the previous chapter, behave pretty much as anticipated. An encounter of note is just one of assumptions. The behaviors chosen for the pain reflexes were chosen independently of each other,

initially, and based only upon the desired 'move-away' response. It didn't require much testing to ascertain that conflicting directions on two adjacent pain responses resulted in an embarrassing cessation of all motion with continual screams from SENSI. A re-analysis and re-wiring of whisker responses, taking into account the likelihood of simultaneous contact by adjacent whiskers, helped to resolve this problem.

In the end, this problem was really due to the non-functional directional competition, and the resulting motor lock-ups. While it is unknown how the original pain responses would now operate, the current setup was still capable of locking SENSI in position if three whiskers concurrently registered pain. With the directional competition in place, however, this is no longer a concern. Should opposing directions be indicated for the same motor, the worst case scenario is rapid switching of motor direction.

Since recoil results in straightforward activations of the motoneuromimes (already shown), no specific recoil diagrams are shown. Behaviorally, sensing of pain results in a fairly immediate response by SENSI away from the contact. Conflicting inputs as to the desired motion direction are competitively handled.

An instance of the neuromime response associated with the startle reflex is shown in Figure 23. In this case, SENSI was moving forward and was touched on the rear. The startle reflex neuromime fires and refreshes the halt signal with each excitation. This diagram illustrates the necessity for the motion sensor signal and the rear touch signal to overlap in order for the startle reflex neuromime to fire.

The mechanical nature of the motion sensor does cause SENSI to exhibit the startle response innappropriately. In tight corners, SENSI will recoil from an obstacle only to make contact on the opposite side. If there has not been enough motion for the sensor to switch indicated direction, startle will occur upon the second touch. The behavioral result is a slow, almost 'considered' motion from SENSI until the motion sensor registers correctly, the touch goes away, or pain and recoil recur.

Drive Networks and Interactions

The drive networks are the most complex networks covered so far for several reasons. First, they contain feedback, which affects their activity after cessation of initial excitatory input. Second, they can interact (compete) with each other, much like the directional neuromimes discussed in the previous section. Third, their activity depends on environmental stimuli, internal stimuli, and the current fatigue




Figure 23. Excitation of the startle reflex neuromime by means of a touch on the robot opposite the primary direction of motion.

and status of competing neuromimes. The diagrams presented here, showing drive network activity, are the results of carefully controlled experiments. In most cases, a single stimulus has been presented or removed, and the results documented on the drive of interest.

Figures 24 and 25 show the activities of the boredom network, a pain neuromime P_x , and the fear network. In this experiment, a whisker was bent to elicit pain, with SENSI initially in the bored state. See figures 9 and 11 for the connections among the sampled neuromimes.

Figure 24 illustrates the introduction of pain. Recall that B is active when there is a lack of pain stimuli, and that it serves to excite the boredom network. When pain commences, B is inhibited. The boredom competition neuromime (BDC), active at the onset of the experiment, is quickly inhibited by the activity of FDR (see figure 9). Without excitatory feedback from BDC, the activity of BD lessens, and halts with the cessation of B's activity. The boredom drive will not become active again until sensory inputs which inhibit B cease. In this manner, the fear drive has become the dominant drive (over boredom). For SENSI, it means that any motor activities associated with the fear drive are now in complete control of its behavior.




Figure 24. Interruption of the excitatory input to the boredom drive, showing the inhibition of the boredom drive neuromimes and the response of the fear drive network to pain.

This fact becomes more important when we examine the learning networks circuitry and behavior.

In figure 25, the cessation of the pain signal is seen; one brief contact results in a pain signal of a little over a second. FDR and FDC quit firing immediately,

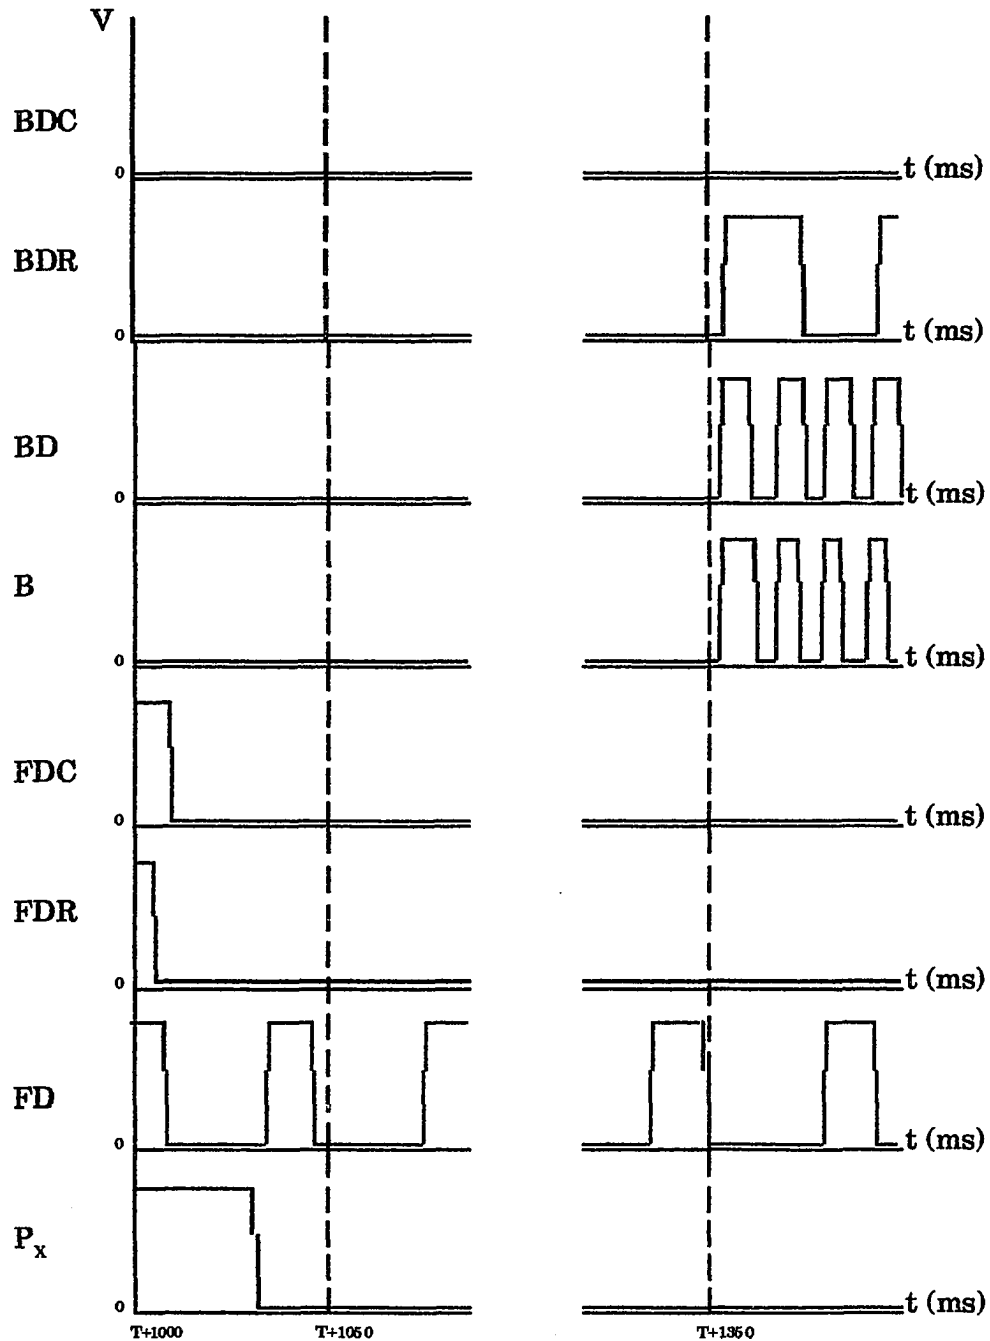


Figure 25. The end of the pain signal and the re-activation of the boredom drive. The subliminal fear (FD activity) continues for several seconds. Times shown are actual times from the start of the experiment, shown in figure 24.

because pain is required for FDR to fire, and FDR must fire for FDC to fire. FD continues to fire due to its self-exciting feedback (figure 11); this subliminal fear will continue for almost thirty seconds. We can see that B recovers from inhibition in

approximately 300 milliseconds. Although BDC's recovery is not shown, it begins to fire about 120 milliseconds after the onset of B.

The behavior of SENSI associated with this experiment is as follows. SENSI is initially bored, and running forward at a slow rate. It experiences pain and immediately responds with the reflex action of moving away from the area of pain. After this reflex has run its course, SENSI stops briefly. It then resumes its bored, forward crawl.

Figures 26 through 28 illustrate hunger drive behaviors and competition between the hunger and boredom drives (see figure 10). In this next experiment, we began with SENSI bored, hungry, and in the dark. There is fatigued activity in the




Figure 26. Fatigued hunger (but no light) while the boredom drive is active. Without light, the hunger drive network cannot become active in the drive competition.

first stage of the hunger drive network (the HD neuromime), as shown in figure 26. The rest of the hunger drive network remains quiescent due to a lack of light (both hunger and light must be present to elicit hunger drive responses). The boredom drive network remains active, and SENSI maintains its bored crawl.

We then added light to the environment. Figure 27 shows the same neuromimes as figure 26. The effect of the light being present can be seen in the activities of the HDR and HDC neuromimes. When HDR becomes active, it inhibits the boredom drive neuromime BDC, as can be seen. Note, however, that BD is still firing. This means that the rest of the boredom drive network is still active; thus if




Figure 27. The stimulus of light is added to the experiment that began with figure 26. The hunger drive becomes dominant over the boredom drive. There is insufficient light for feeding, and the boredom network remains partially active.

the light faded or the robot moved past it, the boredom drive would regain control fairly immediately. Behaviorally, with the hunger drive dominant (HDC is firing), SENSI moves toward bright light. The boredom response of slow forward motion is inhibited with the inhibition of BDC.

The behaviors associated with figures 26 and 27 indicate that SENSI is not particularly elegant in its innate response to hunger. If it cannot sense light at all, it continues to crawl along as if it wasn't hungry. If it can sense light, it moves toward what it perceives to be the brightest light. Often the brightest light it sees is a reflection of a light source off a wall or some other body. It will move toward such a bright obstacle and usually collide with it. Then, unless its reflex motion orients it so that it perceives a brighter light source elsewhere (or it doesn't perceive it at all), it will proceed to move toward the obstacle again. In cases where the reflected light is enough to charge SENSI's batteries, it will succeed in feeding for a while. But the above scenario will also occur if it becomes charged enough that it requires more light to charge further and it is still in a state of hunger. (SENSI has adequate light to charge if the voltage on the photocells is some fixed amount greater than the voltage on the batteries. As the batteries charge, the light present may not generate adequate voltage on the cells, causing SENSI to seek a brighter light source.)

In figure 28, SENSI begins to charge its batteries (we increased the light intensity to provide adequate light). Feeding excites the BC neuromime, which in turn inhibits B, thereby BD quits firing. It can be seen that this result is not immediate; the inhibition takes a while to build up. BC also, in combination with excitation from HDC, fires the halt interneuromime, and SENSI's motors stop. As long as SENSI is hungry and its batteries are charging (and nothing bumps into it painfully) it will remain stationary.

The batteries charging neuromime, BC, may fire any time that SENSI is charging, and SENSI can charge at any time regardless of hunger. However, such charging is 'on-the-run'. SENSI only *stops* to feed when the above conditions are met.

We have biased the network such that both fear and hunger are much stronger drives than boredom and therefore can easily dominate it. As mentioned earlier, the boredom network doesn't actually enter into competition with hunger or fear, because the inhibition is not mutual. The interaction between hunger and fear is a competition. There are several ways this competition could have been arranged. If we had always wanted fear to be dominant, for example, we could have biased the network so that fear was much stronger than hunger, and always inhibited it. We

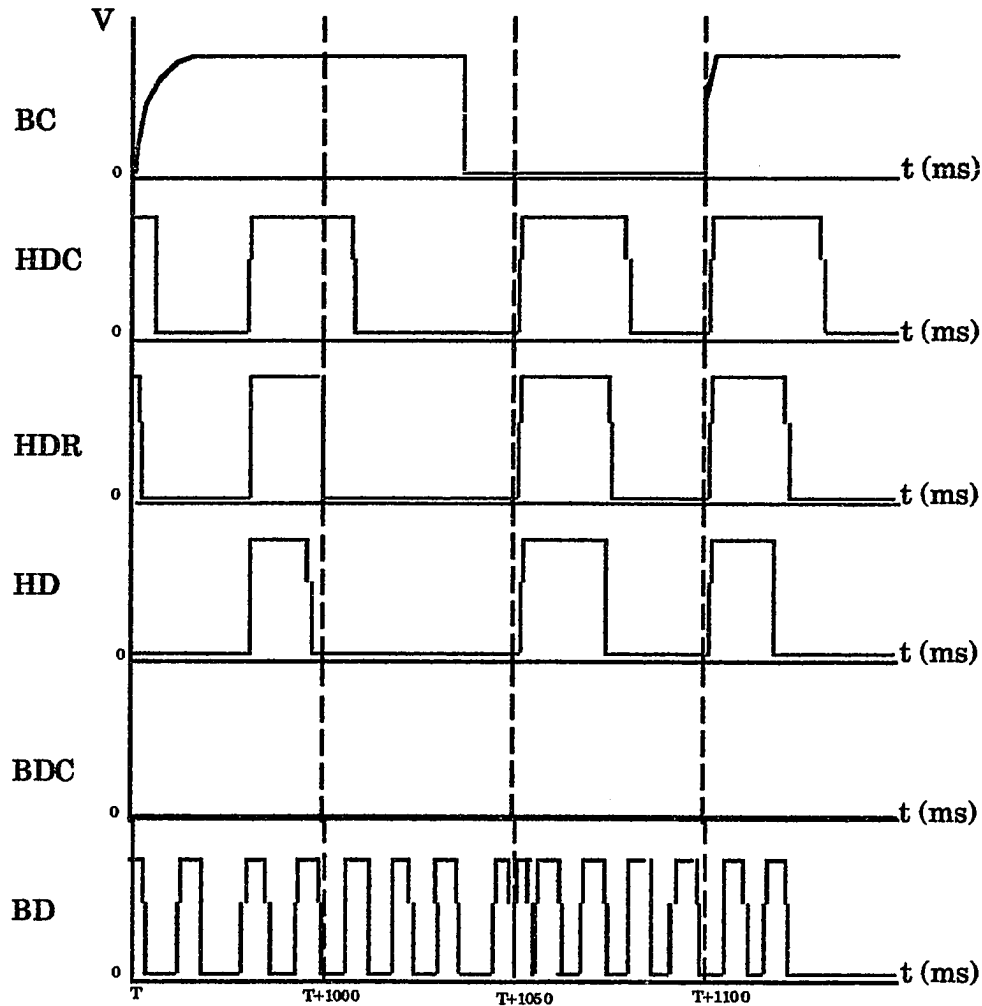


Figure 28. SENSI begins to charge its batteries, inhibiting the boredom drive.

could even have removed the hunger network's capability to inhibit the fear network. This route would have defined a strict behavior ordering: fear if present, else hunger if present, else bored. No neural network is required to achieve this form of behavior!

Mutual Inhibition Revisited – Drive Networks

What we want is competition like that between the directional neuromimes. But the drive networks make competition much more complex than the one-on-one directional neuromime competition. As in the directional case, fatigue plays a critical role; also, because networks are competing, not just neuromimes, it turns out

that feedback (within each network) is equally critical.

Figure 29 reminds us of the network in question; shown are the hunger and fear networks, including their internal feedback and mutual inhibition. Competition between the hunger and fear networks operates as follows. Suppose the hunger drive network is active, i.e. HD, HDR, and HDC are firing. Now suppose fear becomes active (FD and FDR fire). FDC is being inhibited by HDR, and cannot fire immediately. However, since the hunger drive network has been active, it has fatigued. Thus FDR can excite FDC more than HDR is inhibiting it, and eventually




Figure 29. The hunger drive network and the fear drive network, as directly involved in their competition for dominance.

FDC fires. Meanwhile, FDR's activity has inhibited HDC more than the fatigued HDR has excited it, and HDC ceases firing. In this manner FDC fires, HDC is inhibited, and fear becomes the dominant drive.

Continuing, the cessation of HDC's activity lowers the excitation to HD because it eliminates the positive feedback, while the activity of FDC raises the excitation of FD. It follows that HDR remains less active, and FDR increases in activity. As this state persists, the hunger network's fatigue level drops, and the fear network's fatigue level rises. Eventually, HDC manages to fire due to decreased inhibition from FDR. When it fires, it sends positive feedback to HD, and HD, HDR, and HDC receive increased excitation. HDR's increased activity results in increased inhibition of FDC and increased excitation of HDC. The result is the activation of HDC and the inhibition of FDC; hunger becomes dominant. In our scenario, the competitive toggling cycle is completed.

The balances among fatigue, fatigue recovery, feedback, and inhibition are critical in this scenario. Our first conclusion, based on the scenario above, is that HD and FD must fire more frequently in the presence of feedback. Now, it is en-

tirely likely that HD and FD will be fatigued. After all, they both continue firing throughout the competition. We make use of the facts that fatigue is relative, and that the rate of fatigue buildup and decay are adjustable. We reach our second conclusion: HD and FD must be tuned such that the fatigue buildup under non-feedback conditions is low enough to be overcome by positive feedback. This allows us to attain our first goal.

With feedback, then HD and FD will fire more often. For this to affect the competition, we conclude that more frequent firing of HD and FD must result in increased activity in HDR and FDR. This means that HDR and FDR must not operate in the fatigued state most of the time (in our model, if they were fatigued, increased frequency of input would not affect their behavior, only increased level of input). This increased firing rate accomplishes several things: it increases the inhibition of the opposing drive, decreasing its positive feedback; and it increases the activity of the related drive, increasing its positive feedback.

This brings us to the point where one drive has managed to dominate the other. Now, to toggle, we must introduce our last conclusion, about fatigue: at the increased firing rate HDR and FDR, or HDC and FDC, must eventually demonstrate fatigue, or reduced firing due to fatigue of HD and FD. It is this reduced firing which permits the other drive to resume activity. If HDC (FDC) fatigues and fires less often, then HD (FD) receives less in the way of excitatory input, and fires less often. This has the same net effect as HDR (FDR) fatiguing, namely, the inhibition on the other drive from HDR (FDR) is reduced, allowing it a chance to fire and become dominant.

The timing of the fatigue parameters determines the toggle rate, as in the directional neuromime case. The faster a drive network fatigues, the faster the other drive can dominate it. Also, as in the directional case, unless the rates of fatigue and recover are exceedingly well balanced, either the toggling will occur more rapidly until each drive network can only fire a time or two before it is inhibited, or the likelihood of toggling at all will be extremely low.

In the experiment documented in figures 30 through 32, the robot was hungry and moving toward light. We then introduced pain. Figure 30 shows the time period in which the pain began. (The pain signal is not shown, but its commencement coincides with the start of FD's activity.) In figure 30, the hunger drive is seen to be strongly active. The introduction of pain results almost immediately in partial activation of the fear network. FDC cannot fire, initially, due to a buildup of inhibition from HDR. FDR also, in figure 30, has not yet inhibited HDC.

In figure 31, FDR has inhibited HDC, and neither drive is dominant. The effects of HDC's feedback to HD wear off in this image, and we can see that HD is firing less often without it. This reduced firing rate shows up even further in figure 32, where HD is seen in its full fatigued state. The reduced firing of HD, and consequent weakened firing of HDR reduces the inhibition on FDC, and it becomes active. The inhibition from HDR causes it to miss firing once in figure 32, but fear is dominant. In this experiment, fear remained dominant until the pain signal ended. After cessation of pain, it took about 100 ms for the hunger drive's HDC to recover from inhibition and become active.

The usual state for the robot is fatigued hunger and fresh fear (recall that hunger, once active, will remain active until satisfied, but fear fades after the fear-




Figure 30. Pain excites the fear network while the hunger network is dominant.

inducing stimulus ends). Thus, fear normally dominates hunger quite easily. With repeated pain stimuli, fear will fatigue, and both its activation and its inhibition of the hunger network will be reduced. In this event, hunger can dominate fear. This opposing story is shown in the experiment of figures 33 through 35. This time a hungry SENSI was kept in pain and darkness until the fear network was well fatigued. At that point, light was provided. In figure 33, the onset of light can be determined from the start of activity in HDR. As before, HDC cannot fire due to a buildup of inhibition. Note that FDC is well fatigued. The effect of the positive feedback from FDC on the activity of FD is also evident.

Figure 34, which follows figure 33 after a gap of 400 milliseconds, shows




Figure 31. This real-time continuation of figure 30 shows the fear network inhibiting the hunger network. Neither drive network is dominant. The reduction of HD's activity after loss of positive feedback is evident.

fatigue in HD. The hunger drive has still not attained dominance, in fact, the fear network is still in control. But a careful look at figure 34 shows that FD is now becoming fatigued: even with positive feedback, it is unable to fire. This provides the opportunity needed for HDC to overcome its inhibition and activate, as shown in figure 35. Figure 35 also graphically illustrates the effect of feedback on the activity level of HD, and the effect of the lack of feedback on the activity level of FD.

A realistic behavioral scenario for the competition between fear and hunger is that the stronger of the two networks, when both are active, becomes behaviorally dominant. In other words, if SENSI is *really* hungry, fear and its responses should be suppressed. And, if it is badly frightened, hunger should be suppressed. In




Figure 32. This continuation of figure 31 illustrates the fear network becoming dominant. HD is quite fatigued, and HDC remains inhibited.

actuality, we cannot achieve quite this behavior. One reason is the difference between the two networks regarding behavioral response. The hunger network has a distinct response of moving toward any available light. The fear network has no direct innate behavioral response, but does serve to inhibit other behaviors. (We will see a response become associated with fear in the chapters on learning. With a physical response to fear present, it will be possible to observe the suppression of hunger.) Also, while fear network activity may be suppressed or short-lived, the reflexive response to pain, which is hardwired, cannot be eliminated.

What we observe, given extreme hunger and a pain stimulus, is the appropri-




Figure 33. The introduction of light to a hungry SENSI which has felt pain long enough for the fear drive network to be fatigued.

ate reflex, quickly followed by a re-establishment of moving toward light or of feeding. The rapid re-establishment of the hunger response is likely due to the fact that hunger is the dominant drive in the network. In cases of marginal hunger and pain, we observe the associated reflex followed by a pause, as the hunger network recovers from its inhibition by the fear network.

Overall Innate Behavioral Analysis

The general behavior of SENSI on the floor is that of a determined, but not too bright, organism. It plods straight ahead until it hits something or gets hungry.




Figure 34. After 400 milliseconds, HD has fatigued without feedback, but FD is fatigued despite feedback. Fear is still the dominant drive.

If it runs into something, it recoils and heads forward again. Whether or not it hits the same object again depends on the extent of the object and the amount SENSI's heading changed with the reflex. Definitely, SENSI's innate behavior would be much more interesting if some random directionality had been included.

The nature of its recoils is quite satisfactory. Initially, SENSI was spry enough to actually hit an obstacle firmly before the recoil response began. The delay added to the frustrated motion network has eliminated this problem during normal activity. Without the delay, SENSI would begin any motion with all motoneuro-mimes active, and proximity to an obstacle meant a potentially damaging encounter. The delay permits motion to begin slowly, and, providing proper operation of the

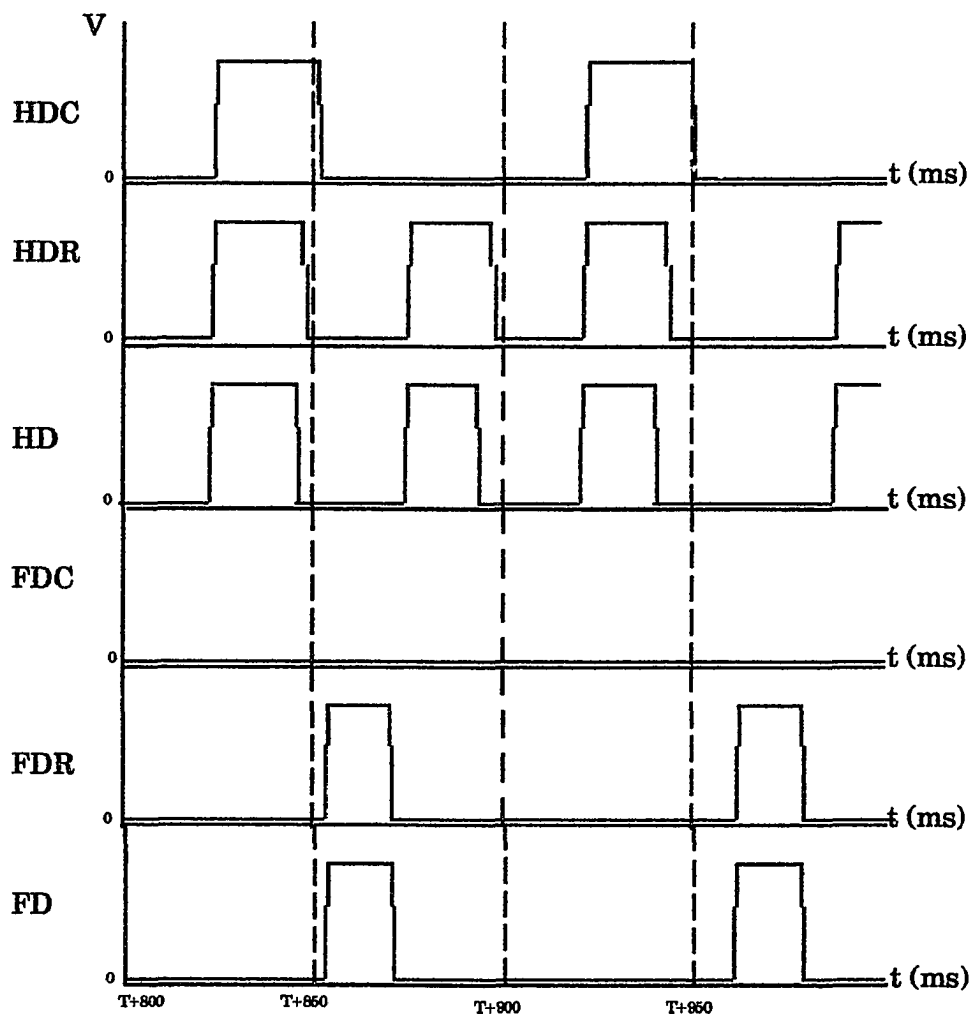


Figure 35. The hunger drive network gains control over the fatigued fear network. The effects of feedback, and lack of it, are evident.

motion sensor, it remains slow enough for recoil to engage before damage. In corners and other tight spaces, whisker damage is still possible. In these cases, opposing reflexes may compete, and by the time a second recoil engages, the first may have caused a firm collision. Despite the occasional damage, we are pleased to note that SENSI has been able to extricate itself, eventually, from the tight corners in which it has placed itself.

There is still a possible problem with the sensitivity adjustment of the light sensors. It is difficult to adjust them to function as desired in all lighting circumstances. If adjusted to room light, they saturate in daylight. If adjusted to daylight or spotlight, they cannot sense room light. Making these sensors more robust is an ongoing concern. The current status is acceptable in a limited way. SENSI can sense light adequate to charge, but only very locally. This keeps it from heading for walls and windows, for the most part, but makes it look rather dense in hunger situations. It currently needs to be within about eighteen inches of the light source in order to 'see' it when hungry.

Once the light is sensed, the observable behavior is reasonable provided there is no obstacle between SENSI and the light source. SENSI turns slowly toward the perceived source, and creeps forward until the light is adequate for charging. It then stops and charges. If the light source is barred from SENSI by an obstacle, its tendency is to alternate between recoil and approach, and any progress around the obstacle seems to be accidental.

Also, unfortunately, the charging time is impractically long. We just can't mount enough solar cell surface area to charge SENSI's batteries in a reasonable amount of time. We also feel that the solar cells are a bit fragile for a clumsy robot. In future models we'd like to incorporate a homing mechanism that permits the robot to plug directly into a wall socket for recharging.

Summary and Utility of Results

The implementation of SENSI's innate network has resulted in several conclusions relevant to others interested in practical implementations of heterogeneous, biologically-based neural networks. The effects of neuromime characteristics and parameters, the utility of two types of neuromime interconnections, and requirements for successful mutual inhibition are reviewed here. Ways in which these results relate to some recent theories and implementations are discussed.

The first conclusion regarding the neuromime model is the usefulness of a

stretched output signature, as opposed to square or pulsed, to obtain desired network activity. We found the stretched output significant in this simple (minimal) network, in areas where each excitatory/inhibitory signal is crucial for a specific behavioral result. These critical areas include mutual inhibition and positive feedback implementations. In essence, we found stretched output necessary to ensure overlap of pulses so as to either avoid or ensure postsynaptic activity.

For our purposes this result indicates that spike-and-decay type output signatures lend stability to simple neural architectures. But (Softky, 1995) offers an interesting alternative perspective, by raising the question of whether neuromime outputs serve to reflect the timing of presynaptic events. In one example he notes:

... the cell tends to fire when it receives those few random fluctuations (extra excitation and missing inhibition) which drive the membrane voltage toward threshold. As these fluctuations are the simultaneous (albeit random) coincidences and anti-coincidences from many different cells, the cell performs a kind of statistical coincidence detection on its inputs; therefore, the output spikes do indeed reflect the timing of presynaptic events.

The stretch we have implemented serves to reduce network sensitivity to presynaptic timing. It was added to minimize the very effect Softky describes. Softky proposes that sensitivity to coincidence provides a potentially major increase in synaptic processing capability by allowing precise, complex, single spike computations. Our work indicates areas where less sensitivity aids in network stability. We conclude that the shape of the output signal is an important parameter affecting neuromime interactions within minimal neural network models.

Fatigue has proven indispensable in this work. This neuromime characteristic, not often incorporated in neural network architectures, is becoming recognized (again) as a useful component of neuron models. In SENSI's network, fatigue and inhibition work together for behavioral choice when behaviors are mutually exclusive. Inhibition results in a 'winner', and fatigue permits other behaviors to opportunistically dominate motor responses.

Recent work acknowledging the utility of explicit inhibition and fatigue is described in (Blumberg, 1994). As in our model, Blumberg's network is hand-tuned to obtain, as he puts it "persistence of behavior balanced with opportunism." He states that persistence is modulated by means of tuned inhibitory gains. In SENSI's network, inhibitory synapse strengths are fixed, and the persistence/opportunism balance is obtained via two fatigue parameters: rate of fatigue, and rate of recovery from fatigue. Each of these affects behavioral choice in a foreseeable fashion. The

rate of fatigue affects the length of time a competing behavior remains dominant; the recovery rate helps define the period during which a recently active behavior can be overridden by some other valid activity.

(Sahota, 1994) investigates the use of interbehavioral bidding, in a non-biologically oriented system, to better reflect the dependencies of actions on the current situation. This is achieved by the combination of 'goal' and 'situation' inputs before bidding, which is very similar to the function of the drive-representation neuromimes we have implemented. In Sahota's work, the persistence/opportunism dilemma is referred to as a starvation vs. stability problem. The use of 'impatience' and 'boredom' is suggested for the avoidance of starvation, and positive feedback is advocated for stability.

Sahota's bidding was hand-tuned to achieve the desired behavioral dynamics, but as with Blumberg, he doesn't discuss the details of the tuning which achieved his results. Our research has led to several conclusions regarding competition via mutual inhibition which serve as guidelines for other research. These results are based upon the premise that mutually inhibitory behaviors should toggle under conditions of persistent, comparable excitation. This premise requires that each behavior is capable of inhibiting the others under comparable excitation, and that as the dominant behavior fatigues, the others are capable of opportunistic firing (due to reduced inhibition from the fatigued behavior).

When the behaviors involve networks, i.e. the mutual inhibition is not directly between neuromimes, an additional component is necessary to ensure the desired behavioral toggling. This component is positive feedback within each behavior, which serves to produce greater activity in the 'winning' behavioral network. The details of the feedback interaction are given in a previous section. This element of drive-based behavioral networks was incorporated by Grossberg (Grossberg, 1987a) to enhance persistence in the absence of competition, but we've found it plays a vital role in any competition which may occur.

Two useful types of neuromime connections were developed during implementation of SENSI's innate network: a delay-inducing connection, and an input-limiting connection. A 'delayed' reaction is obtained by requiring excitatory input to build up to a suprathreshold level. We used it to represent a buildup of frustration. The resulting delay is dependent upon the excitation accrual and decay rates, as well as upon the firing characteristics of the presynaptic cell(s).

To limit excitation, which we required when implementing polymorphic neuromimes (those requiring two or more distinct types of inputs to become active),

we utilized a simple form of dendritic tree. (Lynch, 1986) states that biological evidence suggests cortical cells "are designed to respond to combinations of synaptic contacts from many different afferents," and implies that the design is manifested in longer, more elaborate dendrites. We implemented this concept by having each input of a given type converge within a 'tree' prior to the postsynaptic neuromime. The synapse weight at the postsynaptic neuromime (the 'trunk') limits the total excitation presentable for the type of input entering the tree. Contributions to the tree can be biased by varying local weights at the junctures of presynaptic neuromime and input tree.

The elements of SENSI's innate network function together for behavioral choice at each instant. Others have referred to this decision process as action selection. (Brooks, 1994) claims that action selection has been a focus of research, but coherence, related strongly to attention, is a larger issue to be addressed. Curiously, his discussion of coherence lists three components, two of which are already integral parts of SENSI's network. His list includes internal drives, which he claims are a primary mechanism for action selection, internal rewards, and mutual exclusion. He includes internal rewards as a means of providing additional hysteresis to active, successful behaviors. His hardware implementation is processor-based, with message passing. He does not describe the form the listed elements will assume, and it will be interesting to note if he eventually incorporates some type of timing or fatigue mechanism in order to achieve mutual exclusion which balances behavioral hysteresis and opportunism.

Each of these characteristics and parameters: stretched output, fatigue, positive feedback, simple dendritic trees, and a delay mechanism, have proven vital for some aspect of SENSI's integrated innate behaviors. The concepts behind them are generalizable to any network requiring similar interaction properties. They also serve to indicate the complexity extant in a simple network, and the variability which can result from simple parameter changes.

CHAPTER VI

NETWORK IMPLEMENTATION OF CLASSICAL CONDITIONING

We're still learning about learning, and the complete story
will not be written for some time.

Fred S. Keller

The capability for learning which is implemented in SENSI is restricted to the form of learning described in Chapter II as simple, or classical conditioning. Briefly, this type of learning involves the creation of an association between a stimulus (or combination of stimuli) that initially has no innate response, and a response or reaction which occurs in the presence of some other stimulus. Building of the association requires a number of instances in which both stimuli are active in close succession, and the effect is the elicitation of the response by either stimulus alone.

The connectivities described in this chapter are based on Stephen Grossberg's work, in particular, on (Grossberg, 1987a), in which he discusses the interdependencies of reinforcements, drives, motivation and attention. Much of the terminology used in this chapter is taken from his work, as well. It should be noted, however, that this implementation represents a relatively small fraction of the theory developed by Grossberg. In the following sections, the structure of the learning network is detailed and the basics of its operation are given. Chapter VII continues the discussion with an analysis of the learning network in action, the consequences of its limitations, and the applicability of the results.

Hardware Plasticity

The first requirement for a network that learns is that its connection parameters be 'plastic' or capable of change under controlled circumstances. It is generally accepted that this plasticity is found in the synaptic connections between neurons, and involves adaptability in the influence one neuron is capable of exerting on

another. This influence is referred to as synaptic 'weight'. In SENSI's network, synaptic weights are associated with resistors; the higher the resistance of a synapse, the lower the weight, or influence, of pre-synaptic signals on the post-synaptic neuromime. This resistance effectively controls the voltage presented to the post-synaptic neuromime.

Once the synapse is viewed as a modulator of voltage, it follows that plasticity is obtainable by the replacement of the resistor with a variable, or controllable voltage modulator. The plastic synapse we have designed for SENSI is based on this concept. It is schematized and described more fully in Appendix D, for those who wish to see the design details. It is sufficient to note here that the design is flexible in its learning rate, decay rate, and ultimate influence with respect to the post-synaptic neuromime.

If we could just replace our resistor synapses by plastic synapses and obtain learning, we'd be almost finished, but there is (much) more to the learning circuitry than that. Following Grossberg, cells capable of learning (i.e. those which have outputs connected to plastic synapses) must be buffered so that learning occurs only under appropriate circumstances. The network components, called *sampling cells*, which provide the buffering of learning are discussed next.

Sampling Cells

The buffering of plastic synapses is accomplished by isolating them in the network from spurious activation. They are restricted to being the *output* synapses of designated cells called sampling cells. A sampling cell is the same as any other neuromime except for its synapses. The isolation, therefore, occurs from the careful selection of the *inputs* to the sampling cell. Only when a sampling cell is active, or firing, can its plastic synapses sample and adapt to the excitation levels of the post-synaptic neuromimes.

When the sampling cell is active, its output synapses incrementally adjust their weights to reflect the excitation levels of the neuromimes to which they connect. Upon repeated sampling of consistently active post-synaptic cells, the plastic synapses become capable of providing the pattern of activity they have sensed on those cells. (Note that the *pattern of activity* is *not* absolute excitation, but rather a fraction thereof. Grossberg maintains that it is the spatial pattern which is important; absolute excitation levels are governed independently. The absolute possible influence of each synapse in SENSI's network is dependent upon an associated

resistor value.) The result of this process is that the active sampling cell becomes capable of reproducing that pattern of activity by itself; the network has *learned* to produce the pattern of activity under new circumstances (those which activate the sampling cell). The design issue, then, becomes one of determining the proper inputs for excitation of the sampling cell. As will be seen, these inputs involve not only external sensory cues but internal drive (motivational) cues as well.




Figure 36. The initial state of the network showing the UCS and its response, the CS, and the missing connections where learning circuitry is to be inserted.

Learning Network Basics

Some terminology is necessary to simplify the following discussion. As is common in the literature, we refer to the external cue which has an innate response as the unconditioned stimulus, or UCS. The innate, or unconditioned response, is called the UCR. The cue which we want to become associated with the UCS's response we refer to as the conditioned stimulus, or CS.

Figure 36 illustrates the initial situation, where the UCS excites its UCR. The CS connections are unknown, but do not involve direct excitation of the desired response. Our goal is to fill in the 'learning' box.

The first condition is that the network should only learn to associate the response with the CS when both the UCS and CS activated cells are excited. This is the time when response neuromimes can be productively sampled (due to their excitation by the UCS). The consequence of this condition is the necessary excitation of a sampling cell by outputs of both the UCS and CS cells. The synapses of this sampling cell are connected to the response neuromimes. This network, shown in figure 37, is not the entire solution, however. The desired result of conditioning is

that the CS be able to elicit the response on its own (in absence of the UCS). In the network of figure 37, however, both the UCS and the CS must be active to fire the sampling cell. Thus, while the sampling cell can learn to produce the desired response, the presence of the CS alone is insufficient to activate the sampling cell and elicit the response.




Figure 37. A first attempt at the learning mechanism connections. The shaded cell is the sampling cell; plastic synapses are shown as filled circles.

Recall that the sampling cell's purpose is to buffer the learning (i.e., to restrict the learning that occurs to instances of UCS and CS overlap), so it must be excited by the connections shown. The only way out of the dilemma is to admit other neuromimes into the learning mechanism (Grossberg, 1987a). In particular, the CS must learn to provide the excitation to the sampling cell that is initially provided by the UCS. This implies there is more than one sampling cell involved: one to sample the desired response, and one to sample the excitation from the UCS. Figure 38 shows the network given by Grossberg as a solution to the problem.




Figure 38. A revised learning network whereby the CS becomes conditioned both to the desired response and the arousal (labeled 'A') initially provided by the UCS.

In the network of figure 38, the first conditioning cell excited by the CS samples and learns an activation caused by the UCS. At the same time, it also passes the CS activation to the second sampling cell. When the UCS and CS are simultaneously active, the CS learns to provide the activation required to fire the second sampling cell, which concurrently learns the desired response. Before conditioning, the CS alone cannot fire the second sampling cell. After conditioning it can, because it has learned to provide the needed excitation to the arousal neuromime.

The story of figure 38 is still not complete, however. The arousal neuromime is strongly connected with a drive of the network. The connection between the arousal needed for learning, and the drive networks we have already established, is explored in the next section.

The Arousal Role of Drives

With respect to the drive networks, it is helpful to consider the internal stimulus which excites a drive as a *context* for action. In the Pavlovian conditioning experiments, the context of hunger together with the presence of food resulted in the (innate?) response of salivation. In the drive framework developed in chapter IV, hunger is an internal cue, the presence of food is an external cue, and the combination serves to excite the hunger drive representation neuromime.

Now suppose we consider the conditioning of some CS, say a tone, to the UCS of the presence of food, and its UCR of salivation. So, the food is presented some short interval after a tone is sounded. In what *context* does this experiment result in the desired outcome? Only when hunger is present! In the context of hunger, the animal learns to associate the sounding of a tone with the presence of food. This is a point of importance because it allows us to pinpoint which portion of the drive network provides the arousal sampled in figure 38. It cannot be just the internal drive! The arousal source is the drive *representation* portion of the network. What the network is learning is an additional external stimulus which, in the event of hunger, produces the salivation response. In essence, the tone representation learns to provide the excitation (innately provided by the presence of food) which, combined with hunger, excites the drive representation.

There is a complication associated with this conclusion however, which serves to illustrate the special nature of unconditioned responses. In our understanding of drives, drive competition, and responses from Chapter IV, we would expect the salivation response to result only in the event of dominance of the hunger drive

network. In figure 38, however, it appears that the UCR is excited directly by the UCS. In other words, figure 38 fails to illustrate the context we have just determined to be relevant. This leads us to consider the possibility of two forms of innate responses. The first, as just discussed, requires a drive context to give meaning to the stimulus (the meaning is indicated by the response). A second, which is herein assumed to exist, is the direct connection of a stimulus to a response with the concurrent excitation of a drive neuromime (by the stimulus). This type of connection seems likely in the cases of reflex responses, where the drive neuromime excited is fear. Figures 39 and 40 illustrate the learning networks (so far) for each case.

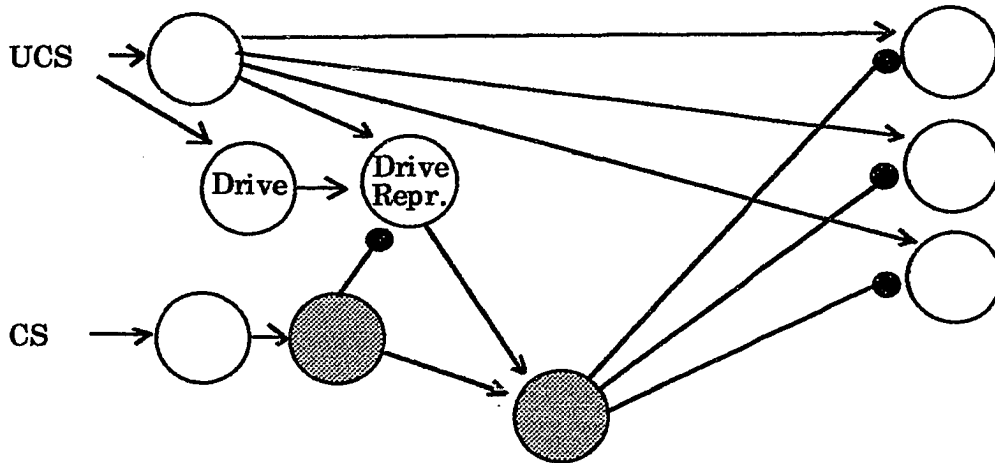


Figure 39. The learning network for the case of a stimulus conditioned to an innate reflex response. The conditioned response cannot occur in the absence of drive activity (which may have other sources of activation in addition to the UCS shown).

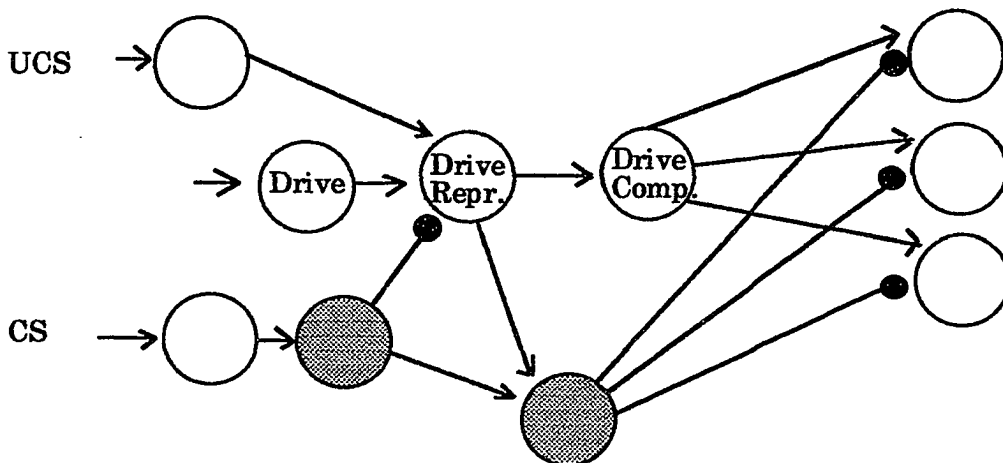


Figure 40. The learning network for a stimulus conditioned to an innate, drive-dependent response (note the drive competition neuromime, as described in Chapter IV). Here also, the drive must be active for the conditioned response to occur.

There are two differences in the networks: The drives are activated by differing sources, and the excitation locus of the UCR varies. Perhaps both of these types of networks exist biologically. Grossberg indicates that the innate UCS/UCR networks may differ from CS/CR networks in that the innate UCS provides both cue and arousal inputs as opposed to a CS providing only a cue input. More important, however, is the common effect of the learning. In either case the CS, by virtue of learning to provide both cue and arousal information (in the context of a given drive), can serve as a UCS in further experiments.

Other Drive-Related Considerations

Before discussing how this all ties in with SENSI's actual network, there are a couple of aspects of figures 39 and 40 which need to be touched upon. In relation to figure 39, which has been presented as a possible network for fear-induced responses, there is a subtle issue of *what* is learned with respect to fear. A natural expectation, when considering learning and fear, is that an organism learns to be afraid of certain stimuli or events. In essence, there is the learning of fear on one hand (called a conditioned emotional response) and on the other hand, the learning of a response to a cue when fear is already present (i.e., in the context of fear). The difference is one of learning an emotional, internal state versus learning a motor response. This latter type of learning is the simple conditioning being focussed upon here. (As an aside, however, a cue could be conditioned to result in the excitation of fear itself; one way to do so is to connect the network such that the cue samples and learns the arousal from pain, not at the drive representation neuromime, but earlier in the network, so that learned arousal excites FD and FDR like pain does.)

An issue which concerns both types of network is drive competition. From chapter IV, we have that drive competition neuromimes function as a choice mechanism for behavior. First, active drives and cues interact at the drive representation neuromimes. Active drive representation neuromimes then enter into competition with each other. The 'winner' of this competition gains control of the motor response. Now notice in each of figures 39 and 40, that there is a sampling cell which is learning a motor response. We need to define how this cell and its learned response fit into the competition for motor control.

A look at the inputs and outputs of the second (rightmost) sampling cell of figure 40 shows us that it has the most in common with drive competition neuromimes: it must receive excitation from a drive representation neuromime, and its




Figure 41. The additional inhibitory connections required to maintain the competition for motor control which is the basis of behavioral choice.

output can excite motor activity. It becomes necessary, then, to include the second sampling cell in the drive-related behavioral competition. The required connections are illustrated in figure 41. (The connections shown are with respect to the scenario of figure 40. Similar connections are required for the network of figure 39.) With the additional connections, new learned behaviors compete with other possibly relevant actions. The behavioral choice mechanism, based upon prevailing internal and external conditions, is maintained.

Application of Learning Networks to SENSI

To incorporate simple conditioning into SENSI's nervous system, we need to identify the unconditioned stimulus and response, the conditioned stimulus, and the arousal elements. The decision was made to provide the capability for SENSI to learn to respond to touch in the presence of fear.

For our purposes, pain is the unconditioned stimulus, with the recoil reflex (or combination thereof) being the response. Touch sensation on some whisker is the conditioned stimulus. The arousal element, identified as a drive representation neuromime in Grossberg's work, is that drive activated innately by pain, which is fear. Figure 42 is the learning network (of figure 39) redrawn, with applicable




Figure 42. The learning network showing the involved components of SENSI's network for two arbitrary whiskers. Analogous wiring is required for the remaining whiskers.

neuromimes of SENSI's network inserted. Also shown are the drive competition components introduced in figure 41.

On a practical note, we must explain the manner by which the common interneurinomimes can be used in the learning of the desired motor response. Recall that there are two sets of interneurinomimes: one of which is used for the reflexes, and the other (common) set which is used for all other motor connections. As noted in chapter IV, this division is not biologically correct, and now it has come back to bite us. For we want the touch to produce a response via the common interneurinomimes, yet it is the set of reflex interneurinomimes which is responsible for that response. It was noted in chapter IV that this division of interneurinomimes was for testing and robustness purposes. These purposes are not defeated by the compromise that must now be made.

In order for touch to learn to excite the common interneurinomimes, there must be a connection between the excitation that the reflex interneurinomimes receive and excitation of common interneurinomimes. This one-way connection has been made (see Appendix E). The result is that the common interneurinomimes receive *all* excitatory interneurinomime inputs. There is no effect on the reflex interneurinomimes.

The immediate behavioral effect is that all recoil responses have twice the neuro-mime activity driving the motors (both reflex and common), so, if anything, they will be a little brisker. With respect to learning, it is now possible to sample, and learn to excite, the common interneuromimes based on reflex behaviors.

From figure 42, it can also be seen that each whisker, or learning point, requires the addition of two sampling cells and nine plastic synapses. The number of synapses is due to the generality of learning – we cannot assume which directional and interneuromimes will be active, and must sample all eight of them. We decided to group touch together, in order to reduce the amount of hardware required for the learning circuitry. In the end, the whiskers were paired, giving five sets of sampling cells and synapses (10 sampling cells and 45 synapses). Ideally, we'd like to see SENSI learn a balance of the possible motor responses for each whisker pair, biased by which occurs more often in a learning context.

Consequences of Generalized UCS Arousal

The labeling of pain and touch in figure 42 is important. They are labeled 'X' and 'Y' to point out an effect of the current learning network. Because arousal from the UCS is sampled (nonspecifically) at the drive representation neuromime, it cannot be assumed that pain and touch occur on the same whisker during any instance of learning. It is therefore possible for SENSI to learn a response to a whisker's touch which is very different from the recoil response associated with that whisker's pain sense. All that is necessary is that pain exist *somewhere* for any concurrently touching whisker to attempt to learn the current motor action.

Is there a way we can reconstrain learning to avoid this? Suppose pain *and* touch were required on a given whisker in order for that whisker to learn a response to touch. Would this eliminate the possibility of 'innappropriate' learning? Unfortunately, it won't. Any time SENSI gets itself into a corner, or even up against a wall, it can simultaneously sense pain and touch on several whiskers. The motor responses compete in such cases, and what is incrementally learned by all contacting whiskers is the dominant motor response at each moment. Therefore, even the most stringent of constraints upon the learning circuitry itself cannot eliminate inappropriate learning. There is a refinement which helps to reduce it, however; we discuss it in the following section.

Learning Refinement: Sensory Focus

Until now, we have not needed to introduce sensory focus into SENSI's network, primarily because of the robot's limited sensory capabilities. But learning has made it necessary. Each pain and touch sensor, when stimulated, has an innate hysteresis which constitutes short-term memory in this simple nervous system. Now we see that this short-term memory must be limited in scope. What is required is a limit to the number of sensory inputs which can concurrently affect the greater network. Competition among sensory events becomes as important to SENSI's behavior as competition among drive representations.

We are not in a position to fully implement sensory competition, but for the sake of more competent learning, competition was added among the touch neuromimes. With this competition in place, at most two touch neuromimes can be simultaneously active for an extended period of time, and those two must be adjacent on SENSI. The competitive connections rely upon inhibition and fatigue, and are shown in the network diagram of Appendix E.

These additional connections ensure that only the most recent touch(es) sensed are active after any quick succession of whisker contacts. It also results in the toggling of activation of sumultaneously contacting, mutually inhibiting touch senses. The effect during learning is a reduction of learning to at most two whiskers at any given time, which reduces the scope of inappropriate learning during competing reflexes. During recall, this addition helps keep learned reactions current to the most recently occurring touch. Ideally, it can help SENSI continue to turn away from an obstacle as adjacent whiskers make contact, allowing the robot to clear the obstacle without incidence of pain.

Learning Refinement: Synapse Specifics

As we implemented our learning network, we found a lack of literature relating to the placement of the plastic synapse on the sampled cell. This placement, variable in several respects, turns out to critically affect learning results. Our plastic synapses were designed to sample net input, meaning they were intended to sample the input to the postsynaptic cell after all excitatory and inhibitory inputs had combined. This setup allowed us to include the synapse's input in the sampled sum. We quickly discovered a problem with this inclusion. If any net excitation was present on the sampled cell, the synapse would sample it and learn to produce some

fraction of it. The next time, assuming the same net external excitation, but with the added (learned) input from the plastic synapse, the synapse was exposed to a higher net excitation, so its capability to produce excitation increased a little.

The problem with this scenario is that the original, net external excitation, which was being sampled, could have been inadequate to fire the postsynaptic neuromime. But as the synapse strength increased, eventually the net excitation from the external source, combined with the learned excitation, would be enough to fire the postsynaptic cell. Thus we were faced with the network learning to fire a neuromime *that had never fired during learning trials*.

This problem arose with the network of figure 42. Given touch and fear (but no pain), the excitatory input to FDR from FD was sampled and learned by the synapse. After several trials, during which FDR never fired, it suddenly did under the combined inputs from FD and the plastic synapse!

Naturally, the first thing that came to mind was that the synapse was learning too great an influence from the input of FD alone. Here is the desired scenario: if only FD input is present when net input to FDR is sampled, then the synapse should eventually learn some fraction of it, *but not enough to fire FDR* together with it. Only if both pain and FD inputs are present during sampling, should the synapse be able to learn a strength adequate to fire (in concert with FD) FDR. The problem? Because the synapse samples its own input summed with that of FD, it will continue to increase in strength, up to its limit, regardless of the existence of the pain input on FDR at the time of sampling. And the limit strength of the synapse must be high enough to fire FDR when combined with input from FD, for this is the purpose of the learning.

The next natural conclusion is that we are sampling the wrong excitation, which is the conclusion we eventually reached. The goal in SENSI's network is for the first sampling cell to learn the excitation to FDR normally provided by the UCS (pain). It is not the goal to learn the entire excitation present on FDR. (We differ from the philosophy offered by Grossberg on this point; he specifically states in (Grossberg, 1987a), Chapter 1, that the pattern learned at the drive representation is a mixture of both drive and cue inputs.)

What we have done, therefore, is changed the location of the plastic synapse from the 'body' of the sampled cell to the incoming input from the UCS. This 'synapse-on-a-synapse' structure is not unheard of (Sheppard, 1990), but no documentation of its existence in this type of circumstance was found. Our pictorial terminology for the synapse-on-a-synapse connection type is given in figure 43.




Figure 43. The diagram representing our synapse-on-a-synapse structure.

We must also be careful, under these learning circumstances, as to the placement of the input to the postsynaptic cell from the synapse. If the input is placed on the cell body, FDR is once again susceptible to inappropriate activation. This time, it's because the pain input and the learned input can combine to fire FDR, without input from FD. (This can't really happen in SENSI's network, but it is possible in other learning situations where the UCS does not directly excite the drive neuro-mime.) Biologically, we'd expect the locus of both functions (sampling and input) to be the same, and it turns out to be necessary in our implementation. The input sampled and the input learned must converge within a dendritic tree prior to arriving at the postsynaptic neuromime.

Consequences of Synapse Placement

The placement of a plastic synapse's sampling point and output point with respect to the postsynaptic neuromime has been shown to have an effect on learning results. The placement of a plastic synapse's sampling point and output point *with respect to each other* is also important. If the sampling cells do not sample their own excitation to the sampled cell (sampling point is before output point), then learned excitation decays with use. If they do sample their own input (sampling point after output point), the learning becomes self-reinforcing and permanent, once at a suprathreshold level. The nature of the desired motor learning led us to decide against self-sustained learning of motoneuromime excitation. For SENSI, this means each time it uses such a learned connection, it will lose some of the connection strength. The way this follows from the learning model used is described in the

next chapter. For now, we note that periodic refreshing, or relearning, is required to maintain learned connection strength.

The nature of the desired motor learning, as just mentioned, is that of procedural learning. It is meant to be adaptive over time. Self-sustained learning, and 'permanent' long-term memory are more appropriate for declarative learning. We stated earlier that we have paired the touch inputs for learning, and wish to see SENSI learn a combined response where indicated to avoid pain. If the synapses are self-sustaining, SENSI cannot shift the learned balance of motor response as learning conditions change. For the motor interneuromimes, we opted to abandon 'permanent' long term memory for 'timely adaptability'. Is there a way to have both adaptability and stability in learning, given our model? The answer lies in reinforcement, and is discussed next.

Learning Refinement: Simple Reinforcement

Because of the synapse arrangement, SENSI's learning is lost when it is actively used. We decided to implement simple active reinforcement to counter this active decay. What we want is a method which allows SENSI to maintain 'good' learning, while permitting 'bad' learning to fade away naturally.

The simplicity of the desired learning allows us to pinpoint a practical evaluation of its 'goodness'. If SENSI is performing a learned response to touch, and pain occurs, the response is considered to be inappropriate and should be subject to decay. As long as pain does not occur during the performance of a learned response, the response should be reinforced. We accomplish this active, selective reinforcement for the motoneuromimes in SENSI's network by the addition of the network shown in figure 44. Such a network has been implemented for each directional and (common) speed neuromime, namely, the neuromimes which are sampled by the S2 sampling cells.




Figure 44. A method for implementing active, selective, positive reinforcement on the motodirectional neuromimes. See text for an explanation.

In the network an additional neuromime, labeled R^+ for positive reinforcer, is added for each motodirectional neuromime. R^+ becomes active when one or more S2 sampling cells are active simultaneously with the motodirectional neuromime. Its output excitatorily influences the motodirectional cell, after the sampling point, so its effect can be sampled (which provides the reinforcement). R^+ is inhibited if pain occurs, allowing active decay. Thus when the response learned by any S2 is active, it is positively reinforced as long as pain is avoided. This mechanism results in adaptive motor learning without active decay of useful learned responses. Inappropriately learned responses actively decay with pain.

We've also added a comparable reinforcement network for the synapses on FDR. The learning is reinforced when activation is received from one or more S1 sampling cells and FDR. Reinforcement is inhibited by pain, to be consistent.

One final network change incorporated for the benefit of learning has been to disable the startle reflex for touch on the back given forward motion. The reason for this is that in almost all cases where back whiskers touch, the motion is predominantly forward (usually SENS1 is turning away from a side contact). In these cases, learning could not be utilized because startle became active and inhibited the common interneuromimes. Recall that it is activity upon the common interneuromimes which is sampled and learned. The disabling of this portion of the startle reflex permits the availability of learned responses on the rear four whiskers.

The network embodying all of the final learning connections is complex, and is not shown here. The connections may be seen in the network diagram of Appendix E.

CHAPTER VII

ANALYSIS OF THE LEARNING NETWORK

My emphasis on anger and fear...is not because they are more important but because they provide most of our evidence. They are easier to observe and record because they often appear as a sharp break in the ongoing flow of behavior.

Donald Olding Hebb

Having added the connections discussed in the previous chapter, we are now in a position to determine if SENSI can learn the desired behavior. Recall that our aim is for SENSI to learn to respond to touch when touch occurs in the context of fear, and via the learned responses, to avoid painful contact when afraid. Before discussing the experimental results, however, there are loose ends to be tied, that have been mentioned in previous chapters: the self-excite, or lasting excitation of the fear neuromime, and the competitive relationship between learning and other behaviors. Following the conclusion of these topics, the behaviors of the neuromimes and synapses involved in learning are documented, and the behavioral effects of the learning network are discussed. Finally, a summary of conclusions and results is presented and related to other research in artificial neural networks and learning.

Lasting Fear

It was noted earlier that the fear neuromime (see figure 11) is unique in its property of self-excitation. This property allows the fear neuromime, once active, to maintain its activity for a short period after its external excitatory input has ceased. We have earlier termed this property 'subliminal fear'. The existence of this property was motivated by (Hebb, 1980): "In fear and anger the sympathetic nervous system is excited, and the aftereffects of the arousal last for some time." And, for our purposes, subliminal fear is crucial.

The goal is for SENSI to learn a response *in the context of fear*. For this response to be meaningful and observable, SENSI must operate in the context of

fear. But the innate system is direct: fear exists in the presence of pain. Recoil reflexes occur in this context also. We use this to our advantage as it provides the ingredients for conditioning. In order to observe the results of conditioning, indeed for them to be of any utility to SENSI, it must be possible for SENSI to be afraid, and active, without pain and recoil.

And thus we come to the need for subliminal fear. After pain has ceased, and the recoil has terminated, SENSI maintains activity in the FD neuromime. From figure 43, the diagram of the learning network, it can be seen that this activity in FD, together with touch, can elicit any existing, learned behavior. From this, we can deduce the effects of learning on SENSI. If it has experienced pain often enough to have learned a response to touch and fear (both of which are active with pain), then, when it experiences touch alone in the presence of subliminal fear, it will react with the learned response.

The last sentence requires explanation. There are four distinct circumstances in which touch can occur without a response. The first is if no learned response exists in the network, i.e., SENSI has not yet learned to respond, or has forgotten its response, to touch. The second occurs if the touch stimulus is presented some time after the last incident of pain. In this case, subliminal fear will no longer be active, and SENSI's learned response is, networkwise, inaccessible except in the context of fear. Third, if SENSI experiences pain for a prolonged time, the fear drive neuromime FD will fatigue, and its self-excitation will fail to keep it active beyond the pain input. Lastly, as can also be seen in figure 43, if SENSI is very hungry, the learned response can lose the competition for dominance of the motors. This issue of competition among behaviors is discussed next.

Learning and Behavioral Choice

In chapter VI, and figure 42, the placement of the learning network within the innate network was described. This placement permits the competition between learned and innate behaviors. What was not mentioned there, but which is equally of importance, is that the setting of the learning network affects the learning itself. What I mean by this is that whether learning occurs, and even what is learned, can be affected by the activity of SENSI's innate network.

Suppose, for instance, that both fear and hunger are active and competing. When touch occurs under these circumstances, the learned effect will be quite dependent upon which of the innate drives is dominant. If fear is dominant, then

learning will proceed much as previously inferred. If hunger is dominant? Well, there are several consequences. First, with hunger dominant, FDR's firing will be limited, while HDR's activity will be quite strong. From this, and the diagram of figure 43, it can be seen that the sampling cell associated with the motor responses may well be inhibited, and learning will not occur. This is not unreasonable; many learning opportunities have been thwarted by hunger or other motivational distractions!

Next consider the consequences of the sampling cell not being totally inhibited. In this case, learning may well occur, but with the recoil response completed, FD still active, and fear subsumed by hunger, any opportunistic firing of the sampling cell will be sampling the motor responses of hunger rather than of fear. In the best case, where a strong learned response is already present, this will serve to adjust it somewhat.

The above considerations quite naturally indicate parameters to which the learning network and SENSI's resulting behavior will be sensitive: the rate of learning, the rate of active decay (or relearning) and the rate of passive decay (forgetting). The balance of these parameters will determine if SENSI can learn what we intend it to, namely, to turn away from obstacles upon touch, if pain has recently occurred.

Learning Network Activity

Now for a look at the activation of the sampling cells, and the plasticity of the synapses. In this section, a picture is given of the parameters mentioned above, namely, the rates of learning and decay. The following diagrams focus on the learning, by the first sampling cell, of the arousal at FDR from pain.

Figures 45 through 47 show the operation of part of the learning network before complete conditioning, during conditioning, and during a trial in which the learned response is activated. In figure 45, the level of the plastic synapse's output, labeled FDR_L , is not high enough to excite FDR. Without FDR, the second sampling cell, $S2_{TFL}$, cannot fire and release the learned response. Under these conditions, SENSI will continue whatever behavior is dominant at the time of the touch. (For these tests, we scare SENSI with a rapid change in light. This allows us to activate FD in the absence of pain.)

The trial shown in figure 45 is a learning trial (because the first sampling cell is active), as are all trials where S1 and/or S2 are active. We have wired the sam-




Figure 45. Touch and fear combined before SENSI has learned to respond. Although the plastic synapse shows some output, it is not enough (together with FD) to fire FDR. The second sampling cell (which has synapses to the motor response) cannot fire without FDR active.

pling cell so that it does not sample its own input; since FDR is not active in this trial, decay of the learned level occurs. In other words, the level of excitation sampled at FDR is zero, because there is no pain, so the learned amount providable will decrease (this is the active decay discussed in the previous chapter).

In figure 46, a portion of a positive learning trial is shown. We can see the activation of pain (which excites FD and FDR), and the activation of touch. When touch and FD overlap, the first sampling cell, S1_{TFL}, can fire. When it does so, it




Figure 46. Pain on whisker 1 results in fear, sampling of FDR, and sampling of the motor reflex response.

samples the excitation present at that time on FDR. The combined excitations of FDR and $S1_{TFL}$ permit the firing of $S2_{TFL}$, and the consequent learning of the reflex motor response (not shown).

After several learning trials, the level of excitation provided by the first sampling cell's learning synapse is adequate to fire FDR in the absence of pain, as shown in figure 47. In this trial, fear becomes dominant over boredom, and the learned motor response is excited by the second sampling cell.




Figure 47. Pre-existing fear, combined with touch, elicits learned activity on FDR and the motor response excited by plastic synapses emanating from $S2_{TFL}$.

Our experiments have indicated that it takes approximately six learning experiences for SENSI to become capable of responding to the combination of touch and fear. The exact number of trials can depend on component variation and on the fatigue levels of all involved neuromimes during each learning trial. The response will occur at a level of learning somewhat below the maximal; additional learning opportunities will maximize the synapses' responses. The learned influence will persist, due to the reinforcement, as long as the response is used regularly. Without use of the response, the synapses will retain usable levels for approximately twelve

hours. This is the measure of passive decay: how long SENSI can retain unused, unrefreshed memory.

The next diagram illustrates the flexibility of the synapses which learn the motor response, as well as demonstrates how inappropriate learning can occur. The

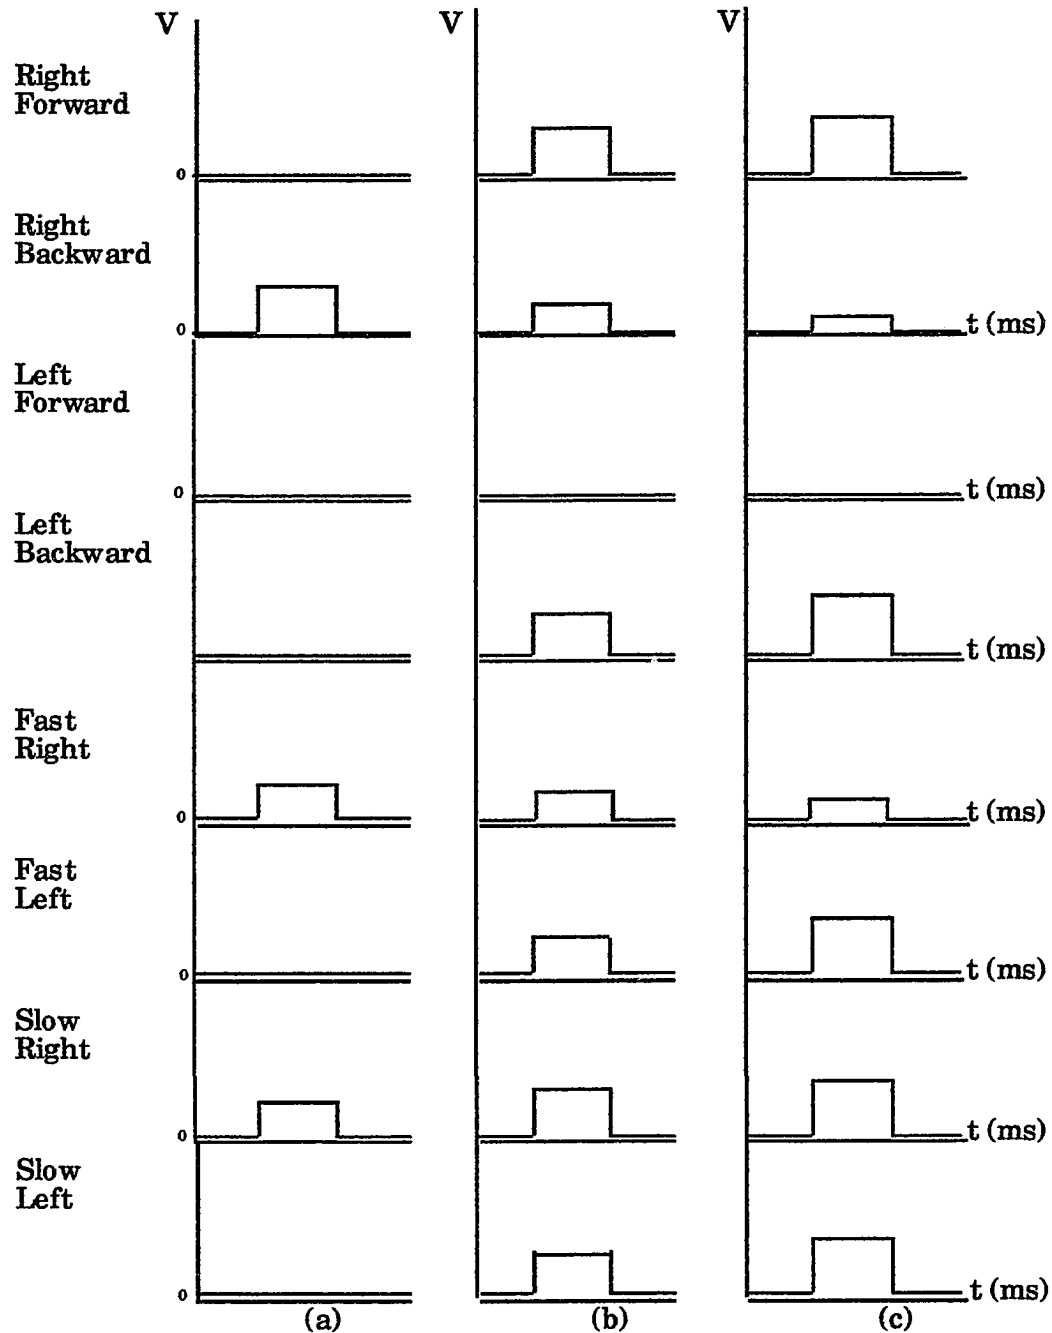


Figure 48. An illustration of retraining the motor synapses of the front, right learning network. (a) synapse output after four learning trials using whisker 1; (b) synapse outputs after four relearning trials using whisker 3; (c) synapse outputs after three additional trials with whisker 3.

figure, divided into three time periods, shows an initial response learned to a single (inappropriate) whisker, and the progressive loss of that response as we train the synapses to an appropriate whisker. The number of learning trials associated with part (a) of figure 48 is four; these trials were obtained by eliciting pain on whisker 1 concurrent with touch on whisker 2. Because touch on whiskers 1 and 2 are not mutually inhibitory, the synapses for whisker 2 can learn the pain response to whisker 1, namely, fast reverse on the right motor. Part (b) reflects the status of the motor synapses for the front right after four relearning trials, where pain and touch occur simultaneously on whisker 3, and part (c) is the output picture after three additional relearning trials.

From figure 48, we can tell that the directional learning proceeds a bit faster than the speed learning. This is because synapses connected to the directional neuromimes are sampling an input that comes directly from the pain signal. This input is quite strong. The synapses connected to the speed neuromimes, on the other hand, sample an input which comes from the reflex speed neuromimes, an indirect, and weaker, pain-related signal. So, even though the same sampling cell is producing the sampling signal for all of the synapses in the figure, the result of the sampling depends upon the overlap of the sampling cell's output and the sampled cell's input.

It is also plain in the diagram that the unlearning (active decay) occurs more slowly than the learning. This is an effect due to resistor choices and characteristics of capacitor charge and discharge. But we can see that it will result in concurrently viable signals on two conflicting directional neuromimes. In fact, if both whisker's were to contribute equally to the learning, both of these signals could be maximally learned. If this occurs, the competition between the directional neuromimes would determine the motor response on recall trials.

Of the five pairs of whiskers, only one pair (the rear two whiskers) has matched responses. The remaining four learning networks can (ideally) learn a pain response that reflects one or both of the associated whisker responses. The learned response, at any given time, will be dependent upon the relative use of each whisker in learning, and the number and nature of inappropriate learning encounters which have occurred in the near past.

Actually, the above paragraph is not the whole story, as indicated earlier in this chapter. It would be true for an isolated learning network, but in our case, the activities of other drives may affect both when learning occurs, and the availability of the learned response. The next diagrams focus once again upon drive interac




Figure 49. A sample of reduced learning and/or response due to inhibition of $S2_{TFL}$ by HDR. FDC is also being inhibited, but it is less fatigued. Fear is dominant.

tions, this time in the context of how another active drive (hunger) affects the learning and recall processes.

Figure 49 illustrates the activities of hunger, fear, and learning neuromimes during active competition between the drives. This diagram is a fairly representative of both learning and recall trials, where the hunger drive has been active long




Figure 50. A sample of competition between hunger and fear, where hunger is regaining dominance. Learning and/or the learned response is inhibited.

enough to be fatigued. We can see that the fear drive is dominant, and that the sum of inhibition (from HDR) and fatigue is taking a toll on the second sampling cell (S2_{TFL}).

If this was a learning trial, SENSI would still be executing its reflex response, but the learning of that response by the second sampling cell would be

significantly reduced. In a recall trial, figure 49 corresponds behaviorally to a reduced learned response.

Figure 50 illustrates a trial in which fear had fatigued, and hunger was regaining dominance. FDC and $S2_{TFL}$ are unable to fire due to fatigue and increased inhibition from HDR, and HDC is firing more often as the inhibition from FDR lessens. When this occurs during learning, the learning of the motor response will not occur. During a recall trial, the learned motor response to fear will be inhibited, and the hunger response of movement toward light will be the dominant behavior.

Behavioral Analysis

After the necessary number of learning trials, the learning network results in an obvious improvement in obstacle avoiding behavior. SENSI's responses to touch when afraid are readily observable; it tends to stop upon touch and then move away from potential collisions. If the robot drives itself into a corner and learns inappropriate touch responses, these can be seen to persist until enough relearning extinguishes them. Inappropriate responses are observable as learned movement into the object touched.

SENSI's behavior prior to the addition of reinforcement was very disappointing. In order to accommodate reasonable adaptivity, the synapse decay rate was set only slightly lower than the learning rate. Learned responses could readily change, but they were also quickly lost when used. Once suprathreshold levels were reached on the plastic synapses, the learned response would occur. This weaker-than-maximum learning would decay by use to subthreshold levels within about three recall trials. The learned touch response is a much slower motion than the reflexes, so there is less motion away from the obstacle per touch. This means multiple touches often occur before an obstacle is cleared. What usually happened was several adjusting movements away from an obstacle terminating in a painful encounter due to loss of the learned response.

The reinforcement network added to counter this problem permits SENSI to complete adjustment movements (as long as subliminal fear persists) so that it resumes forward motion without experiencing pain. (Pfeifer & Vershure, 1992) utilized post-synaptic based reinforcement within Lola's network which also served to counteract active decay.

Unfortunately, SENSI still will exhibit several touch responses only to finally experience pain. Sometimes this is due to a bit too vigorous of a touch response, but

the biggest culprit is the eventual fatigue of fear. When fear fatigues, boredom becomes dominant, and SENSI moves straight ahead. It must experience pain or startle again before the learned responses can become available.

Several options exist for 'fixing' the loss of fear. One approach would be to develop a network which could learn the conditioned emotional response of aversion to touch. Rather than learn to be afraid of touch, however, ultimately we'd like SENSI to learn a response to touch given boredom. A discussion of how such learning could be implemented in our network is given in the next chapter as an example of future directions to be taken.

The adaptivity/permanence dilemma has been satisfactorily resolved with the active reinforcement circuitry. We observe continued learned responses in the continued presence of touch, and active loss of inappropriately learned responses, as desired. In the next section, we analyze why the dilemma exists.

Plasticity and Permanence

Our experimentation with learning hardware was interesting to say the least. We modeled our plastic synapse after equations given by Grossberg, and expected to achieve synapses which "learn a time average" of signals present upon the sampled cell, yet, which can also "learn and later perform... without destroying the memory" (Grossberg, 1982). The dilemma: the equation which learns a time average seems, by its very nature, to discourage self-sustained or permanent synapse strength.

This equation, from (Grossberg & Levine, 1988) is of the form

$$dz/dt = -Az + f_1(x)f_2(y) \text{ or } dz/dt = [-Az + f_2(y)]f_1(x)$$

where z is the synaptic efficacy or associative strength,

x, y are the correlated cell activities

f_1, f_2 are monotone, nondecreasing, non-negative signal functions, and

A is a (slow) decay function.

Our model uses a step function for f_1 (it's a bit more complex since the threshold varies with fatigue) where activity below the threshold of the presynaptic cell yields $f_1(x) = 0$, and suprathreshold activity yields $f_1(x) = 1$. Given that the presynaptic cell is active and sampling, which is the only time the strength, z , changes, both forms of the equation become

$$dz/dt = -Az + f_2(y).$$

At the postsynaptic cell, we sample some part or all of the excitatory input strength,

and we use $f_2(y) = By$, where B is the learning rate. In our model, $B > A$.

Now suppose the sampling cell samples its own input to the postsynaptic cell. If no other input is present at the time of sampling (recall) we have

$$dz/dt = -Az + Bz = (B - A)z$$

and since this rate of change is positive, the associative strength is permanent and will increase to its limit. This is the self-sustained learning we speak of.

If the synapse does *not* sample its own input to the postsynaptic cell, and no other input is present at the time of sampling, we obtain

$$dz/dt = -Az,$$

the active decay which permits adaptability (or the learning of a time average of presented signals), but results in learning loss with use.

What our reinforcement does, quite simply, is add an input to the postsynaptic cell, during recall trials (no pain) where the synapse has learned enough strength to fire the cell. This input is then a sample-able excitation to the postsynaptic cell, which permits the synapse to maintain its learned strength.

Had we chosen to utilize self-sustained learning, the only way we could see of reducing the amount of learning, once learned, would be to counter it with inhibition where the (net) input was being sampled. If an excitatory input had been learned, but was summed with some inhibitory input at the sampling point, then the overall excitation sampled would be less. Our synapse model would react to that by incrementally decreasing its learned response strength. With continued inhibition, the synapse could gradually lose its learned strength, in much the same way as it gradually learns it now.

Recall that our synapse was initially intended to sample net input. In our neuromime model, however, fatigue is not separable from the sum of excitatory and inhibitory inputs, because it is added into the inhibitory inputs before net input summation (see Appendix A). This created a real problem in sampling the motoneuromimes' activities, and we ended up sampling net *excitatory* input. Were it not for the problem with fatigue, we could have obtained the desired permanence and flexibility on the motodirectional interneuromimes by sampling their net inputs. These neuromimes actively inhibit each other, so the excitation for learning and the inhibition for relearning are already in place. For the motor speed interneuromimes, the needed inhibition is not currently present.

What about cases like FDR, where the input is multi-modal, and single mode sampling is required? We already noted that in these cases, the sampled input and the input from the synapse need to be combined within a dendritic tree. Then to

sample net input for that modality, we need a dendritic tree model which can sum both excitation and inhibition. Though not presently used that way, our model has that capability. As in the case of motor speed, the question becomes: what is the source of inhibition?

Although we dealt with the adaptability/stability dilemma by permitting active decay, but implementing active reinforcement, there are unanswered questions to this approach also. It implies that we can reinforce 'good' learning, while permitting 'bad' learning to fade away naturally. It then follows that there must be a way of evaluating the goodness or badness, *within the network*, of a particular learned response. We could readily identify the goodness or badness in this learning instance, but it is not likely to be so obvious in other learning situations, nor is it desirable to hardwire such reinforcement. In general, the source of positive or negative reinforcement will be based upon the expected outcome of the action in question, and this expectancy will probably also have been learned from experience. In such cases, it may be either the expectancy or the motor action which is in error. This issue is clearly one which must be addressed within a more sophisticated network.

Summary and Utility of Results

As with the innate network, the experimentation with learning has illustrated architectural aspects which are generalizable to other research efforts in networks and learning. In particular, one form of attentional focus is found to be a natural property of the network structure. A design principle which aids scalability has been determined, and should prove useful for planned expandability. Finally, questions raised by the effects of specific learning connections indicate issues which require further research.

SENSI's learning network follows (Grossberg, 1987a) in that all learning is associated with some (innate) drive. Learning associated with a drive only occurs, and can only be used, when that drive is dominant. This provides a focus to learning, one which is similar to what (Foner & Maes, 1994) refer to as cognitive selectivity. Foner and Maes discuss the use of focus of attention to improve unsupervised learning, and describe two methods of attaining focus: perceptual selectivity and cognitive selectivity. Cognitive selectivity is some method which limits attended internal structures; drive networks and competition provide this function in SENSI's network.

Perceptual selectivity, or limiting the sensory stimuli attended at any given time, has been included (by necessity) in our work. Foner and Maes indicate that it puts limits on what might be learnable, and point out that it can be spatial or temporal. We have found this limitation useful in reducing inappropriate learning, as described in the previous chapter. Our implementation of sensory competition (among the touch inputs) limits active touch by both spatial and temporal criteria. One or two touch stimuli may be active at any given time; the most recent touch and the preceding touch may be simultaneously active if they are on adjacent sensors. In SENSI's network, touch itself competes, thus sensory input is limited, which may be different from limiting *attended* sensory input.

Foner and Maes also indicate that the current 'goal' could be used to help select relevant facts, and they express a desire to incorporate a mechanism for learning 'what to pay attention to.' Although we did not implement the network components which could have achieved these capabilities in SENSI, they are included in networks presented by Grossberg, in particular, (Grossberg, 1987a, Chapter 1), which describes plasticity from drive to sensory events via feedback connections.

The addition of plasticity from drive to sensory neuromimes may provide yet another benefit, occurring naturally in classifier systems with genetic algorithms (Holland, 1975), but not within SENSI. This benefit is the capability to generalize situational, or sensory elements, during learning. The sampling cell structure utilized in SENSI's network naturally generalizes learned actions by means of maintaining only those synapse strengths leading to neuromimes that are consistently active together with the sampling cell. Plastic feedback from drives to sensory neuromimes could differentially strengthen relevant sensory input (by means of the same consistency in activity), resulting in both sensory focus and generalizability.

Scalability -- the ability to expand small networks with new sensory or behavioral capabilities *without loss of functionality* -- has become an important issue in recent research (Cliff, et al., 1994). Emphasis on maintaining coherence in expandable systems is given in (Brooks, 1994), although it is unclear where learned behaviors fit into his proposed architecture. In SENSI, learning is the method of choice for expanding initial system capability. (Aitken, 1994), states that modularity and unsupervised learning are essential to enable neural networks to scale to larger applications.

SENSI's network incorporates both unsupervised learning and a great deal of modularity. Additional learning networks, in particular, could be quickly added

between sensory and drive components by following the connection example given for touch learning in the previous chapter. Added learning would not disturb existing functionality because learned behaviors in SENSI are automatically included in the competitive structure responsible for determining dynamic behavior. The manner in which learning occurs with the context of a drive also results in learning, itself, functioning as a behavior. In other words, the activation of learning (or its utilization, as they are equivalent) depends upon the same criteria and competitive interactions as the innate behaviors. We feel that any scalable system must likewise incorporate some mechanism for automatic mediation among innate and learned (learning) behaviors.

In SENSI, automatic mediation is achieved by the use of dendritic tree structures to limit the application of learned excitation, and by mutual inhibition among innate and learned behaviors. Other recent research on learning does not directly address the limitation/normalization issue which is raised as viable connections increase in number, but do utilize competition. (Millán, 1994) implements dynamic module generation in his interesting simulation of learning. Competition occurs among existing modules, and therefore incorporates learned modules as they appear. (Gaussier & Zrehen, 1994) uses winner-take-all groups, with added noise to reduce the likelihood of equivalent excitation. Choice is made via a modified Maes network in (Gizster, 1994). (Aitken, 1994) focuses upon modules for sensory, motor, and association functions, and the architecture is designed specifically for modularity. Competition for representational space is mentioned in this work, but the discussion revolves around a single sensory and motor module, leaving issues of mediation among greater numbers of modules for future work. Interestingly, none of these approaches that focus on learning utilize the fatigue we found so necessary for effective competition. Its inclusion in SENSI's innate network naturally carries over to the competition involving learned and learning behaviors. We feel that this mediation issue must be addressed *especially* in systems which evolve in complexity by means of learning.

Also not found in the literature is any mention of the permanence/adaptability dilemma described earlier in this chapter. Weight changes in Millán's work are permanent, and determined by reinforcement. Gaussier and Zrehen use Hebbian learning modulated by pain in their sense-think-act-evaluate-adapt cycle. (Bersini, 1994) modulates weights based upon maintaining endogenous variables within some pre-determined 'viability zone'. These efforts seem to be concerned with procedural learning, which we believe should be continuously adaptive. Our use of reinforce-

ment serves to maintain useful learned behaviors in SENSI's network, and was needed due to the implementation of active decay which permits continuous adaptivity.

The approach closest to ours is given by Gaussier and Zrehen, who have implemented both positive and negative reinforcement signals based upon pain and pleasure. In their system, strengths influencing competition increase slightly if an action neither starts nor ends with pain, increase significantly if an action begins with pain present and ends with it absent, and decrease if pain is active at the end of an action. Theirs is a strategy of interest to us because active, nonspecific decay is replaced with active, specific, negative reinforcement. However, the approach is better suited to their simulation cycle than our dynamic network because it relies upon discrete actions.

It is our hope, as biologically-based networks become more prevalent, that the adaptability/permanence dilemma we encountered will be further addressed. Is this a biological phenomenon? Do biological connectivity differences account for differences in procedural and declarative learning? Does the connection locus of plastic synapses play as important a role in natural systems as it has in SENSI's network? The results presented here offer interesting connection variabilities and effects which deserve further investigation.

CHAPTER VIII

CONCLUSIONS AND FUTURE DIRECTIONS

The greater our knowledge increases the more our ignorance unfolds.

John Fitzgerald Kennedy

One never notices what has been done; one can only see what remains to be done....

Marie Curie

The work done with SENSI has reaffirmed aspects of current theories regarding how behavioral choices are made in a neural network. It has also illustrated ways in which learning and the choice mechanism affect each other. Several components utilized in this network offer alternate strategies to the plethora of competition-based choice mechanisms currently being investigated. Numerous questions, primarily in the area of learning networks, were raised by this research. This chapter begins with a brief review of the research goals, summarizes our results, conclusions, and their applicability, and expresses unresolved issues and potential enhancements for future investigation.

Review of Research Goals

This project began with the primary goal of investigating the nature of neuronal connections that result in behavioral choice. The intent was to explore a network structure which embodied available behaviors, a method for dynamically choosing among them, and a manner of attaching available behaviors to different environmental conditions, through learning, without violation of the choice mechanism. By building this network entirely from electronic components, we further sought to investigate these structures within an inherently real-time and dynamic network. This network was to function as the sole controller of an autonomous

robot, and was to be reactive, directed, and adaptive.

Accomplishments and Conclusions

The current result of our research is an electronic network which almost meets the above goals. It consists, as illustrated in Appendix E, of 87 neuromimes and 342 synapses. There are 297 fixed strength synapses, of which 183 are excitatory and 114 are inhibitory. All of the 45 plastic synapses are excitatory. Of the neuromimes, 30 are sensory and 19 serve to regulate motor function. Thus a little over one half of the network is devoted to sensorimotor activity. 19 neuromimes form the basis of the drive networks and the innate behaviors. The remaining 19 neuromimes are the sampling and reinforcement cells of the learning network.

SENSI could be thought of as only having variations on a single behavior: that of motion. If we consider as distinct those directional and speed variations available to SENSI, our network embodies 13 innate behaviors. 10 of these are reflex responses and 4 are associated with the boredom and hunger drives. An additional 5 behaviors may be acquired by the learning network, but since they are learned from the innate behaviors, they are not necessarily distinct activities.

The choice mechanism utilized in SENSI's network is dynamic competition. It usually occurs among the drive networks, but competition also exists at the motodirectional and sensory interfaces. Learned behaviors exist in the context of an innate drive, so the choice between a learned and an innate behavior is also made through the drive interactions. (If both the learned and innate behavior are in the context of the same drive, the choice falls to the sensory and motoneuromime competitions.)

We have therefore succeeded in creating an electronic network embodying innate behaviors, learned behaviors, and a method of choice effective over both behavioral sources. This network is the 'nervous system' of our robot SENSI. SENSI is reactive due to reflex behaviors, directed through the maintainable activities of the drives, and adaptive in that it can learn to apply behaviors to circumstances other than innately provided for.

Despite our desire for SENSI to be autonomous, it is clear that we cannot claim it. In some senses, it has the potential to be autonomous. It is battery powered and supposedly capable of recharging its batteries given an intense source of light. However, with the power requirements of the complete network, it doesn't have enough solar cell area to both run itself and charge. Also, it isn't too capable of

survival given its behavioral repertoire. Its sensory limitations prevent it from avoiding a sudden demise in all but the most friendly of environments. We consider it to approach the autonomy of an amoeba: in its intended environment, it can move about, away from noxious stimuli once it has sensed them (for SENSI, by contact), and toward perceived sources of food when hungry. Outside of its intended environment, (and indeed in it, due to its charging limitations) extended survival is unlikely.

In the process of building SENSI's network, we learned a few things about neuron models. We discovered facets of our neuromime and network design which are required to obtain dynamic competition. We ascertained the placement of learning networks which results in learned behaviors (automatically) becoming a part of future competitions among relevant behaviors. And, we found features of plastic synapses and learning networks which raised questions about the nature of long term memory. Each of these findings will be discussed next.

Neuron models. SENSI's neuromimes are designed to emit square pulses when active, as opposed to the spike and decay activity of biological neurons. This on-or-off feature of our model has created problems for us in at least two ways. The problem is rooted in the all-or-none influence a presynaptic neuromime has upon a postsynaptic neuromime.

When the presynaptic neuromime's influence is inhibitory, it is capable of inhibiting the postsynaptic neuromime only when the timing of incoming inhibitory pulses results in complete overlap of incoming excitatory pulses. (It is necessary to assume a single source of excitation, otherwise the stronger excitation should be able to overcome the inhibition.) At any rate, we found numerous instances of opportunistic postsynaptic neuromime firing, which occurred due to some portion of an excitatory pulse exciting the postsynaptic neuromime between inhibitory pulses.

We believe that part of this problem is due to the reduced density of neuromimes in SENSI's network; the 'one-for-many' simplification does not permit the statistical overlap of inhibitory and excitatory pulses to result in limited and harmless opportunistic firing. However, even within such a simple network, we eliminated the problem by the insertion of a capacitor in the inhibitory synapse. This component more closely mimics the biological signal decay, resulting in a stretched inhibitory pulse which fades from full to zero intensity.

The second area in which all-or-none influence failed us was in positive feedback. The feedback in our drive networks allows an active drive to remain

active, despite the loss of the stimulus which initially excited it, as long as nothing actively inhibits it. This positive feedback must overlap with an existing sensory input for drive network activity to continue (see Chapter IV). We found that the feedback input also required a capacitor stretch to ensure reasonable overlap with incoming sensory pulses.

As networks become more complex, it is possible that the opportunistic firing we have suppressed will become important (Softky, 1995). Complex networks may be dense and repetitive enough that the signal shape becomes less important than the statistical properties of overlap. However, for simple, exploratory networks, where each connection serves a (possibly unique) functional purpose, signal shape greatly affects network stability. Future work on SENSI's network will incorporate stretch on all neuromime outputs.

We have found that fatigue should be a parameter of all neuromimes. Initially we did not include it for some sensory cells, in consideration of SENSI's limited sensory input sources, and the concept that the network needed to be aware of whatever input it could. This is not valid. Fatigue does not prevent a neuromime from firing, it just reduces its firing rate. As will be reviewed in the following subsection on competition, fatigue is a crucial ingredient of neuromime and network interactions.

A final conclusion regarding neuron models is a need for separation of fatigue and inhibition. Although each has a similar effect on the neuromime, namely, to raise its threshold so that greater levels of excitation are required to fire the neuromime, our model grouped these inputs (external inhibition and internal fatigue) together before comparing them to existing excitation. There are circumstances where the inhibitory and fatigue inputs need to be available separately. For example, in learning, we have concluded a need to totally separate fatigue from the inhibitory input, so the inhibition present on a neuromime may be sampled independently of the level of fatigue.

Competition ingredients. Because effective competition among neuromimes and neuromime networks was such a critical part of this research, we reached several conclusions with respect to the nature of competition and its necessary components. Regarding the nature of competition, we found a limit condition which helped us ensure that valid competition could occur at all. This limit condition relates to the worst possible scenario for competing network elements: the existence of continuing, comparable, excitatory input to the competing entities.

Under these conditions, we conclude that the competing entities should 'toggle' in their activation. (Comparable input implies not necessarily equivalent input, but equivalence of input and internal bias. In other words, higher priority entities need less excitatory input than lower priority entities in order to have comparable input with respect to activation.)

With this as our guiding criterion, we attained effective competition at the neuromime level with the implementation of three additional constraints. First, it must be possible for a competing neuromime to inhibit its competitors. This implies no neuromime involved in competition can be capable of firing regardless of extant inhibition and fatigue. Second, for neuromime-neuromime competition, we found it necessary to make this inhibition complete, namely, an unfatigued inhibiting neuromime should be able to prevent any opportunistic firing in the inhibited neuromimes (see the discussion on inhibition and opportunistic firing in the previous subsection). Third, and most important, fatigue must be a feature of the competing neuromimes. The buildup of fatigue in the dominant neuromime is required to lessen its inhibitory effects on the suppressed neuromimes. When this occurs, the suppressed neuromimes begin to opportunistically fire, inhibiting the dominant, fatigued neuromime (which further limits its inhibitory effect on the competition). When a different neuromime becomes dominant, the toggling requisite of competition is fulfilled. As might be expected, the rates of fatigue and recovery from fatigue are critical to the behaviors of competing neuromimes. We found slower fatigue and/or quicker recovery to be a viable method of introducing bias into our network.

The constraints on effective competition among networks are similar to those for competition among neuromimes. The only network competition we have worked with so far is among the drives, so this discussion relates directly to the drive networks (refer to Chapter IV). The limiting condition we used for neuromimes is applied to drives: we expect continuing, comparable, excitatory input to result in toggling of drive activation. For this to occur, we found the effects of both fatigue and feedback to be critical.

A drive network produces internal, positive feedback only when it is completely active (fires at the final, competition, stage). The first condition we found is the existence of this feedback must affect the activity level of the drive. When a drive is dominant, it will produce regular feedback signals; these signals must serve to increase the drive's activity level. In other words, the drive neuromime must fire more frequently with feedback than without it. This constraint results in two additional conditions relating to individual drive neuromimes.

One is that fatigue buildup in the drive neuromime (the first stage of the drive network) under non-feedback conditions, must be surmountable by feedback. However long the drive neuromime may have been active before the drive network produces feedback to it, the additional input must positively affect its activity level for the first condition to hold. Naturally, if the first stage of the drive network cannot increase its activity, the rest of the network will have no basis to do so.

The other condition relates to the drive representation neuromimes of each drive network. These neuromimes must not be fatigued under normal operating conditions. If they become fatigued without the drive ever having been dominant, then they cannot increase activity when feedback does occur, and the drive neuromime output increases in frequency. (This is because increased overlap is required to overcome fatigue; increased frequency of input due to increased activation of the drive neuromime would have no effect on the network if the representation neuromime had already fatigued.) Again, our first condition would be violated.

The last requirement for our drive competition is that fatigue buildup eventually occur in some portion of the drive network. The increased activity which occurs with feedback marks the dominant period of the drive in question. While it is dominant, it inhibits other drives' competition neuromimes so that no other drive receives positive internal feedback. As with neuromime competition, fatigue buildup in the dominant drive network reduces the overall inhibition of other drives, allowing each the chance to opportunistically produce feedback. The resulting increase in activation which follows feedback then permits some other drive to gain dominance.

Given our neuron model and network structure, the choice mechanism of dynamic competition requires the existence of fatigue in the neural model. Also, in the drive structure proposed by Grossberg and implemented here, positive internal feedback is necessary not only to provide continuity in drive-related behaviors (robot directedness), but also for the effective implementation of dynamic behavioral choice.

There are numerous implementations which employ competition to mediate among behaviors. What we contribute is a limit condition applicable to any such model, and examples of meeting this condition in a continuous, dynamic neural network implementation. The limit condition of behavioral toggling given comparably desirable actions results in systems which balance persistence with opportunism (Blumberg, 1994), a goal of any system which must mediate among conflicting behaviors to attain its objective(s).

Integration of Learning. From the diagrams of figure 41 and 42, and the

related discussion, we determined the placement of learning networks within the competition framework. The placement results in the automatic inclusion of learning and utilization of learned behaviors in the innate competition framework. This aspect of network design is critical for systems which learn, so that learning does not disrupt innate competence.

The relation of learning networks to specific drive networks provides the automatic mediation. It also results in cognitive selectivity (Foner & Maes, 1994): a focus within the system that limits attention to internal structures. Foner and Maes advocate the use of both cognitive and perceptual selectivity to improve unsupervised learning. We have found sensory competition, which provides perceptual selectivity, important for both behavioral choice and effective learning. Competition among sensory events limits the scope of learning and recall, by limiting the activity of networks (both innate and learned) to those excited by the most relevant sensory inputs. Sensory competition is also useful for determining the current behavior when more than one behavior is possible for the dominant drive.

Two other discoveries were made about proper integration of learning into the innate network. The first, with respect to the plastic synapses, is that the input sampled at the postsynaptic cell is not necessarily the net input to that cell. Especially in cases of cells like the drive representation neuromimes, where excitatory inputs come from both external and internal sources, our model requires that the synapse be connected so that it samples only the source that its learned input is intended to 'replace'. If this is not done, the synapse can learn enough excitation from the 'wrong' source to cause the postsynaptic cell to fire during recall trials, when it had never done so during training.

The other discovery has to do with the plastic synapse's output to the postsynaptic neuromime. When only one of the available sources of input to the postsynaptic cell is sampled (for neuromimes with more than one source), the learned synaptic output and the source of the learning must be joined within a dendritic tree. Recall that dendritic trees can serve the purpose of limiting the total excitation presented to the destination neuromime. In this case, we need this limitation so that the learned input cannot combine with the sampled source to provide a suprathreshold excitatory input to the postsynaptic cell. Without this restriction, learning can again result in improper activation of the postsynaptic neuromime.

The above two conclusions regarding plastic synapse connectivity are applicable primarily to biologically-oriented network models. For such models, they illustrate the care that must be taken to clarify the intent of plasticity.

The Nature of Learning. Our research has shown us that the nature of learning is dependent upon the structure of the sampling cell's connections to the sampled cell. Basically, the sampling cell can learn from excitation present on the sampled cell, which either includes or excludes the output from the sampling cell itself. Inclusion of the sampling cell's output results in self-sustained learning; exclusion results in learning which exhibits active decay when used.

We have determined several approaches to dealing with the dilemma of permanence versus flexibility. One is to develop the network with self-sustained learning, but incorporate the necessary elements to also respond to the effects of inhibition if the response becomes inappropriate. Another is to develop reinforcement sub-networks, which serve to positively reinforce appropriate learning while allowing inappropriate learning to actively decay. It is possible that both of these approaches exist in animate networks; we feel that the former is more applicable to declarative learning, and the latter to procedural learning.

Areas for Continued Research

There are two main areas in which we'd like to continue this work: enhancements to the current network, and elaboration of the current network. Along the lines of enhancement, there are changes we'd like to make to the neuromime model, the synapse model, and the learning circuitry. We need to upgrade our neuromime model to separate fatigue and inhibition, if only for the purpose of optimizing learning for the directional neuromimes. (However, as they are no doubt the result of distinct mechanisms in animate neurons, we'd probably encounter other reasons for separating their effects on our neuromimes.) We'd also like to implement 'spike and decay' output pulses throughout the network, since there are at least two areas already where such an output shape is beneficial to overall behavior.

An Alternative Neuromime Model. Before the whole network is altered at the neuromime level, however, there is a distinct neuromime model to be investigated. The model we'd like to take a closer look at is one which utilizes a bi-polar power supply. With bi-polar power available, we could use negative voltages for inhibitory inputs, and positive voltages for excitatory inputs. One potential benefit to such a model is the ease with which net input learning could be implemented (whether excitatory or inhibitory) via each plastic synapse.

According to (Kandel, 1979), whether a synaptic connection is excitatory or

inhibitory is dependent upon the receptor at the synapse. Thus, it is feasible to have a sampling cell learn connections which will ultimately be excitatory on some neuromimes and inhibitory on others. SENSI's neuromimes are currently modeled in such a way as to make this possible. The type of action resulting from a synaptic input depends on its location on the neuromime. Unfortunately, there is no way of predetermining whether or not a learned connection should be inhibitory or excitatory. Therefore we don't know how to attach the plastic synapses *output* to the sampled cell, so that the learned (excitatory or inhibitory) result occurs.

Even with the changes suggested in the previous chapter, namely, the implementation of inhibitory sources by which learned responses could be unlearned, we still cannot, via the synapse, learn to inhibit a sampled neuromime. With our model as it stands, the only way a sampling cell can result in either excitation or inhibition of the sampled cell, is by locating a synapse on each input location (excitatory and inhibitory) of the sampled cell. This isn't a problem, but it at least doubles the number of synapses which must be available.

With different voltage polarities representing inhibition and excitation, any synapse, anywhere, could learn to provide whatever net input it was exposed to during learning. Any polarity changes resulting from future network changes would eventually be reflected in the plastic synapse's output. This seems to be an elegant and minimal way to achieve permanence and flexibility in long term memory. Instead of the synapse receptor site determining the input polarity, the information would just be a part of the signal. When designing networks by hand, where connection strengths can change over time, but connection loci are fixed, this type of architecture could greatly simplify the design process.

Learning and Unlearning. Whatever neuromime model is implemented, there is a need to explore sources of inhibition and/or reinforcement which can be used to inhibit and/or cause unlearning of behaviors which become inappropriate. Our research network has greatly increased in inhibitory connections during its development. However, they exist primarily in the innate competitions defined between sensory inputs, innate drives and motor actions. The only exception to this is the set of inhibitory connections added with the reinforcement networks. (Shepherd, 1990) indicates that 20-30% of all neurons in the cerebral cortex and thalamus are inhibition-releasing cells. "They are instrumental in defining and confining the response properties not only of single neurons, but also of large neuronal circuits." Sensory processing, in particular, appears to be dependent upon feed-forward inhibi-

tion for the enhancement of contrast which aids in sensitivity to changing stimulation.

We are more interested in applications of feedback (both inhibitory and excitatory) which may be related to reinforcement. (Routtenberg, 1979) speculates that "the pathways of brain reward may function as the pathways of memory consolidation." We have implemented positive motor reinforcement as described in the previous chapter, and it does serve as a form of memory consolidation. Unfortunately, our current knowledge regarding implementation of expectancy mechanisms is too incomplete to generalize that design.

When this research began, we were interested in the widespread effects of pain and pleasure systems. We'd like to investigate the literature on the topic more fully, to determine if the existence of widespread and subtle inhibitory and excitatory effects might be reasonably linked with our network, especially for the reinforcement or extinction of learned behaviors. Such a network presupposes the presence of an evaluative system made more complex by the probability that evaluations are also learnable and subject to reinforcement or extinction. The first step will be to use variables already present in the innate system to provide a simple evaluation of behaviors. (Gaussier & Zrehen, 1994) use pain and pleasure (defined as absence of pain) to positively and negatively affect weights to actions. The effect of behaviors on internal variables, combined with an evaluation based upon maintaining the values of those variables within a given range, is the basis of reinforcement in (Bersini, 1994). This elemental approach to reinforcement is promising for network models with unsupervised learning.

The capability to form expectancies and evaluate existing behaviors will play a role in future implementations of other types of learning, as reviewed in Chapter II. Habituation appears to be an extinction based on disconfirmed expectancy. The basic structure for habituation exists in our network as active decay. Edelman's theory of development by selection implies that flexible plastic connections capable of active decay can occur anywhere the initial network contains numerous feasible activity options, and experience is used to select those which are to remain viable. Our research indicates, then, that this experience may be used in forming reinforcing expectancies. This results in selection by reinforcement.

Grossberg proposes the initial existence of plastic connections between drive representations and cue representations. Initially, all such connections are viable and uniformly strong. He claims that "conditioning can change this uniform LTM pattern by differentially strengthening the LTM traces ...during learning trials". It

seems likely, since all of these connections are viable to begin with, that differential *weakening* of LTM traces is a possible selection mechanism here. Either way, this set of connections, not currently implemented, holds promise for attentional mechanisms and situational generalization, as discussed in Chapter VII.

Clearly, there is much to be learned about expectancy and reinforcement before the implementation of trial and error learning and cognitive map formation can even begin. Interestingly, however, our network contains the fundamental ingredients for latent learning, as described in the following section.

Latent learning. In Chapter II, we described latent learning as learning devoid of specific motivation, resulting in a training later applicable to motivational needs, and enabling an animal to selectively attend to a previously learned whole, without trial and error, in relation to current needs. We can obtain this form of learning in SENSI's current network, by adding only a few more cells, synapses, and connections.

SENSI could learn a latent response to touch given boredom, which then could be applied to avoid collisions even in the absence of subliminal fear. The added sampling cells would be placed somewhat differently, due to the differences between the fear and boredom networks. First, no cue is associated with boredom, so the first set of sampling cells is somewhat extraneous. (No needs exists to provide cue arousal to boredom.) The second set of sampling cells, excited by the first set and the drive representation neuromime in the fear network, would be excited by cue information and the BDR neuromime. The plastic synapses would be attached to the motodirectional neuromimes as in the fear network.

The anticipated result? Pain inhibits boredom, so any learning regarding touch during pain would be limited to the fear network. Touch occurring during bored motion, prior to any other learned response, would result in learning the 'normal' bored crawl activity. But after a response is learned for touch and fear, a whole new possibility arises. The fear-touch response inhibits BDC, but not BDR. Thus every fear-touch response (unaccompanied by pain, which does indirectly inhibit BDR) is learnable by the boredom-touch sampling cells. After an adequate number of successful fear-touch avoidance responses, touch during boredom would be capable of eliciting similar avoidance maneuvers! The requirement for the presence of subliminal fear in order to respond to touch would be lost (but the response would remain available if fear occurred).

This effect closely approaches latent learning because the response to touch

and boredom is learned without specific motivation (SENSI is not dominantly bored when it is learned), it is later applicable to motivational needs (it can be used when SENSI is dominantly bored), and it allows SENSI to use a previously learned response without trial-and-error. We would like to incorporate this learning, and if it generates the anticipated results, this learning structure could be incorporated into a drive related to exploration.

Grossberg's writings contain much more in the way of theoretical network structures for learning, expectancy, and reinforcement than we have been able to realize or try in our network. Some of this literature clearly points to other future expansions we'd like to implement within SENSI's network.

The primary area we need to expand is the sensory area. We have seen how dependent behavior is on what can be sensed and differentiated in the environment. We have also experienced firsthand the necessity of sensory focus, and its resulting effects on learning and behavioral choice. Finally, evidence indicates that sensory processing forms a large part of expectancy-based reinforcement. A large portion of Stephen Grossberg's theories deal with sensory input handling, including pattern matching and expectancy. We hope to investigate and implement these theories and explore the behavior of the resulting networks.

Within the limitations of our network, we have found basic principles which establish dynamic behavioral mediation. With these properties in place, it remains to expand the network, in any or all of its input, internal processing, and behavioral areas. We have created a testbed which will permit further behavioral robustness testing as network complexity increases.

APPENDICES

APPENDIX A

NEUROMIME MODEL

The neuron model, or neuromime, is based upon the operation of a 555 timer IC. The timer can be wired as a voltage-controlled oscillator and supports either monostable or astable operation. The schematic for our neuromime, with labels corresponding to those used below, is given in Figure 51. The 555 timer is wired to fire in astable mode. It is normally quiescent, but releases a pulse given a voltage applied to the not-reset pin. The resistors DCH and DCL, and the capacitor DCC control the baseline duty cycle of the timer's output.

Input to the neuromime is excitatory (I_{E0} , I_{E1} , I_E) or inhibitory (I_{I0} , I_{I1} , I_{IT} , FAT). FAT is a fatigue input, which builds as the neuromime is active. I_{IT} is a total-inhibit feature; a voltage applied to this input effectively negates any and all excitation, due to the lack of input resistance for this input. The effects (weights) of all inputs may be altered by changing the input resistors associated with them. IRC, IRD, and IRL form an inhibition stretching circuit, extending the effects of inhibition. Adjusting IRC and IRD affects the duration of the stretch.

THI, TLO, and INH provide the threshold for the neuromime. Since it is effectively presented as an inhibition input, this is the voltage level which the excitatory inputs must overcome for the neuromime to fire in the absence of other inhibition. The threshold is easily varied by changing the values of THI and TLO.

The input resistor network and 324 op-amp constitute a summing, offset differential amplifier. When the excitatory inputs are greater than the inhibitory inputs, a positive voltage is shown to the 555, causing it to produce an output. This voltage is also presented to the voltage control pin of the 555, via a simple stretching circuit, and serves to modulate the duty cycle of the output according to the analog

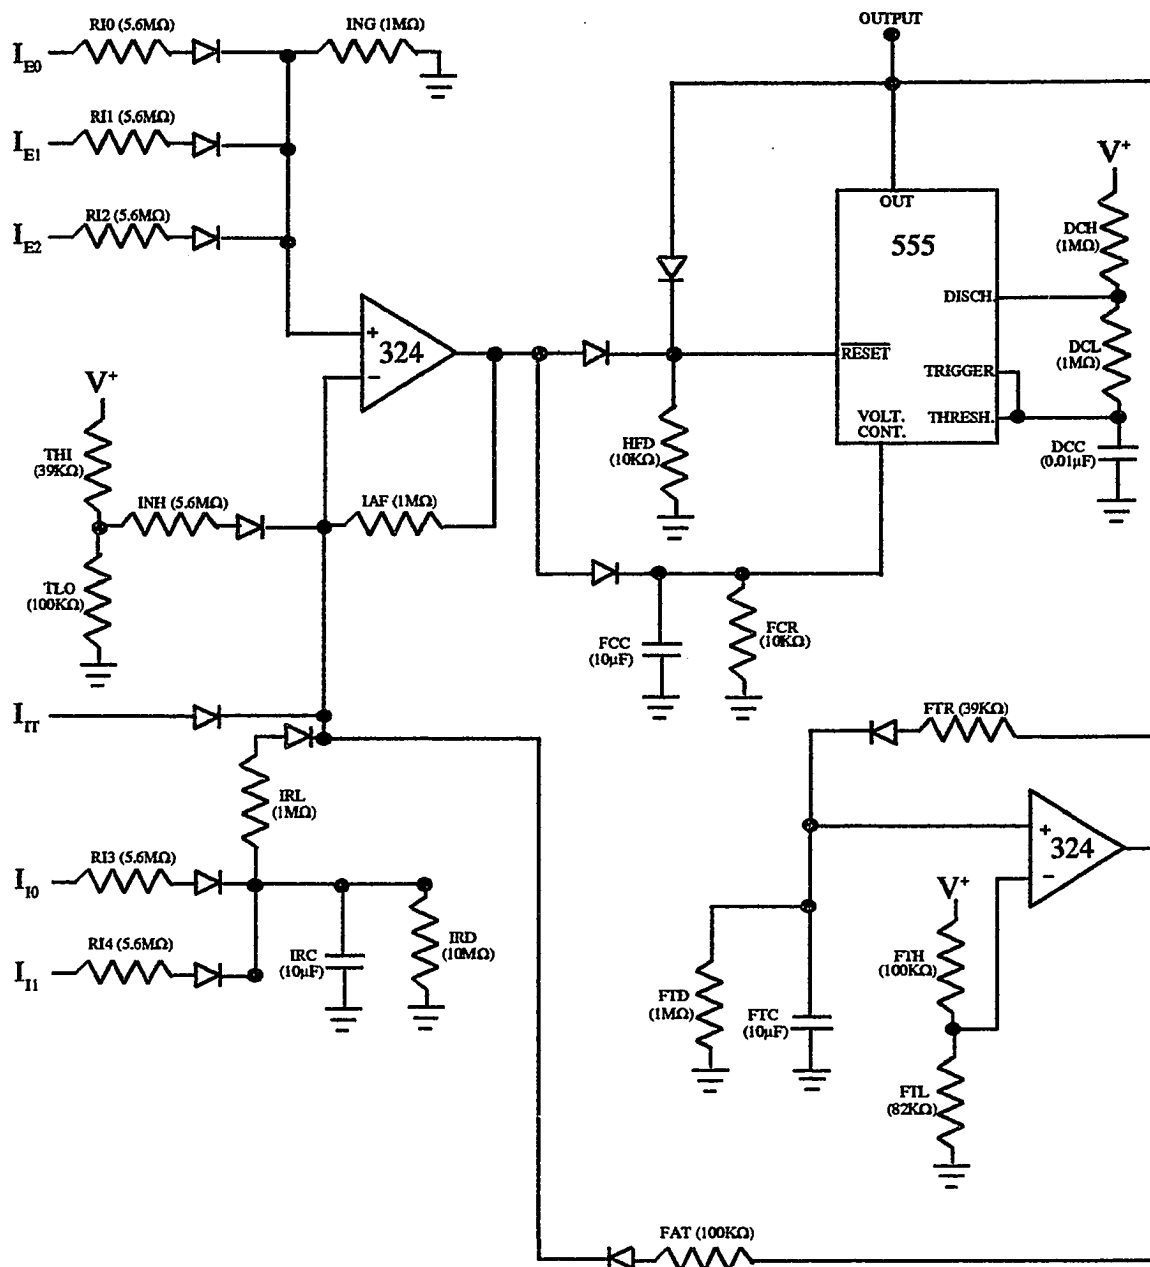


Figure 51. Schematic of the general neuromime, showing resistor and capacitor labels and their default values. Power connections and pin numbers (for the timer and operational amplifier ICs) are not shown. A description of the circuit is given in the text.

voltage present (the amount of relative excitation). This is the manner in which we translate intensity of input into frequency of firing.

Feedback from the 555's output, when active, is used to clamp the reset input in the high state to insure the generation of a full pulse when the input sum is very close to the reset threshold. HFD and the two diodes in its vicinity effectively eliminate high frequency oscillation characteristic of the 555 in this configuration. The output voltage is also presented to the fatigue subcircuit. The output builds up voltage in FTC, which, when higher than that given to the comparator by FTH and FTL, sends a fatigue (inhibitory) input to the circuit. The rates at which fatigue arises and decays can be varied by altering the values of FTR and FTD, respectively. The threshold for actual fatigue effects is set by FTH and FTL.

Wired as shown in Figure 51, the neuromime requires at least two active excitatory inputs in order to fire. Each of the inhibitory inputs cancels the effect of an excitatory input (their input weights are equivalent). The firing rate is approximately 22 Hz with a 50% duty cycle. Onset of fatigue reduces firing (with a maintained input) to 17 Hz with a 8% duty cycle.

The neuromimes are wired with 10% precision resistor components, so all measurements should be taken as approximations.

There are many possible variations to the general neuromime which alter the firing characteristics. Some variations and their effects are listed in Table 2. The component names in Table 1 follow the labeling given in Figure 51. Duty cycle and frequency data are the approximate values for normal excitation adequate to fire,

Table 2. Component values and characteristics of several neuromime types.

Neuron	FAT	DCH	DCL	DCC	FTR	FTH	FTL	FTD	FTC	DUTY	FREQ	Fatigue Onset
GENERAL	100K	1M	1M	0.01 μ F	39K	100K	82K	1M	10 μ F	50%	20 Hz	1 sec.
S (MOTOR)	100K	100K	1.5M	0.1 μ F	39K	100K	82K	470K	0.1 μ F	40%	20 Hz	immediate
FR "	1.5M	100K	3.6M	0.1 μ F	200K	100K	47K	470K	10 μ F	60%	45 Hz	5 sec.
FF "	1.5M	100K	10M	0.1 μ F	100K	100K	47K	470K	10 μ F	90%	10 Hz	< 1 sec.
I _p , I _r "	none	1M	1M	0.01 μ F	none	none	none	none	none	75%	40 Hz	never

in all cases except for the S neuromime. The S neuromime fatigues almost instantly and always runs that way; this is the only way for us to obtain a duty cycle less than 50% with our model. The S, FR, and FF neuromimes are the motoneuromimes described in Chapter IV. The I's are the motor interneuromimes, and the D's are the directional neuromimes. Note the higher duty cycle and frequency for the I's as compared to the general neuromime. This is due to the sustained high input they

receive from the pain neuromimes. These neuron types were developed to work together for the recoil motor reflex. Table 2 serves to illustrate both the variety of behaviors obtainable from the general model, and the tedious nature of the hand tuning required for the development of this type of network.

APPENDIX B

MOTOR CONTROL CIRCUITRY

The motor control interface, detailed in figure 52, consists of a power pulse delivery circuit which controls the speed for each motor, and directional control circuits which vary the polarity of current delivered to the motors. Two operational amplifiers provide high impedance to reduce loading on the motor speed neuromimes. The high output of the amplifiers drives the power transistors, sending V^+ raw battery power to the motor relays.

Motor direction pulses are stretched through the input capacitor-driven resistor network. Directional pulses activate the appropriate relays for motor polarity by way of LM339 comparator circuits which trigger from the start of the pulse until it decays to the threshold set by a voltage divider.

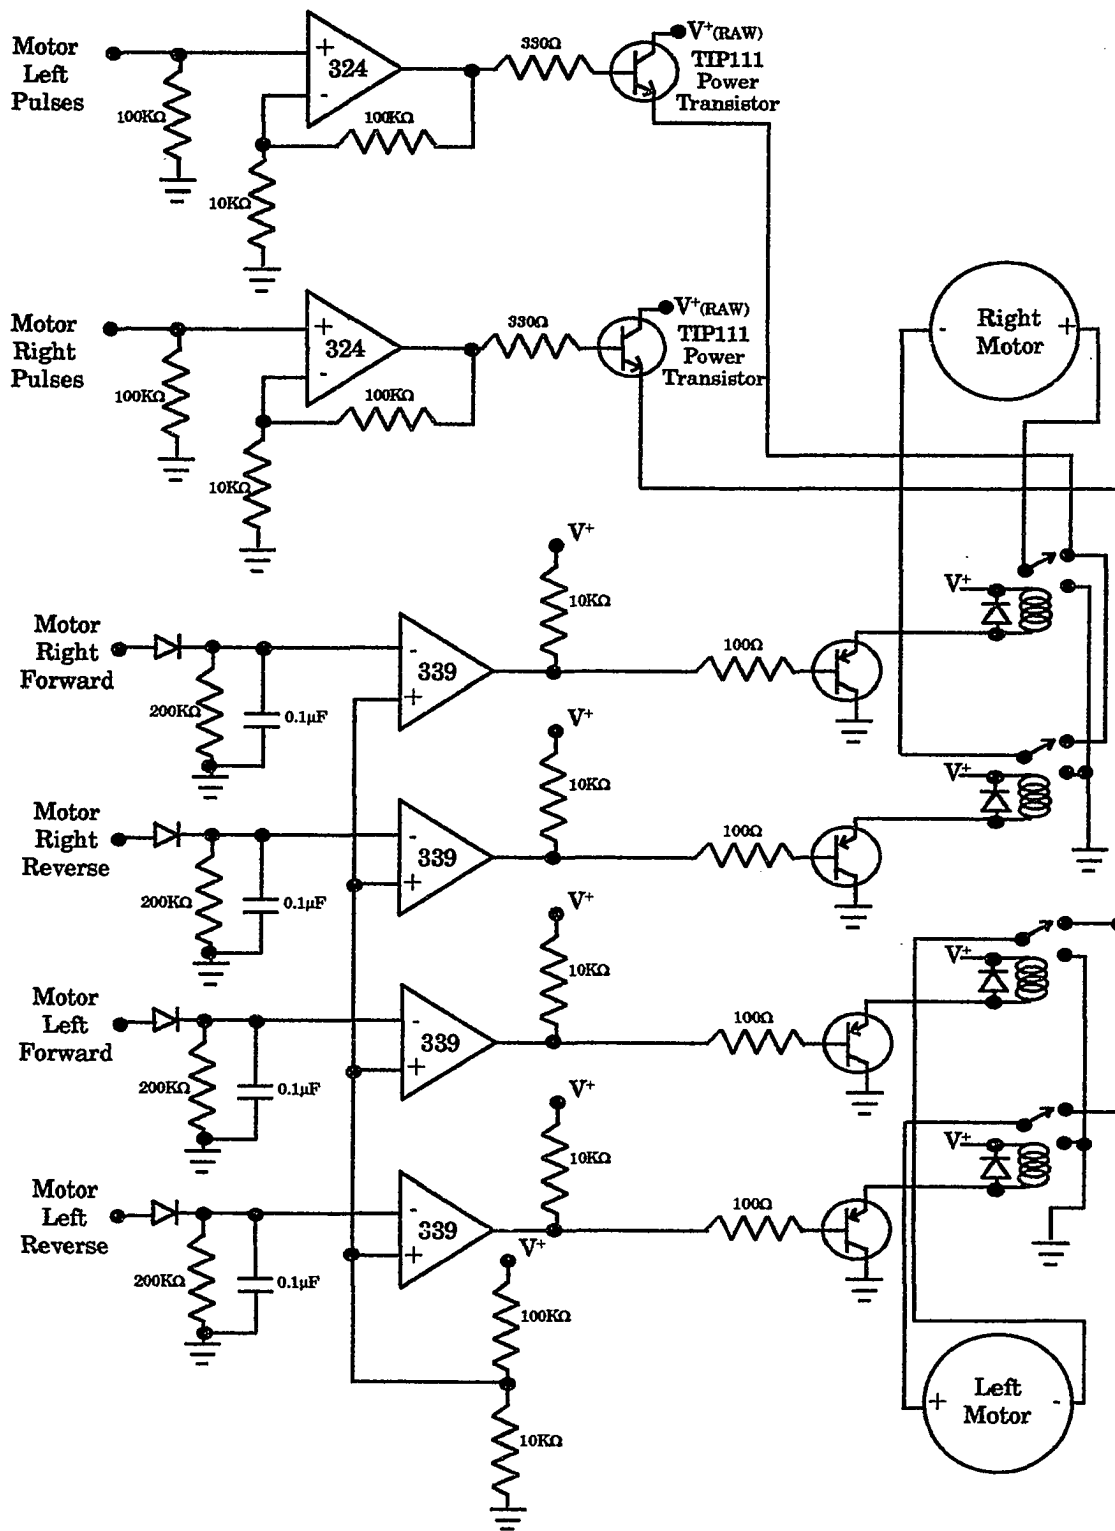


Figure 52. Schematic of motor control circuitry. All diodes are of type 1N914; all unlabeled transistors are type 2N3906.

APPENDIX C

SENSORY CIRCUITRY

Motion Sensor

The motion sensor, as drawn in Figure 4, produces separate signals for forward and backward motion. The signal frequency is proportional to the speed at which SENSI is moving. As SENSI moves, the motion sensor wheel rotates due to ground friction. (The entire assembly pivots on ball bearings to allow for maximum freedom of movement and rapid indication of motion direction.) This rotation causes the metallic bars to pass the photodetector; each pass registers as a pulse.

The circuitry which receives the sensor's signals and presents the motion signals is shown in figure 53 below. Reflected infrared from the LED produces a high level pulse which is converted to 5 volts by the 324 comparator. Forward or reverse is selected via the split ring and wiper. The 1000pF capacitor and diode result in a differentiated trigger pulse to the appropriate 555, which is wired in monostable, or one-shot mode. The result is one output pulse per wheel stripe.

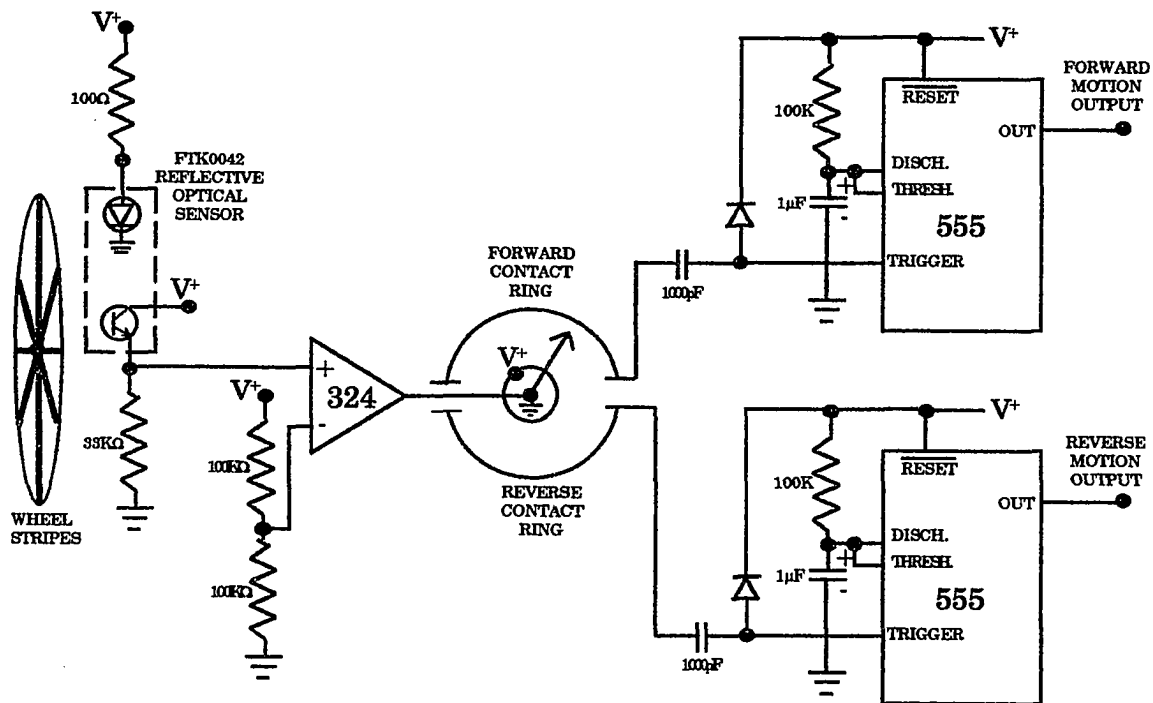


FIGURE 53. Motion sensor schematic. Refer to Figure 4 for actual sensor diagram.

Pain and Touch Sensors

The pain and touch sensors, or whisker assembly is diagrammed in figure 54. The circuitry producing the signal is a very simple implementation of the neuro-mime which receives its trigger signal from the whisker. When the whisker touches one of the two rings, it essentially closes a switch and presents a V^+ signal to the 555's trigger, shown in figure 54 as input. The small table in the figure lists the component values for the two types of neuromimes.

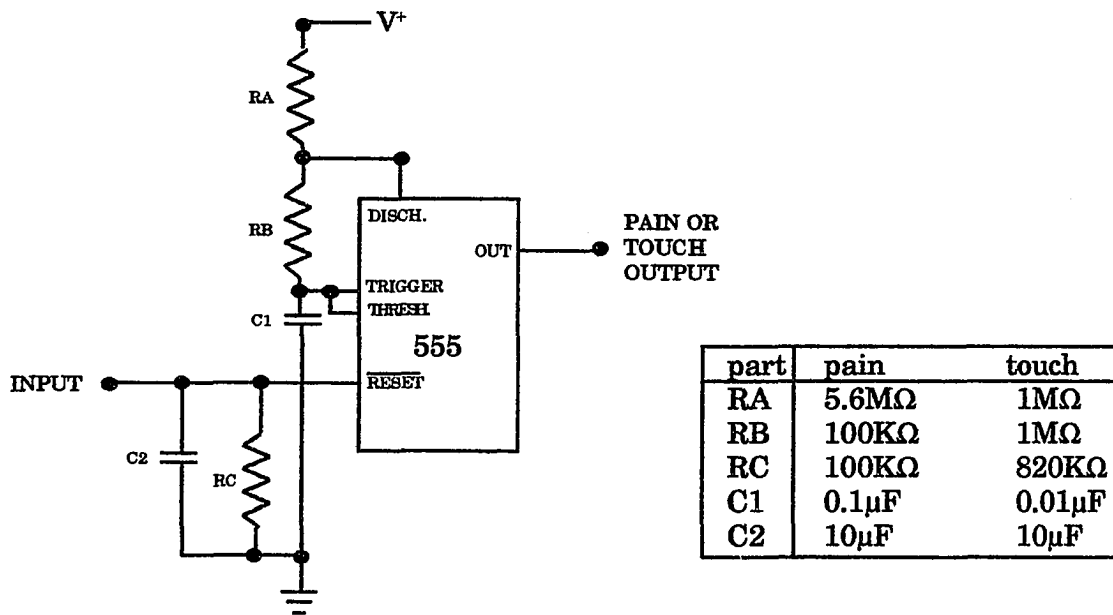


Figure 54. Schematic of the pain/touch neuromimes, with a table indicating the proper component values for each case.

Light Sensors

The photocell eye, or light sensor circuitry is shown in figure 55. Each 555 IC is the locus of a specialized neuromime producing one of the light sense outputs. To aid in explanation, the operational amplifiers have been subscripted.

Light detected by the photo transistor in the lensed eye housing of either eye sensor produces a small positive voltage which is amplified. This initial amplifier (A) has a variable gain to allow for sensitivity adjustment. The output of amplifier (A) provides the firing signal for both the sense neuromime and the light greater

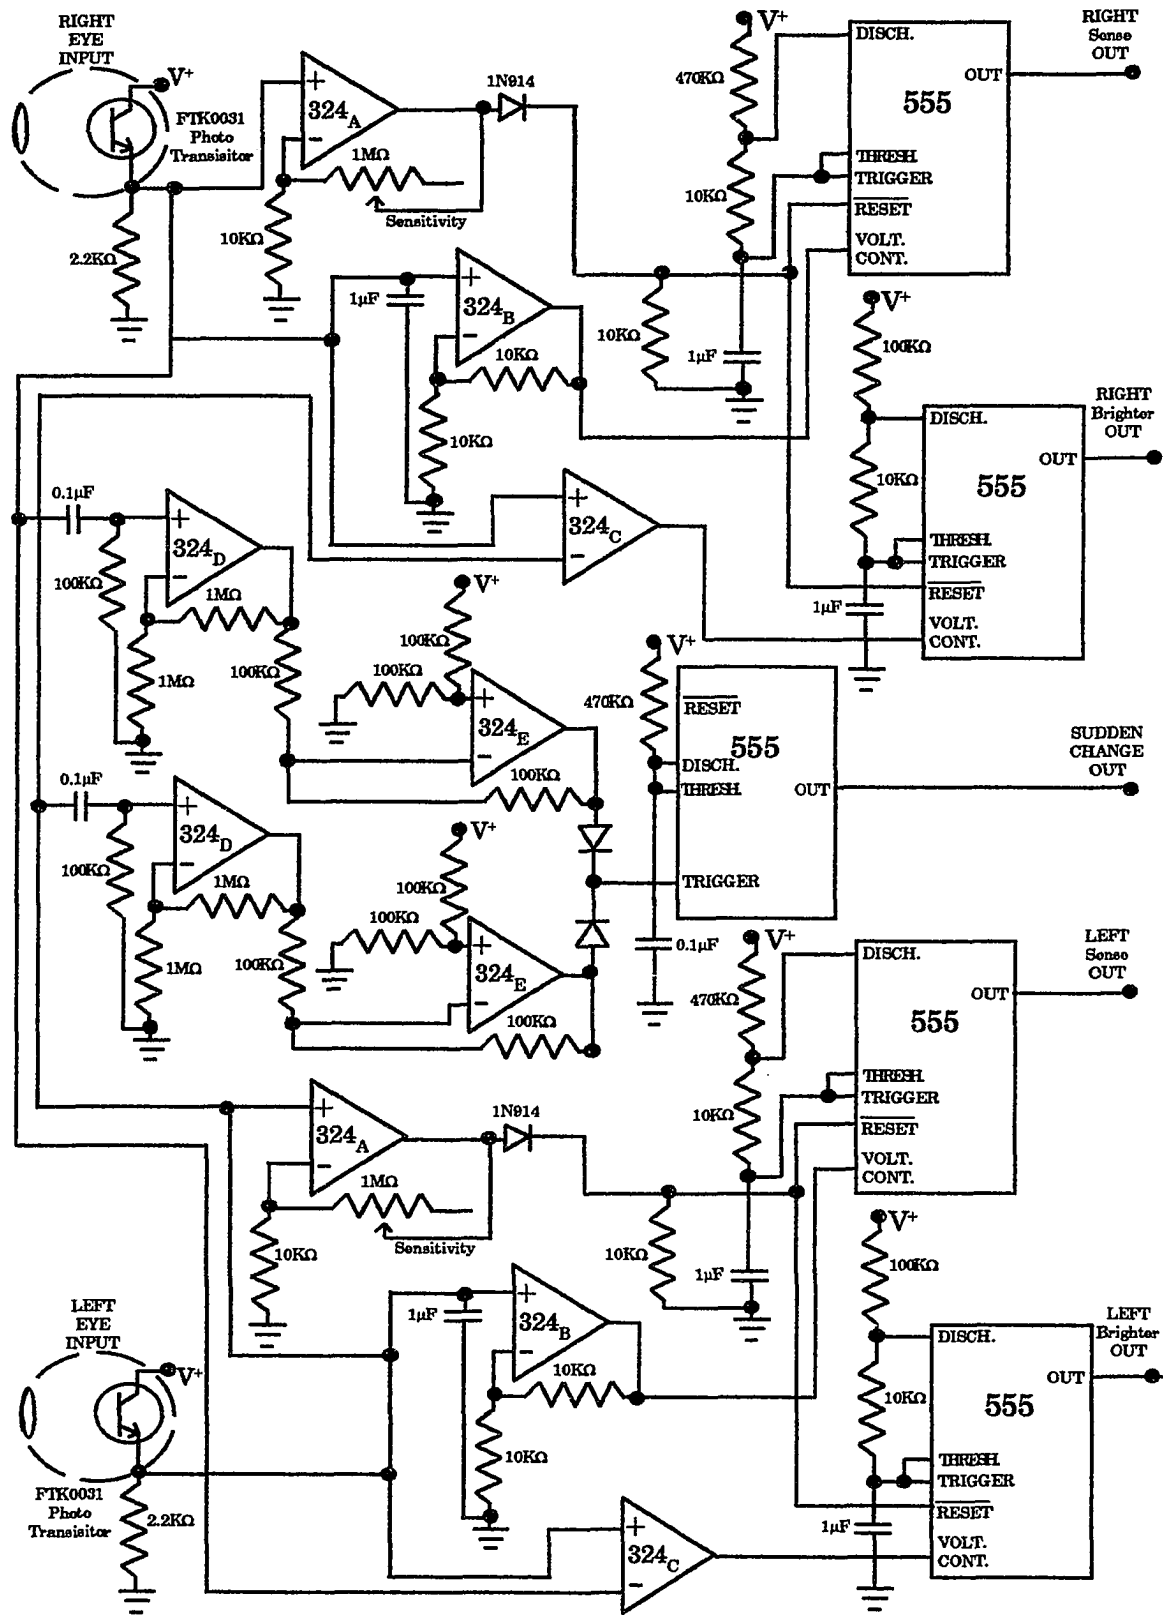


Figure 55. The light sensing circuitry. An explanation of function is in the text.

neuromime. The eye signal also passes to amplifier (B) on the same eye, which provides a gain of 100% and controls the frequency and duty cycle of the sense neuromime output. A third amplifier (C) acts as a comparator, allowing the light greater neuromime to fire only for the eye which has greater illumination. The interaction of amplifiers (A) and (B) prevents firing of the light greater neuromime for an eye when that eye is not registering illumination.

Finally, two remaining amplifiers, (D) and (E), combine to create a differentiator which generates a firing signal on the sudden-change-in-light neuromime when there is a significantly rapid change in illumination to that eye. This neuromime is wired to produce a single pulse when triggered.

Power Sensors and Systems

Power related senses include the batteries high, batteries low, and batteries charging sensors. The batteries charging, or feeding, circuit is given in figure 56. In figure 56, the IN4148 diode provides a voltage drop which is detected and amplified by the operational amplifier. The power capacitor and filter capacitors on the inputs prevent transient signals (from brief light flashes or heavy battery drain) from

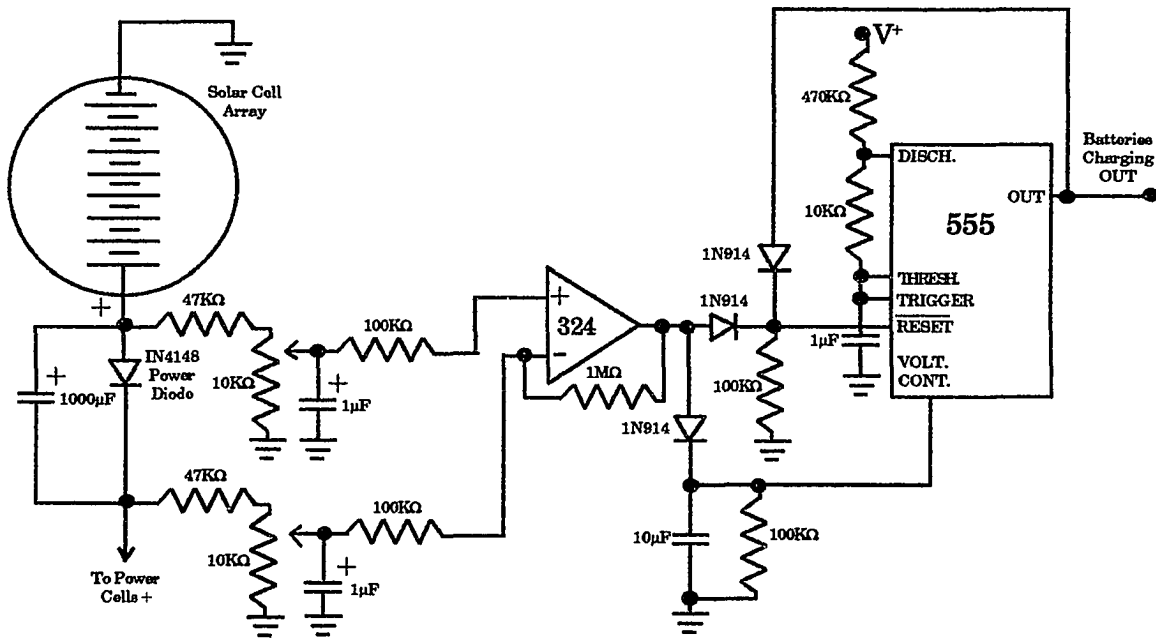


Figure 56. SENSI's feeding circuit. A batteries charging signal occurs when the current flow from the solar cell array to the batteries is strong enough.

generating spurious pulses. Potentiometers on the inputs allow for adjustment of sensitivity to the level of charging current. The output of the operational amplifier provides the trigger and controls the duty cycle and frequency of the 555 IC output.

The hunger and sated neuromimes are contained in a single circuit, given in figure 57. Unregulated battery voltage passes through an adjustable voltage divider and filter network to an inverting amplifier. A zener diode to regulated battery power supplies a stable, three volt reference. When the filtered battery voltage falls below the reference, the amplified difference triggers the hunger, or batteries low

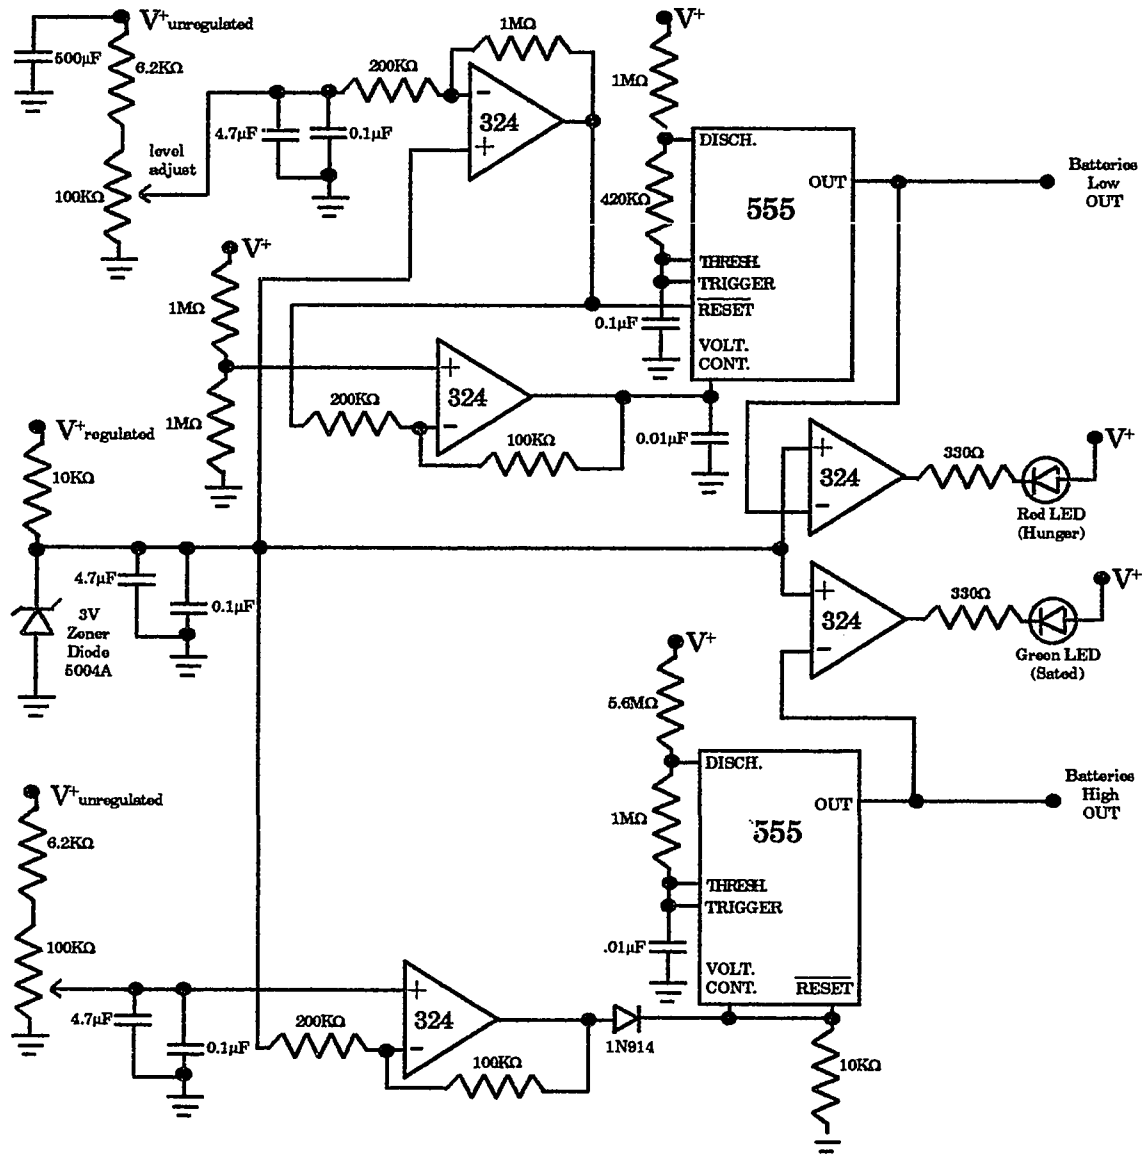


Figure 57. The batteries low (hunger) and batteries high (sated) circuitry.

neuromime. This trigger signal is inverted again, and its gain is adjusted by a second amplifier, to allow the hunger neuromime to respond differently over a wide range of voltage levels (range of hunger). Increasing hunger results in higher frequency, lower duty cycle response on the hunger neuromime. An operational amplifier used as a comparator drives an LED to visually signal the presence of hunger.

A non-inverting operational amplifier triggers the sated, or batteries high neuromime when the filtered battery voltage exceeds a preset reference level, in other words, when the cells are adequately charged. As with batteries low, a comparator drives a separate LED, which illuminates when the batteries high neuromime is active.

Because of its intimate connection with the charging circuitry, and to provide completeness, the power bus is diagrammed in figure 58. SENSI carries sixteen C-cells which provide power to the general system. The motor power system (MPS) supplies unregulated battery voltage to the motors through TIP111 power transistors and electromechanical relays. An MPS tap on the chassis supplies raw battery voltage to the feeding and hunger circuits and to the solar cell array.

The electronics power system (EPS) receives power from the batteries through a blocking power diode to a 914 5-volt regulator which is filtered by a large storage capacitor. EPS taps around the edge of the chassis supply regulated, filtered power to all electronic boards and components. Each board has individual additional filtering and despiking capacitors at the EPS tap point.

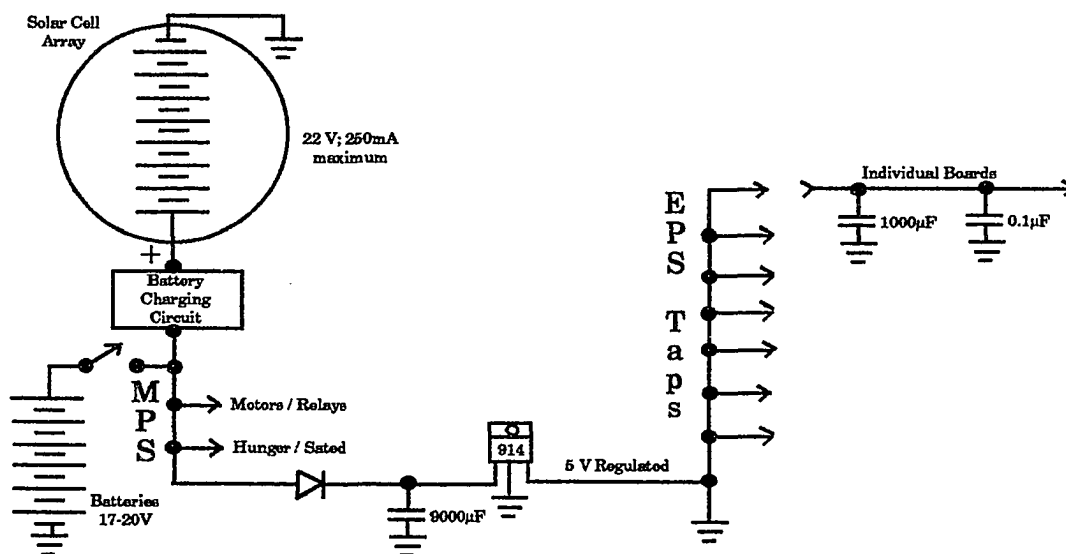


Figure 58. SENSI's power bus configuration.

APPENDIX D

SYNAPTIC PLASTICITY

The plastic synapse, which makes possible SENSI's learning network, is based on the functionality depicted in figure 59. We see various inputs to a certain neuromime, which are innate, nonplastic connections. The plastic synapse, represented by the box, connects a sampling cell to the aforementioned neuromime. When the sampling cell is active, or firing, the input weight of the plastic synapse to the postsynaptic neuromime should incrementally adjust to mimic some fraction of the innate input to the postsynaptic neuromime.

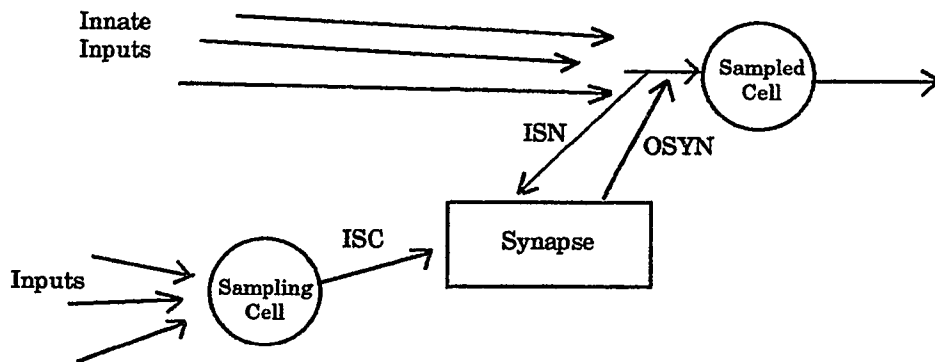


Figure 59. A preliminary synapse diagram, indicating function. An explanation of terms and mnemonics is in the text.

In figure 59, we observe that there are three interfaces with the synapse: ISC, which stands for Input from the Sampling Cell; ISN, meaning Input from the Sampled Neuromime, and OSYN, for Output of the SYNapse to the sampled Neuro-mime. Briefly, ISC must be present for the synapse to learn, ISN is what it is attempting to learn, and OSYN is the effect of the synapse on the postsynaptic cell. Naturally, OSYN is only present when ISC is present.

Figure 60 presents the electronics of our plastic synapse. Note the location of the three interfaces, ISC, ISN, OSYN. The circuit joining these points, the synapse, consists of four functional blocks: the input buffer, the learning rate network, the sample and hold circuit, and the output buffer / scaler.

The input buffer is a high impedance, slightly greater than unity gain amplifier, which receives the summed input voltage from the input network on the

sampled neuromime. The buffered voltage is passed to the learning rate network, where the input voltage (positive only) is reduced and passed to the sample and hold circuit.

The sample and hold block is a large valued (slow-changing) capacitor which charges when the sampling cell activates the synapse's 4066 switch. The switch isolates the sample and hold capacitor when the synapse is not active. When active, the stored excitation level is passed to the output stage, which consists of a unity gain amplifier to sample the capacitor's voltage and pass it on to the postsynaptic cell with minimal drain. The 4066 switch isolates the synapse from the postsynaptic cell except when it is active (which occurs when the sampling cell is firing). The NSR resistor performs two important functions. It clamps the amplifier low in the absence of voltage on the capacitor SHC, and it allows unlearning, or drain on SHC, in the absence of input to the sampled cell at the time of sampling.

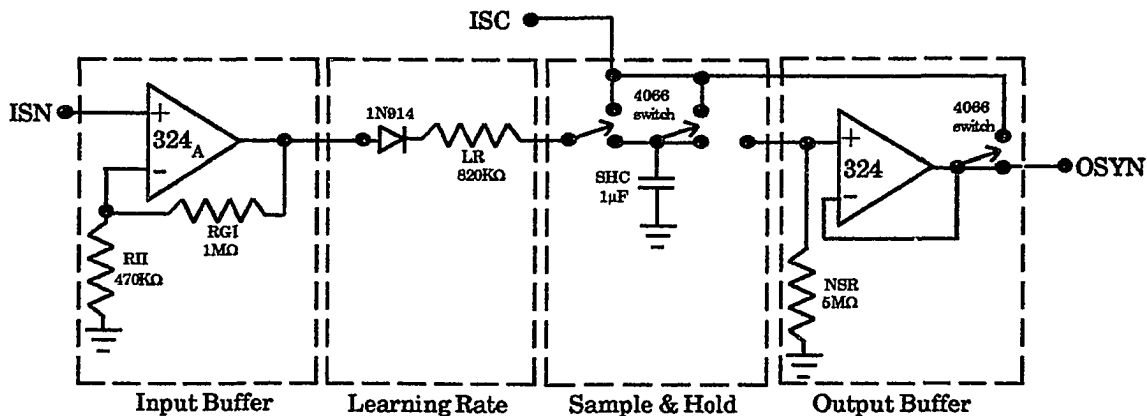


Figure 60. The plastic synapse. An explanation of the circuit is in the text.

APPENDIX E

OVERALL NERVOUS SYSTEM

The complete neuromime connection diagram for SENSI's nervous system is given in figure 61. Table 3 re-lists the neuromime mnemonics and their meanings.

Table 3. Neuromime Lables and Their Meanings

AL	Adequate Light
B	Boredom
BC	Batteries Charging
BD	Boredom Drive
BDA	Boredom Arousal
BDC	Boredome Drive Competition
BDR	Boredome Drive Representation
BL	Batteries Low
BH	Batteries High
Dxx	Directional Interneuromimes
ESC	Eye; Sudden Change
FD	Fear Drive
FDC	Fear Drive Competition
FDR	Fear Drive Representation
FFx	Fast Fatigable Motoneuromime
FM	Frustrated Motion
FRx	Fast, fatigue Resistent Motoneruomime
FTx	Forward Turn Neuromimes
HD	Hunger Drive
HDC	Hunger Drive Competition
HDR	Hunger Drive Representation
IH	Halt Interneuromime
IxxC	Common Motor Interneuromimes
IxxR	Reflex Motor Interneuromimes
LSx	Light Sensor Neuromimes
LGx	Light Greater Neuromimes
MSx	Motion Sensor Neuromimes
Px	Pain Neuromimes
Sx	Slow Motoneuromimes
S1xxx	Sampling Cell; First stage
S2xxx	Sampling Cell; Second stage
Tx	Touch Neuromimes
R+	Reinforcement Neuromime

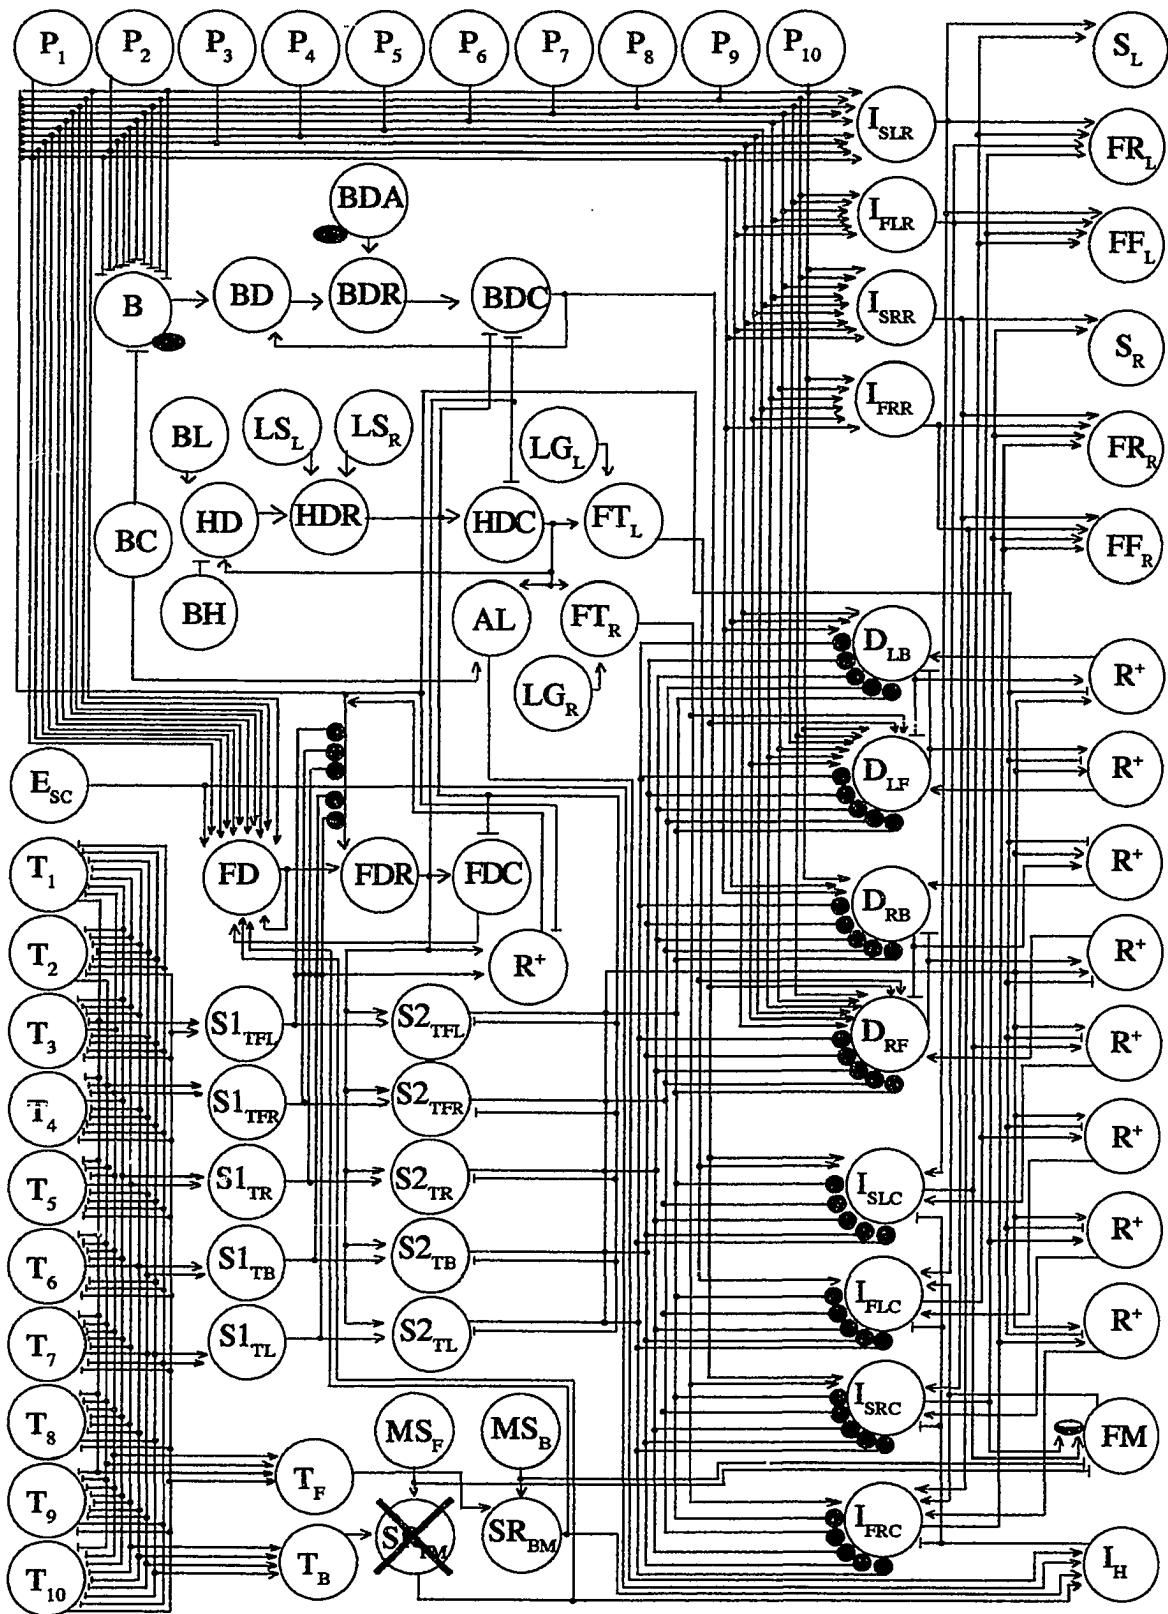


Figure 61. SENSI's connection overview. See table 3 for a list of neuromime labels.

BIBLIOGRAPHY

- [Aitken, 1994] Ashely M. Aitken. An architecture for learning to behave. In Cliff, et al., editors, *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press, 1994.
- [Anderson & Rosenfield, 1988] James A. Anderson and Edward Rosenfeld, editors. *Neurocomputing: Foundations of Research*. MIT Press, 1988.
- [Anderson, et al., 1990] James A. Anderson, Andras Pellionisz, and Edward Rosenfeld, editors. *Neurocomputing 2: Directions for Research*. MIT Press, 1990.
- [Beach, et al., 1960] Beach, Hebb, Morgan, and Nissen, editors. *Neuropsychology of Lashley – Selected Papers of K. S. Lashley*. McGraw-Hill, 1960.
- [Beer, 1990] Randall D. Beer. *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology*. Academic Press, 1990.
- [Bersini, 1994] Hugues Bersini. Reinforcement learning for homeostatic endogenous variables. In Cliff, et al., editors, *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press, 1994.
- [Blumberg, 1994] Bruce Blumberg. Action-Selection in Hamsterdam: Lessons from Ethology. In D. Cliff, et al., editors, *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press, 1994.
- [Booker, 1982] Lashon Bernard Booker. *Intelligent behavior as an Adaptation to the Task Environment*. PhD thesis, University of Michigan, 1982.
- [Braham, 1990] Rafik Braham and James O. Hamblen. The design of a neural network with biologically motivated architecture. *IEEE Transactions on Neural Networks*, 1(3)251-262, September 1990.
- [Braitenberg, 1984] Valentino Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. MIT Press, 1984.
- [Brooks, 1994] Rodney Brooks. Coherent behavior from many adaptive processes. In D. Cliff, et al., editors, *From animals to animats 3: Proceedings of the Third International*

Conference on Simulation of Adaptive Behavior. MIT Press, 1994.

- [Brooks, 1992] Rodney Brooks. Artificial life and real robots. In F. Varela and P. Bourguine, editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life.* MIT Press, 1992.
- [Brooks, 1991] Rodney Brooks, Intelligence without representation. *AI*, 47:139-159, 1991.
- [Brooks, 1986] Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14-23, March 1986.
- [Bunge, et al., 1978] Richard Bunge, Mary Johnson, and C. David Ross. Nature and nurture in development of the autonomic neuron. *Science*, 199:1409-1416, March 1978.
- [Burke, 1990] Robert E. Burke. Spinal cord: ventral horn. In Gordon M. Shepherd, editor, *The Synaptic Organization of the Brain, Third Edition.* Oxford University Press, 1990.
- [Burrows & Laurent, 1989] Malcolm Burrows and Giles Laurent. Reflex circuits and the control of movement. In Durbin, Miall, and Mitchison, editors, *The Computing Neuron.* Addison-Wesley Publishing Company, 1989.
- [Carpenter & Grossberg, 1991a] Gail A. Carpenter and Stephen Grossberg. ART2: Self-organization of stable category recognition codes for analog input patterns. In *Pattern Recognition by Self-Organizing Neural Networks*, chapter 12. MIT Press, 1991.
- [Carpenter & Grossberg, 1991b] Gail A. Carpenter and Stephen Grossberg. ART3: Hierarchical search using chemical transmitters in self-organizing pattern recognition architectures. In *Pattern Recognition by Self-Organizing Neural networks*, chapter 14. MIT Press, 1991.
- [Carpenter & Grossberg, 1991c] Gail A. Carpenter and Stephen Grossberg. A massively parallel architecture for a self-organizing neural pattern recognition machine. In Gail A. Carpenter and Stephen Grossberg, editors, *Pattern Recognition by Self-Organizing Neural Networks*, chapter 10. MIT Press, 1991.
- [Charlesworth, 1976] W. R. Charlesworth. Human intelligence as adaptation: An ethological approach. In L. B. Resnick, editor, *The Nature of Intelligence.* Laurence Erlbaum Associates, 1976.
- [Cliff, et al., 1994] Dave Cliff, P. Husbands, J. Meyer, and S. Wilson, editors. *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior.* MIT Press, 1994.
- [Connell, 1990] Jonathan H. Connell. *Minimalist Mobile Robotics: A Colony-style architecture for an artificial creature.* Academic Press, 1990.
- [Coward, 1990] L. Andrew Coward. *Pattern Thinking.* Praeger Publishers, 1990.
- [Crick & Asanuma, 1986] F. Crick and C. Asanuma. Certain aspects of the anatomy and physiology of the cerebral cortex. In J. A. Feldman, P. J. Hayes, and D. E. Rumelhardt,

- editors, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. MIT Press, 1986.
- [Dennett, 1985] David C. Dennett, Can machines think? In Micheal Shafto, editor, *How We Know*, Harper & Row, 1985.
- [DeWeerth, et al., 1991] Stephen P. DeWeerth, Lars Nielsen, Carver A. Mead, and Karl J. Åström. A simple neuron servo. *IEEE Transactions on Neural Networks*, 2(2):248-251, March 1991.
- [Downs & Stea, 1973] Roger M. Downs and David Stea. Cognitive maps and spatial behavior: Process and products. In Roger M. Downs and David Stea, editors, *Image and Environment*. Aldine Publishing Company, 1973.
- [Edelman, 1978] Gerald M. Edelman. Group selection and phasic reentrant signalling: A theory of higherbrain function. In *The Mindful Brain*, chapter 2. MIT Press, 1978.
- [Edelman, 1983] Gerald M. Edelman. Cell adhesion molecules. *Science*, 219:450-457, 1983.
- [Edelman, 1985] Gerald M. Edelman. Neural darwinism: Population thinking and higher brain function. In Michael Shafto, editor, *How We Know*, chapter 1. Harper & Row, 1985.
- [Fikes & Nilsson, 1971] R. E. Fikes and N. J. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, Volume 2, 189-208. North Holland Publishing Company, 1971.
- [Foner & Maes, 1994] Paying attention to what's important: using focus of attention to improve unsupervised learning. In Cliff, et al., editors, *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press, 1994.
- [Frye, 1991] Robert C. Frye, Edward A. Reitman, and Chee C. Wong. Back-propagation learning and nonidealities in analog neural network hardware. *IEEE Transactions on Neural Networks*, 2(1):110-117, January 1991.
- [Gaussier & Zrehen, 1994] P. Gaussier and S. Zrehen. A topological neural map for on-line learning: emergence of obstacle avoidance in a mobile robot. In Cliff, et al., editors, *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press, 1994.
- [Gizster, 1994] Simon Gizster. Reinforcement tuning of action synthesis and selection in a 'virtual frog.' In Cliff, et al., editors, *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press, 1994.
- [Gleick, 1987] James Gleick. *Chaos: Making a New Science*. Viking Penguin, Inc. 1987.
- [Grossberg, 1988] Stephen Grossberg. Competitive learning: From interactive activation to adaptive resonance. In Stephen Grossberg, editor, *Neural Networks and Natural Intelligence*, chapter 5. MIT Press, 1988.

- [Grossberg, 1987a] Stephen Grossberg. *The Adaptive Brain I: Cognition, Learning, Reinforcement and Rhythm*. Elsevier Science Publishers B. V., 1987.
- [Grossberg, 1987b] Stephen Grossberg. A psychophysical and pharmacological correlates of a developmental cognitive and motivational theory. In Stephen Grossberg, editor, *The Adaptive Brain I*, chapter 1. Elsevier Science Publishers B. V., 1987.
- [Grossberg, 1987c] Stephen Grossberg. Some psychophysical and pharmacological correlates of a developmental cognitive and motivational theory. In Stephen Grossberg, editor, *The Adaptive Brain I*, chapter 2. Elsevier Science Publishers B. V., 1987.
- [Grossberg, 1976a] Stephen Grossberg. Adaptive pattern classification and universal recoding: I. Parallel development and coding of neural feature detectors. *Biological Cybernetics*, 23:121-134, 1976.
- [Grossberg, 1976b] Stephen Grossberg. Adaptive pattern classification and universal recoding: II. Feedback, expectation, olfaction, illusions. *Biological Cybernetics*, 23:187-202, 1976.
- [Grossberg & Levine, 1988] Stephen Grossberg and Daniel S. Levine. Neural dynamics of attentionally modulated pavlovian conditioning: Blocking, interstimulus interval, and secondary reinforcement. In Stephen Grossberg, editor, *Neural Networks and Natural Intelligence*, chapter 9. MIT Press, 1988.
- [Grossberg & Schmajuk, 1988] Stephen Grossberg, and Nestor A. Schmajuk. Neural dynamics of attentionally modulated pavlovian conditioning: Conditioned reinforcement, inhibition, and opponent processing. In Stephen Grossberg, editor, *Neural Networks and Natural Intelligence*. chapter 10. MIT Press, 1988.
- [Hebb, 1949] D. O. Hebb. *The Organization of Behavior, A Neuropsychological Theory*. John Wiley & Sons, Inc., 1949.
- [Hebb, 1966] D. O. Hebb. *Textbook of Psychology*. W. B. Saunders Company, 1966.
- [Hebb, 1980] D. O. Hebb. *Essay on Mind*. Laurence Erlbaum Associates, 1980.
- [Hart & Moore, 1973] Roger A. Hart and Gary T. Moore. The development of spatial cognition: A review. In Roger M. Downs and David Stea, editors, *Image and Environment*. Aldine Publishing Company, 1973.
- [Holland, 1975] J. H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, 1975.
- [Horwitz, 1983] Barry Horwitz. Unequal diameters and their effects on time-varying voltages in branched neurons. *Biophysical Journal*, 41:51-66, January 1983.
- [Kaelbling, 1992] Leslie Pack Kaelbling. An adaptable mobile robot. In F. Varela and P. Bourguin, editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*. MIT Press, 1992.

- [Kandel, 1979] Eric R. Kandel. Small systems of neurons. In *Mind and Behavior, Readings from Scientific American*, Chapter 1. W. H. Freeman and Company, 1979.
- [Kaplan, 1973] Stephen Kaplan. Cognitive maps in perception and thought. In Roger M. Downs and David Stea, editors, *Image and Environment*. Aldine Publishing Company, 1973.
- [Kaplan & Kaplan, 1982] Stephen Kaplan and Rachael Kaplan. *Cognition and Environment—Functioning in an Uncertain World*. Praeger Publishers, 1982.
- [Kohonen, 1987] Teuvo Kohonen. *Self-Organization and Associative Memory*. Springer-Verlag, 1987.
- [Kosko, 1991] Bart Kosko. *Neural Networks and Fuzzy Systems: A Dynamical Approach to Machine Intelligence*. Prentice Hall, 1991.
- [Lynch, 1986] Gary Lynch. *Synapses, Circuits, and the Beginnings of Memory*. MIT Press, 1986.
- [MacGregor, 1987] Ronald J. MacGregor. *Neural and Brain Modeling*. Academic Press, 1987.
- [MacGregor & Lewis, 1977] Ronald J. MacGregor and Edwin R. Lewis, editors. *Neural Modeling: Electrical Signal Processing in the Nervous System*. Plenum Press, 1977.
- [Maes, 1990] Pattie Maes, editor. *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*. MIT Press, 1990.
- [Maes, 1989] Pattie Maes. *How to Do the Right Thing*. A. I. Memo 1180, Massachusetts Institute of Technology. December, 1989.
- [Mataric, 1990a] Maja J. Mataric. A distributed model for mobile robot environment learning and navigation. Master's thesis, Massachusetts Institute of Technology, AI Lab, May 1990. Tech. REp. 1228.
- [Mataric, 1990b] Maja J. Mataric. Navigating with a rat brain: A neurobiologically-inspired model for robot spatial representation. In *Proceedings Adaptive Behavior Conference, From Animals to Animats*. MIT Press, 1990.
- [Mataric & Brooks, 1990] Maja J. Mataric and Rodney A. Brooks. Learning a distributed map representation based on navigation behaviors. In *Proceedings U.S.A.-Japan Symposium on Flexible Automation*, 1990.
- [Mead, 1991] Carver A. Mead, Xavier Arreguit, and John Lazzaro. Analog VLSI model of binaural hearing. *IEEE Transactions on Neural Networks*, 2(2):230-236, March 1991.
- [Mead, 1989] Carver A. Mead. *Analog VLSI and Neural Systems*. Addison-Wesley Publishing Company, 1989.
- [Millán, 1994] José del R. Millán. Learning efficient reactive behavioral sequences from basic

- reflexes in a goal-directed autonomous robot. In Cliff, et al., editors, *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press, 1994.
- [Murray et al., 1991] Alan F. Murray, Dante Del Corso, and Lionel Tarassenko. Pulse-stream VLSI neural networks mixing analog and digital techniques. *IEEE Transactions on Neural Networks* 2(2):193-204. March, 1991.
- [Nadel, 1980] Lynn Nadel. Cognitive and neural maps. In P. Jusczyk and R. Klein, editors, *The Nature of Thought: Essays in Honor of D. O. Hebb*. Laurence Erlbaum Associates, 1980.
- [Nelson, 1989] Katherine Nelson. Remembering: A functional developmental perspective. In Solomon, Goethals, Kelley, and Stephens, editors, *Memory: Interdisciplinary Approaches*. Springer-Verlag, 1989.
- [Nii, 1986a] H. Penny Nii. Blackboard systems: the blackboard model of problem solving and the evolution of blackboard architectures. *AI Magazine*. Summer, 1986.
- [Nii, 1986b] H. Penny Nii. Blackboard systems: blackboard application systems, blackboard systems form a knowledge engineering perspective. *AI Magazine*. August, 1986.
- [Olton, et al., 1979] David S. Olton, James T. Becker, and Gail E. Handelmann. Hippocampus, space, and memory. *The Behavioral and Brain Sciences*, 2:313-365, 1979.
- [Patel & Schnepf, 1992] Mukesh J. Patel and Uwe Schnepf. Concept formation as emergent phenomena. In F. Varela and P. Bourguine, editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*. MIT Press, 1992.
- [Pfeifer & Verschure, 1992] Rolf Pfeifer and Paul Verschure. Distributed adaptive control: a paradigm for designing autonomous agents. In F. Varela and P. Bourguine, editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*. MIT Press, 1992.
- [Popper, 1968] Karl R. Popper. *Conjectures and Refutations: The Growth of Scientific Knowledge*. Harper & Row, 1968.
- [Rosenzweig, 1972] Mark R. Rosenzweig, Edward L. Bennett, and Marian Cleeves Diamond. Brain changes in response to experience. In *Mind and Behavior, Readings from Scientific American*. W. H. Freeman and Company, 1972.
- [Rumelhart, et al., 1986] David E. Rumelhart, James L. McClelland, and the PDP Research Group. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Volume I. MIT Press, 1986.
- [Routtenberg, 1979] Aryeh Routtenberg. The reward system of the brain. In *Mind and Behavior, Readings from Scientific American*, chapter 3. W. H. Freeman and Company, 1979.
- [Sahota, 1994] Michael K. Sahota. Action selection for robots in dynamic environments through

- inter-behaviour bidding. In D. Cliff et al., editors. *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press, 1994.
- [Sacerdoti, 1974] E. D. Sacerdoti. Planning in a hierarchy of abstraction spaces. *Artificial Intelligence*, (5)115-135. North Holland Publishing Company, 1974.
- [Schwartz, 1990] Tom J. Schwartz. A neural chips survey. *AI Expert*, December 1990.
- [Segev & Parnas, 1983] Idan Segev and Itzhak Parnas. Synaptic integration mechanisms: Theoretical and experimental investigation of temporal postsynaptic interactions between excitatory and inhibitory inputs. *Biophysical Journal*, 41:41-50, January, 1983.
- [Shepherd, 1990] Gordon M. Shepherd, editor. *The Synaptic Organization of the Brain, Third Edition*. Oxford University Press, 1990.
- [Stefik, 1981] M. Stefik. Planning and meta-planning (MOLGEN: part 2). *Artificial Intelligence*, (16)141-170. North Holland Publishing Company, 1981.
- [Sussman & McDermott, 1972] Gerald J. Sussman and Drew V. McDermott. From PLANNER to CONNIVER – a genetic approach. *Joint Computer Conference Proceedings*, 1171-1179, 1972.
- [Thomson, et al, 1983] Richard F. Thomson, Theodore W. Berger, and John Madden IV. Cellular processes of learning and memory in the mammalian CNS. *Annual Review of Neuroscience*, 6:447-491, 1983.
- [Thorpe, 1956] W. H. Thorpe. *Learning and Instinct in Animals*. Harvard University Press, 1956.
- [Tsay & Newcomb, 1991] Suan W. Tsay and Robert W. Newcomb. VLSI implementation of ART-1 memories. *IEEE Transactions on Neural Networks*, 2(2)214-221, March 1991.
- [Uhr, 1966] L. Uhr. Pattern recognition. In L. Uhr, editor, *Pattern Recognition: Theory, Experiment, Computer Simulations and Dynamic Models of Form Perception and Discovery*. Wiley, 1966.
- [Vernon, 1966] M. D. Vernon. The nature of perception and the fundamental stages in the process of perceiving. In L. Uhr, editor, *Pattern Recognition: Theory, Experiment, Computer Simulations and Dynamic Models of Form Perception and Discovery*. Wiley, 1966.
- [Walter, 1953] W. Grey Walter. *The Living Brain*. W. W. Norton & Company, Inc., 1953.
- [Waterman, 1989] Talbot H. Waterman. *Animal Navigation*. Scientific American Library, 1989.
- [Young, 1966] J. Z. Young. Memory mechanisms of the brain. In L. Uhr, editor, *Pattern Recognition: Theory, Experiment, Computer Simulations and Dynamic Models of Form*

Perception and Discovery. Wiley, 1966.

[Zeidenberg, 1990] Matthew Zeidenberg. *Neural Networks in Artificial Intelligence.* Ellis Horwood, 1990.