

INTELLIGENT BEHAVIOR AS AN ADAPTATION  
TO THE TASK ENVIRONMENT

by  
Lashon Bernard Booker

A dissertation submitted in partial fulfillment  
of the requirements for the degree of  
Doctor of Philosophy  
(Computer and Communication Sciences)  
in The University of Michigan  
1982

Doctoral Committee:

Professor John H. Holland, Co-Chairman  
Professor Stephen Kaplan, Co-Chairman  
Assistant Professor Warren G. Holmes  
Assistant Professor Paul D. Scott

THE UNIVERSITY OF MICHIGAN  
ENGINEERING LIBRARY

en sr

UMR 977

v.1

To Doreen,  
for being so extraordinarily patient

## ABSTRACT

### INTELLIGENT BEHAVIOR AS AN ADAPTATION TO THE TASK ENVIRONMENT

by

Lashon Bernard Booker

Co-Chairmen: John H. Holland, Stephen Kaplan

As research in artificial intelligence focuses on increasingly complex task domains, a key question to be resolved is how to design a system that can efficiently acquire knowledge and gracefully adapt its behavior in an uncertain environment. This dissertation argues that examining more closely the way animate systems cope with real-world environments can provide valuable insights about the structural requirements for intelligent behavior.

Accordingly, a class of simulated environments is designed that embodies many of the important functional properties characteristic of natural environments. A new type of adaptive system is then defined that uses pattern-directed, rule-based processing to cope with uncertain information. As a rule-based system, the system presented here is notable in that several rules can be active at once and there are no fixed priorities determining the order in which rules can be activated. Moreover, the syntax of each rule is simple enough to make a powerful learning heuristic applicable - one that is provably more efficient than the

techniques used in most other adaptive rule-based systems.

A simple version of the adaptive system is implemented as a hypothetical organism having to locate resources and avoid noxious stimuli by generating temporal sequences of actions in a simulated environment. Simulation results show that the naive organism quickly acquires the knowledge required to function effectively. Further experiments show that the system is capable of discriminating a large class of schematic patterns; and, that prior learning experiences transfer to novel situations.

The results presented here demonstrate that activity in a collection of simple computational elements - operating in parallel and activated stochastically - can be orchestrated to produce reliable behavior in a challenging environment. The system touches on several issues related to cognitive functioning such as the generic representation of objects and the management of limited processing resources. These issues have been addressed in a way that is computationally feasible and that allows for rigorous testing.



## ACKNOWLEDGEMENTS

"Apart from me you can do nothing"  
John 15:5

Amen.

First and foremost, I thank God for the many opportunities and resources made available to me over the years. The fact that I was able to complete my education to the Ph.D. level is a tremendous blessing for which I am deeply and humbly grateful.

My co-chairmen have been the most important influences on my life as a graduate student. My gratitude for their friendship and tutelage goes far beyond the scope of any brief acknowledgement. John Holland has been a valued source of wisdom and enthusiastic support. His pioneering insights about adaptive systems have obviously had a profound influence on my research. Stephen Kaplan introduced me to the world of psychology and convinced me that a project of this kind would be worthwhile and fun. He was right. I am very grateful for his expertise, encouragement, and advice.

I am also indebted to the other members of my committee. Paul Scott proved to be a reliable source of helpful suggestions and insightful observations. I wish he had come to Michigan sooner. Warren Holmes' comments were a tremendous help in making each chapter readable. I appreciate his many hours of hard work.

This research would not have been possible without the

financial support of the Ford Foundation, the IBM Corporation, the Rackham School of Graduate Studies, and Natural Science Foundation Grant MCS78-26016. I also thank the Logic of Computers Group, directed by Arthur Burks, for its support and the use of its computing facilities.

Finally, I thank my family for their steadfast love and support.

This report is reproduced in two parts:  
 Part I goes to page 150  
 Part II starts at page 151

TABLE OF CONTENTS

DEDICATION . . . . .	ii
ACKNOWLEDGEMENTS . . . . .	iii
LIST OF TABLES . . . . .	vii
LIST OF ILLUSTRATIONS. . . . .	viii
LIST OF APPENDICES . . . . .	x
CHAPTER	
I. INTRODUCTION . . . . .	1
Models of Biological Systems . . . . .	3
Knowledge Engineering. . . . .	5
Functional Constraints on Processing . . . . .	7
Summary. . . . .	11
II. RESEARCH STRATEGY AND SCOPE. . . . .	13
Approach to the Problem. . . . .	14
The Experimental Frame . . . . .	15
Review of Related Studies. . . . .	18
Summary. . . . .	31
III. A SIMPLE MODEL OF INSTINCTIVE BEHAVIOR . . . . .	34
The Organism/Environment Interface . . . . .	36
Cues and their Implications. . . . .	39
Spatial Information. . . . .	44
Implications for an Artificial Environment . . . . .	51
Simulation of a Simple Model . . . . .	57
A Basic Goal-Directed System . . . . .	72

IV.	SOME CONSEQUENCES OF UNCERTAINTY . . . . .	86
	Specifying Stimulus Patterns . . . . .	88
	Identification of Distal Objects . . . . .	102
	Modifying Stimulus-Response Probabilities. . . . .	114
	A Revised Design . . . . .	121
	Implementation as a Classifier System. . . . .	127
V.	THE LEARNING ALGORITHM . . . . .	151
	Criteria for a Mechanism . . . . .	152
	A Theoretical Framework. . . . .	155
	Foundations for an Adaptive Plan . . . . .	160
	Modifying the Genetic Algorithm. . . . .	177
	Testing the Modified Algorithm . . . . .	209
	Summary. . . . .	213
VI.	IMPLEMENTATION OF THE ADAPTIVE SYSTEM. . . . .	216
	Tags and Concept Formation . . . . .	217
	The Hypothetical Organism as an Adaptive System . . . . .	231
	Simulations of the Adaptive System . . . . .	243
VII.	SUMMARY AND CONCLUSIONS. . . . .	269
	Statistical Information Processing . . . . .	271
	More Sophisticated Systems . . . . .	274
	Structuring a Task Environment . . . . .	281
	Implications for Artificial Intelligence . . . . .	284
	APPENDICES . . . . .	288
	REFERENCES . . . . .	318

## LIST OF TABLES

	page
<u>Table</u>	
5.1. Comparison of four match scores. . . . .	184
5.2. Second comparison of four match scores . . .	186
5.3. Performance of G2 on two pattern classes . .	211
6.1. Five concept learning tasks. . . . .	229
6.2. Sample classifiers from a trained organism .	257
A.1. Performance of R0 and R1 on <i>E</i> . . . . .	294
A.2. Performance of R1 and R2 on <i>E</i> . . . . .	297
A.3. Performance of R2 and R3 on <i>E</i> . . . . .	306
A.4. Performance of R3 and R4 on <i>E</i> . . . . .	312
A.5. Performance of R0 and R4 on <i>E</i> . . . . .	314

## LIST OF ILLUSTRATIONS

<u>Figure</u>	page
2.1. The environment for Doran's automaton. . . .	20
2.2. The environment for Holland and Reitman's classifier system. . . . .	29
3.1. The interaction between an adaptive system and its environment. . . . .	37
3.2. Typical distribution of signal intensities in an environment. . . . .	55
3.3. An innate releasing mechanism. . . . .	59
3.4. A simple organism and environment. . . . .	62
3.5. The environment used for organism simulations 1, 2 and 3 . . . . .	63
3.6. Some representative approach trajectories. .	65
3.7. Some typical avoidance trajectories. . . . .	70
3.8. A simple mechanism for motivational control.	77
3.9. The environment used for organism simulation 4. . . . .	78
3.10. A typical component in the organization of an instinct according to Tinbergen . . . . .	82
3.11. The control structure of the goal-seeking organism . . . . .	84
4.1. The structure of a simple classifier system.	109
4.2. The integration of perception, affect, and learning into the model. . . . .	122
4.3. Implementation of the model using two message lists. . . . .	130
5.1. An example of interference matching. . . . .	163
5.2. The crossover operator . . . . .	172
5.3. Producing fewer offspring per generation . .	189

5.4.	Deterministic versus stochastic selection of parents. . . . .	191
5.5.	On-line performance as a function of varying SETSIZE. . . . .	194
5.6.	The performance of G1 for two values of SETSIZE. . . . .	196
5.7.	A comparison of G0 and G1. . . . .	197
5.8.	A comparison of G1 and G2. . . . .	203
5.9.	Improvements to the crowding algorithm . . .	204
5.10.	Performance of G2 as a function of the crossover rate . . . . .	207
5.11.	A comparison of G0 and G2. . . . .	208
6.1.	On-line performance when tag scores are included in payoff . . . . .	221
6.2.	Cumulative error rates when tag scores are included in payoff . . . . .	222
6.3.	On-line performance using actual and estimated tag scores . . . . .	227
6.4.	Cumulative categorization error using actual and estimated tag scores . . . . .	228
6.5.	Learning curves for several categorization tasks. . . . .	232
6.6.	The two channel environment used for the "training" simulations . . . . .	246
6.7.	The behavior of the adaptive system during the training interval. . . . .	248
6.8.	Learning curves for each of the five organisms. . . . .	250
6.9.	The effects of LEARNRATE on the behavior of the organism . . . . .	253
6.10.	Performance of the adaptive system with and without the genetic algorithm. . . . .	258
6.11.	Experiment demonstrating positive transfer learning . . . . .	260

6.12.	Experiment demonstrating negative transfer learning . . . . .	262
6.13.	Average results from the serial reversal learning experiment. . . . .	265
6.14.	The best performance observed in the serial reversal learning experiment . . . . .	266
A.1.	The genetic plan R0. . . . .	291
A.2.	On-line performance of R0, R1, and R2 on F1. . . . .	298
A.3.	Off-line performance of R0, R1, and R2 on F1 . . . . .	299
A.4.	Allele loss for R0, R1, and R2 on F1 . . . . .	300
A.5.	On-line performance of R2 and R3 on F1 . . . . .	303
A.6.	Off-line performance of R2 and R3 on F1. . . . .	304
A.7.	Allele loss for R2 and R3 on F1. . . . .	305
A.8.	Two versions of the generalized crossover operator . . . . .	308
A.9.	On-line performance of R3 and R4 on F1 . . . . .	309
A.10.	Off-line performance of R3 and R4 on F1. . . . .	310
A.11.	Allele loss for R3 and R4 on F1. . . . .	311
A.12.	Allele loss for R0 and R4 on F1. . . . .	313
A.13.	The genetic plan G0. . . . .	315

LIST OF APPENDICES

	page
<u>Appendix</u>	
A. The Reproductive Plan G0. . . . .	289
B. List of System Parameters . . . . .	316



## CHAPTER I

### INTRODUCTION

How can an artificial system be constructed so that it exhibits intelligent behavior? If we assume that "intelligent behavior is always manifest in performance as successful, orderly, adaptive, problem-oriented transactions with the environment" [Welker, 1976, p. 270], then the approach taken to this question will depend on how one characterizes the transactions between a system and its environment. The most widely accepted characterization is embodied in work done in that domain of computer science called artificial intelligence. Briefly, the idea is that most of the transactions between a system and its environment require an ability to store and manipulate symbols.

The central notion is that of the symbol, which is taken to mean essentially what it does in computer science, an entity with a certain functional property, to wit: that when a process has a token of a symbol it has access to information about what the symbol designates (encoded in symbolic expressions). The processes that can be performed on symbols are their creation (and, possibly, destruction), the obtaining of designated information, the creation of symbolic expressions, and the manipulation of these symbolic expressions by insertion, deletion, replacement, and reordering. [Newell, 1973, p. 27]

From this point of view, the way to build intelligent systems is to find powerful ways to store, retrieve, and manipulate symbolic expressions.

This paradigm emphasizes the similarities between the processing capabilities of machines and humans. It is a research strategy that has led to some impressive examples of machine generated intelligent behavior. It has had, moreover, a profound influence on the development of the "information processing" approach to psychology [McCorduck, 1979]. Nevertheless, it is clear that the view of man as an information processor - that is, as a symbol manipulating device - is a very narrow point of view. Animate information processing systems differ in rather striking ways from their nonbiological imitators. While considerable effort has been devoted to understanding the similarities between men and machines, little has been done to analyse the differences. Understanding those differences might lead to new insights about the structural requirements for intelligent behavior.

The most obvious difference between animate and artificial systems is that the former have been designed to survive and perform adaptively salient functions in real-world environments. The transactions between an animate system and its environment can therefore be characterized in terms of their adaptive significance to the system. From this point of view, intelligent behavior is a primary mode of adaptation to the environment.

... Most major characteristics of animals, behavioral as well as morphological, are products of evolution and thereby represent the species' latest attempt at adaptation to its environmental niche. Intelligence can be viewed as representing one such adaptation. [Charlesworth, 1976, p. 148]

Understanding the structural requirements for intelligent behavior in animate systems is aided enormously by an understanding of the functional requirements of the system's environmental niche. Can a similar perspective be helpful in considering intelligent behavior in artificial systems?

### Models of Biological Systems

The functional requirements of the environment are clearly important when the artificial system in question is proposed as a model of a biological system. Consider, for instance, the task of designing an artificial system to process visual information. Thoughtful consideration of the kinds of information directly available in natural environments [Gibson, 1966] has caused many psychologists to question the notion that experience of the visual world is mediated by some internal representation, symbolic or otherwise. It has been argued, for example, that the continuous optical flow from a scene of objects is not likely to be reconstructed from a series of discrete, static representations [Turvey, 1977]. Even those not willing to deny a role for some kind of discrete representation acknowledge the need to consider the salient aspects of a natural visual environment. Haber [1978, p. 3], for instance, has pointed out that "... the study of the

processing of impoverished scenes, such as simple line-drawn objects, may reveal a fundamentally different set of processing strategies than those found for richly informative natural scenes." Though line drawings make sense from the computational point of view as an obvious special case, they may be misleading with regard to the functional roles of natural mechanisms. Similar concerns have influenced Marr's [1979] research. He proposes representations for visual information based on an analysis of the functional structure of the human visual system and the information it has to work with. Marr emphasizes that any computational theory of vision should be constructed by taking a problem faced by natural systems and applying real world constraints to make the problem tractable.

Similar observations have been made regarding models of human cognition. Lachman and Lachman [1979] point out that

... information processing theories can and should take into consideration the evolutionary factors that have impacted human cognition for the following reasons. First, it may be possible to use evolutionary considerations to avoid dead ends - to eliminate hypotheses that are clearly implausible for an evolved organism even though they make sense for artificial intelligences. Second, attention to evolutionary factors may suggest areas in which the information-processing formalisms available from the computational disciplines are patently inadequate. Third, the evolutionary history of the human being may contain the explanation of research results that are puzzling for theories relying too heavily on assumed parallels between human cognition and machine intelligence. (p. 137)

This point of view was then shown to have practical consequences for the implementation of various semantic

network models.

### Knowledge Engineering

There is a sizable portion of the artificial intelligence community, however, that is not necessarily concerned with modeling human intelligence. Feigenbaum [1977] calls these researchers "knowledge engineers" and their research endeavors knowledge engineering. The goal of artificial intelligence, from this point of view, is to discover whatever techniques may be useful in the construction of intelligent systems. Design efforts should be focused on techniques that exploit specific knowledge about the task domain. Indeed, the efficiency of any particular system can be hampered by an emphasis on general techniques and mechanisms.

The fundamental problem of understanding intelligence is not the identification of a few powerful techniques, but rather the question of how to represent large amounts of knowledge in a fashion that permits their effective use and interaction.... The current point of view is that the problem solver (whether man or machine) must know explicitly how to use its knowledge - with general techniques supplemented by domain specific pragmatic know-how. [Goldstein and Papert, 1977, p. 85]

In other words, the most important general principle underlying intelligent behavior is the need for easy access to as much relevant symbolic knowledge as possible. This includes whatever heuristics, advice, intuitions, etc. that an expert might bring to bear on the given task.

There is no doubt that the representation and use of

knowledge is central to any study of intelligent behavior. Here too, though, it seems that there are valuable insights to be derived from considering the way animate systems are designed to manage these issues. Take, for example, the "engineering" problems raised by the need to integrate knowledge from many sources in a large system. There has been some interest in viewing this problem in terms of the interactions that take place among the constituent knowledge modules or "experts" [Lenat,1975],[Hewitt,1977]. This perspective raises many interesting issues about the types of communication mechanisms and control structures that allow the interactions to result in intelligent behavior; and, it is clear that many of these same issues are raised when considering the brain as an interconnected set of neural elements. The specification of a control structure is also a problem during the design of a production system architecture. MacLaren [1978] has shown how biological examples can suggest practical organizational and control structures for production systems - structures that would not be obvious from a nonbiological point of view.

One of the most important aspects of the way artificial intelligence systems solve problems is that they rely heavily, if not exclusively, on search.

Solutions are usually sought by creating symbolic expressions and modifying them sequentially until they satisfy the conditions for a solution. Hence, symbol systems solve problems by searching.... They exercise intelligence by extracting information from a problem domain and using that information to guide their search, avoiding wrong turns and circuitous bypaths.

[Newell and Simon, 1976, p. 126]

This description of the search process is analogous to the way one might describe an animate system as it finds its way through a complex environment - moving sequentially and avoiding noxious stimulation as it seeks its goal. It has been suggested that there is an important link between such overt exploration of the environment and the "exploratory" search process used in problem solving [Kaufmann, 1979].

### Functional Constraints on Processing

Thus it can be argued that even knowledge engineers can profit from certain general considerations about natural systems and their environments. But how can one begin to formulate some of these general principles in ways amenable to a computational implementation? Some have argued that men and machines process information in ways that are so fundamentally different, such a task is hopeless. Dreyfus [1972], for example, argues that while machines deal with discrete symbols and isolated facts, humans process information in the context of a "situation"; in his words,

The basic insight dominates these discussions that the situation is organized from the start in terms of human needs and propensities which give the facts meaning, make the facts what they are, so that there is never a question of storing and sorting through an enormous list of meaningless, isolated data. (p. 174)

Humans have goals, experiences, and expectations that continuously interact with and are modified by the current stimulus configuration. In particular, humans can bring to

bear a seemingly endless amount of common sense know-how to a problem, above and beyond whatever task-specific knowledge is at their disposal. This constitutes a powerful context that, according to Dreyfus, is not realizable by a system using discrete symbolic representations.

The importance of such a context has not been ignored by those working in artificial intelligence however. Minsky [1975], for instance, observes

...that the ingredients of most theories both in artificial intelligence and in psychology have been on the whole too minute, local, and unstructured to account - either practically or phenomenologically - for the effectiveness of common sense thought. The 'chunks' of reasoning, language, memory, and 'perception' ought to be larger and more structured, and their factual and procedural contents must be more intimately connected in order to explain the apparent power and speed of mental activities. (p. 211)

Minsky then proceeds to describe a *frame*, a type of data structure that would be useful in maintaining information about Dreyfus-like situations. There has also been work trying to understand how knowledge about situations can be used to plan actions and infer their consequences.

Artificial intelligence researchers have recognized from the beginning that human reasoning processes are not adequately expressed by simple logical formalisms. Winograd's [1980] review of the work in this area shows that considerable effort is being invested in capturing the common sense aspects of human "logic".

Attempts to incorporate information-rich contexts, common sense knowledge, and common sense reasoning into



artificial intelligence systems have just begun. The progress to date seems to blunt Dreyfus' extreme claims about what machines inherently can't do. Nevertheless, many of Dreyfus' insights into human information processing ring true. Generic, prototypical information about situations helps to organize experience and generate efficient, orderly, consistent behavior. Yet the idea of a structure in which reality fills in the blanks seems to overlook a crucial point. Situations change and the system will quite often be in need of a new and more appropriate frame. It is not clear where new frames will come from, how they are selected, or how the transition from one frame to another takes place.

In the same way, it seems misleading to explain the continuity evident in human behavior by proposing that successive frames share the same "terminals". Residual elements of the previous situation certainly influence the representation of the current one; however, they are not the only elements that play a crucial role. Just as important are the defeated remnants of unsuccessful alternative representations and the hopeful precursors of the next situation. The representation of the current situation emerges as a variant of the previous one; yet it is a variant tempered by competition with alternatives and primed by vague expectations for the future. This ebb and flow of constraints is what Dreyfus describes as the influence of the "fringe". He indicates that this point of

view is not a mere metaphor for what's going on. It is a powerful, though little understood, strategy for processing information. Moreover, it is a strategy that is difficult - if not impossible - to realize using mechanisms that only replace and reorder symbolic expressions.

These characterizations of human information processing have not been overlooked by psychologists. They were eloquently expressed long ago by William James [1892]:

It is, the reader will see, the reinstatement of the vague and inarticulate to its proper place in our mental life which I am so anxious to press on the attention.... Every definite image in the mind is steeped and dyed in the free water that flows round it. With it goes the sense of its relations, near and remote, the dying echo of whence it came to us, the dawning sense of whither it is to lead. The significance, the value, of the image is all in this halo or penumbra that surrounds and escorts it, - or rather that is fused into one with it... (p. 32)

James then proceeds to describe this "fringe" in terms of general mechanisms that determine how activity flows through the brain. Lashley [1951] also points out the importance of understanding information processing in the brain in terms of widespread patterns of activity. These patterns are not the result of activity in discrete, isolated circuits or pathways. Instead, they arise as the statistical outcome of many neuronal interactions. While there is much that can be explained by considering each circuit as a discrete entity, there is also much to be learned by studying the emergent properties of their interactions. Lashley concludes that "only when we can state the general characteristics of this background of excitation, can we understand the effects of a

given input."

### Summary

If intelligent behavior is viewed as an adaptation to the task environment, many of the issues in artificial intelligence can be rephrased in terms of problems handled routinely by animate systems in their ordinary commerce with the real world. From this point of view, the functional nature of a system's transactions with its environment are more informative than the symbolic nature of those transactions. The crucial point is the role of the environment.

... It does not make sense to talk about adaptation without something to adapt to. And if one designates intelligence as an important mode of adaptation, then intelligent behavior has to be viewed in terms of environmentally posed problems. [Charlesworth, 1976, p. 150]

Finding the structural requirements for intelligent behavior requires a clearer understanding of environmentally posed problems and an understanding of how a system can be organized to solve those problems.

This dissertation will develop these ideas into a computational model of how knowledge can be usefully acquired and represented in a complex and uncertain environment. The mechanisms and overall architecture used in the model will be derived by considering the functional criteria for behaving in such an environment. It is hoped that this enterprise will yield new insights about how to construct artificial systems that exhibit intelligent

behavior.

## CHAPTER II

### RESEARCH STRATEGY AND SCOPE

The research issues to be dealt with in this thesis may be restated as follows. Natural systems that generate intelligent behavior are information processing devices whose design reflects the influence of evolutionary pressures. This means that the functional organization of a natural system is intimately related to the demands and constraints of the system's environment. It is the environment that embodies the criteria for whether or not a given system design is adequate. From the information processing point of view, perhaps the most important design criteria are those related to the informational aspects of the environment; namely, what kinds of information are available and what constraints there are on how it can be processed. Natural environments are complex, diverse, uncertain, and dangerous. These factors make information processing in natural environments a very formidable challenge. The problem is to understand how a system can accumulate and use knowledge about such environments in a way that makes efficient and prompt action - and hence survival - possible.

### Approach to the Problem

Given that natural environments embody the desired information processing design criteria, an obvious first step is to determine what those criteria are. Natural environments must somehow be characterized from an information processing point of view. General considerations like complexity and uncertainty must be expressed more concretely. What is needed is a set of dimensions that specify a broad range of hypothetical environments and tasks domains; and, that moreover preserve the relevant aspects of their natural counterparts. These hypothetical environments can then serve as a test bed for studying the behavior of proposed information processing models

It is not enough, however, to merely specify the demands an environment makes on a given system. Kaplan [1973, p. 64] points to the key issue with the observation that, as human beings, "... we are profoundly influenced by the environment, but in ways mediated by our sensitivities, our structures, and our inherited initial condition." The extent of the challenge faced by an information processing system depends in large measure on the structures and processes the system has to work with. The properties of a system - what kinds of information it can detect, how the information is represented, what it deems useful or important, and how it can affect the environment - are obvious constraints on its information processing

capabilities.

It is therefore also necessary to characterize in some way the functional mechanisms of an information processing system. Once again, the goal is to find a set of dimensions spanning a large class of diverse systems. Given characterizations of the environments of interest and the information processing models to be considered, it is possible to study systematically the relationship of a system with its environment. By judiciously changing the parameters of an environment, the demands of that environment can be adjusted to test the limitations of a given information processing mechanism. The success or failure of various mechanisms across the range of environments might then provide some clues as to which mechanisms are needed to function in a given type of environment. By studying and understanding these relationships, we can discern how environmental parameters establish criteria for information processing and how various mechanisms meet those criteria.

#### The Experimental Frame

Simulation is a useful tool for studying how a given model behaves in a given environment. The commerce of a system with its environment is complicated enough, however, to preclude any manageable study of all the issues at once. What is needed, in modeling terminology, is the specification of an *experimental frame* [Zeigler, 1976]. An

experimental frame delineates the circumstances under which a system is to be observed. This means designating a subset of all possible inputs from the environment as "relevant" to the system; and, specifying the system components, structural and procedural, that are under observation. Processes and structures that are not relevant to a given frame can be ignored. Parts that work together as a unit to achieve a given function can be lumped together. The resulting "lumped" model is simple and relatively easy to simulate. Moreover, since the simplifications are made only when they preserve behavior in the chosen frame, the lumped model is a valid way to simulate the more complex system. The task of the modeler is therefore to choose experimental frames judiciously so that simplification is not only possible but also informative.

For this study, the first step in specifying an experimental frame is to choose a task domain. There are several criteria for deciding whether or not a given domain is appropriate. First of all, the domain should be extensive enough to encompass as many "intellectual" functions as possible. It should also allow for tasks of varying complexity that can be simplified or complicated as needed. At the same time, the task domain should present the model with as few distortions of a natural situation as possible. Indeed, an underlying theme of this research is to understand information processing "... as it occurs in the ordinary environment and in the context of natural



purposeful activity" [Neisser, 1976, p. 7]. Finally, the domain should allow for the testing of models that are specified primarily in terms of their general architecture rather than their structural details. It is important that the models emphasize "... control structure and data flow rather than data structure" [Minsky, 1979, p. 18, Note 4]. The hope is, after all, to bring together organization and control principles that are applicable to both natural and artificial systems.

One task domain that satisfies all these criteria is the locating of resources in an environment by generating temporal sequences of actions. This is a domain that can require information processing capabilities ranging from simple reflexive behavior to sophisticated knowledge based problem solving. Consequently, it has the flexibility needed for a systematic study of environmental parameters and information processing mechanisms. The tasks are, moreover, clearly salient in terms of adaptation. An organism searching for resources must manage information about its environment, its needs, and its capabilities. It is in a Dreyfus-like "situation" in every sense of the word. Just as important is the fact that temporal integration is a crucial aspect of many intellectual functions. Planning, the use of language, and motor coordination are all examples of activities requiring some kind of serial ordering capability. Lashley [1951] has pointed out that "temporally integrated actions do occur even among insects, but they do

not reach any degree of complexity until the appearance of the cerebral cortex. They are especially characteristic of human behavior and contribute as much as does any single factor to the superiority of man's intelligence."

The approach taken here is therefore to design a computational model of a simple hypothetical organism, one that must find resources and avoid noxious stimuli in a carefully chosen simulated environment. Research on this and related problems has been done before. It is instructive to examine these past attempts, both to consider the merits of other approaches to the problem and to understand why the research has not sparked further interest.

#### Review of Related Studies

Much of the work on mechanical devices and robotics is not directly relevant to this study. Research in these areas has been concerned primarily with engineering problems as opposed to general issues related to intelligent behavior. There is one device, however, that does seem relevant here. It is one of the earliest and simplest of the robots; namely, Grey Walter's [1953] mechanical tortoise. In this system there are only two discernable stimuli: light sources and physical contact. These stimuli are used by a feedback mechanism that controls in which direction the device will move. The tortoise sustains itself by recharging its batteries when necessary;

therefore, the power level of the batteries is also monitored by the control mechanism. In spite of this very simple architecture, the tortoise is capable of generating very sophisticated looking behavior. It pushes small obstacles out of the way, finds its way around heavy ones, and avoids bright lights - approaching light, however, when its batteries need charging. Signals from the environment, together with the system's needs, serve as modulators of the internal state of the system.

What is so interesting about this device is its simplicity. Its knowledge of the environment is not stored symbolically in a data base; instead, it is embodied in the pre-wired structuring of the constituent parts. The behavior of such a system "... *for an observer* can be described in terms of representations, but ... can also be understood as the activity of a structure-determined system with no mechanism corresponding to a representation" [Winograd, 1981, p. 249]. There is no doubt that under more demanding circumstances, an information processing system cannot get by with this kind of inflexible, built-in structure. Yet the fact that such a simple mechanism generates such impressive behavior suggests that, even in a complex task domain, much of a system's knowledge can be stored implicitly in its initial structure.

There have been several studies about an organism's interaction with its environment that fall within the symbol manipulation paradigm. One such effort is described by

Doran [1968]. His work considers an automaton that must plan sequences of behavior in a very simple environment to travel from a given point back to a "nest". The motivation to return to the nest is provided by a "desirability" factor, directly perceived at each location, that the automaton tries to optimize. The environments under consideration are simple square areas with boundary and interior "walls" (see Figure 1).




Figure 2.1. The environment for Doran's [1968] automaton, shown on the left as a set of barriers and on the right in terms of stimulus information. The automaton is in its "nest" when it is facing and against the letter A. The desirability of any location is given by the formula  $50 - (d + 3*w)$  where  $d$  is the distance to the wall being faced and  $w$  is 1 if the wall is A, 2 if B, etc. (Adapted with the permission of the publisher)

Each wall is identified by a letter of the alphabet. In fact, all the automaton can detect is the letter identifying the wall it is facing, how far away that wall is, and the desirability of the location. The automaton deals with its

environment in terms of state transitions; that is, in terms of the changes observed in what it detects after applying a particular action. As it interacts with the environment, the automaton stores these transitions in a limited capacity memory. When a perceived state is "recognized" as something previously encountered, a plan of action is formulated. This means a "lookahead tree" of possibilities is constructed and a path through the tree is selected that optimizes desirability. When the automaton has no memory of a given perceived state or the planning process does not generate anything promising, an action is chosen at random.

Several "incarnations" of the automaton were tested repeatedly in the environment shown in Figure 1. Doran reports that "the automaton successfully uses its record of its past explorations to form and implement plans. These plans enable it to find its way to its nest by a much more direct route on the second or subsequent trials than that followed on the first trial" (p. 209). Moreover, the improved routes often included locations the automaton had never visited before. Experiments with the automaton model may therefore be deemed successful. It must be pointed out, however, that the environment is structured to make things easy for the automaton in at least two ways. First of all, the identifiers for the walls in the environment are a small set of discrete symbols. The symbols have no underlying structure or information useful for generalization and abstraction. The automaton is accordingly not equipped to

handle environments having countless numbers of unique states but a manageable number of meaningful patterns that can be inferred from experience. Because the number of symbols is small, the automaton can afford to store a fairly detailed state transition history of its experience. Secondly, each location in the environment provides the automaton with an explicit indication of its desirability. This means the automaton can do a calculation of expected desirability to help choose among prospective plans. Natural environments are not so cooperative. Midgley [1978] argues persuasively that it is very difficult to know what "really pays" in a natural environment. While humans certainly can and do perform payoff calculations, such calculations are not as important in making a choice as a genuine motivation to do one thing instead of another. These considerations raise questions about whether or not Doran's environment has what Brunswik [1956] calls "ecological validity".

Jacob's [1972] simulated creature PERCY is another example of the symbol manipulation approach. PERCY's task is to build nests in an two dimensional environment that contains food, barriers, landmarks, and nest building materials. PERCY is motivated by an internal satisfaction measure. Satisfaction decreases whenever PERCY is hungry or the nest building is progressing too slowly. The system makes decisions so as to keep its satisfaction at acceptable levels. PERCY's behavior is controlled by a hierarchically

arranged set of task oriented components. The top level component establishes a goal for the system. This triggers a sequence of plans to be activated that help to attain the goal. Each plan in turn specifies a sequence of behavioral situations it expects to occur and each situation influences what the system actually perceives. Jacobs claims that the organization of these components is a model useful for describing all kinds of purposive behavior.

PERCY's elaborate knowledge structure is curious because it does not contain any details about the structure of its environment. What it contains, instead, is an exhaustive account of every task related goal, plan, behavioral situation, and object that PERCY will ever need to know about. At each point in its behavioral cycle, PERCY essentially retrieves a description of what to do and what to look for and then follows directions. In this sense, PERCY is little more than an elaborate mechanism for generating instinctive behavior. It has been noted earlier that initial structure is an important influence on the behavior of this kind of system. It must be emphasized, however, that flexibility is a key aspect of the intelligent behavior of natural systems; and, such flexibility cannot be realized by a system that does not use information about the structure of its environment.

The most comprehensive simulation to date is probably the system described by Findler and Allan [1973]. In their environment objects are characterized by features. These

features include visual and auditory stimuli, as well as certain motivationally coded information such as the potential of an object as food, shelter, or a source of danger. Some of this information is transient; that is, objects may suddenly appear or disappear and food sources eventually become exhausted. The impact of any given object on the organism is determined by the results of a feature analysis. An attention-scanning heuristic selects one of the detected objects to focus on, transfers certain information to long term memory, and modifies the overall "goal-state" of the organism. A planning mechanism generates the organism's behavior so as to attain the indicated goal.

The memory structure for this organism is a dynamic network of nodes and arcs. Each node represents a previously encountered object and contains the results of the feature analysis. A node may also contain a detailed list of the properties of the object. Such detail, however, is obtained only via physical contact with the object. Each arc represents the plans and actions previously used to travel from one object to another. In this way memory contains a practical map of the known terrain. The map is improved under the guidance of several learning schemes which decide when a new node is added, when an arc can be made more efficient, and when the potential exists for adding a useful new arc.

Overall, the behavior of the organism is very



impressive. In two test environments it exhibited an ability to explore and become familiar with the terrain, find food when necessary, avoid obstacles or dangers, and improve its efficiency as it gains experience. Yet there is much about the approach to designing the system that seems ad hoc. Consider, for example, the way the environment provides information about objects and their properties. When the organism detects an unfamiliar stimulus pattern, it searches its memory of previously encountered patterns for one that is similar. The organism infers, reasonably enough, that the current object has properties similar to those associated with the pattern stored in memory. There is no meaningful relationship, however, between a stimulus pattern and the properties of an object. Only by direct physical contact with an object are its properties somehow mysteriously revealed. The patterns themselves convey no information at all about the nature of their source. They are mere labels to be used to index the memory structure. This is a severe distortion of the relationship between an organism and its natural environment. Not only are stimulus patterns the only source of information available to real organisms; they are a diverse, complex, and equivocal source of information. Perception under these circumstances is an achievement [Hilgard, 1978]. The structures and mechanisms required by this task provide the framework for all subsequent processing and behavior. Findler and Allan have completely ignored this issue. Consequently, their model

does not address one of the most basic requirements of functioning in a realistic environment.

Some organism/environment simulations have been organized more along psychologically motivated lines. One such approach is found in the model described by Plum [1972]. The environment is a one dimensional runway with an object and light at one end and a shelter at the other end. Sometimes the object delivers food to the organism, other times it delivers pain. The change in state of the object is signaled by a change in the color of the light. The organism is motivated to seek food when it is hungry - that is, when its resource reservoir falls below a certain level - and it must therefore learn which sensory cues indicate food is present. The organism is constructed using principles from the cell assembly theory first described by Hebb [1949]. More specifically, the system processes information using a network of "aggregated cell assembly" units. Each unit in the network has a specialized job: either detecting stimuli or rewards, recognizing situations, predicting sequences of events, or initiating action. The state code for each unit is activity. The level of activity indicates how much evidence a unit has accumulated from other units and/or the environment. All of the units associated with a given function in the system compete with each other for control of that function. In this way, system behavior is generated by the set of units with the highest levels of activity. The types of units in the

system and their interconnections are specified in advance. It is possible, however, to modify the influence one unit has on its neighbors.

Limited simulations of the model generated somewhat disappointing results. The organism never learned where to eat or under what circumstances food would be available. One obvious reason for this failure is the inability of the system to modify its structure in any significant way. There is no way to add a useful new connection or delete an unnecessary old one. Similarly, there is no way to create new units or modify existing units. When the model the system is given proves inadequate for a given environment, the system is doomed to fail. Another difficulty is in the way the system/environment interface is specified. The only things the organism can detect about the object are if it is large or small, indicating that it is close or far away; and, whether or not it is emitting an odor signaling that it will deliver food. The organism has no way of figuring out where the object is or when being close to the object is close enough to eat. The task of locating an object, even in a one dimensional environment, is next to impossible without information of this kind.

A powerful learning algorithm is one of the impressive aspects of the model presented by Holland and Reitman [1978]. Here again the environment is one dimensional, this time with resources at each end. The organism has two needs, one of which is satisfied by visiting the left

resource, the other being satisfied by visiting the right resource. In order to be successful, the organism must learn to move through the environment in a way that keeps both needs under control. The basic unit of structure in the organism is the *classifier*, a condition-response rule that is sensitive to a set of signals from the environment. Any given signal will usually satisfy the conditions of several classifiers in the system. The choice of which of these classifiers will generate behavior is made stochastically, based on each classifier's utility value. Estimates of a classifier's utility for finding a given resource are refined over time; that is, each time a need is fulfilled the credit for the success is apportioned among all classifiers that generated behavior along the way. The overall utility of the system's classifiers is improved via the genetic algorithm [Holland, 1975]. This is a general purpose learning algorithm used here to generate new classifiers that will be more useful to the system than the old ones.

Two experiments were run to demonstrate the potential of this classifier system. In the first experiment the system was placed in a one dimensional space with seven nodes, each labeled with a randomly chosen eight bit signal (see Figure 2). The directed arcs between adjacent nodes were marked with a zero or one, chosen at random, indicating the response the system must make to traverse the arc. There was twice as much of the right resource available as

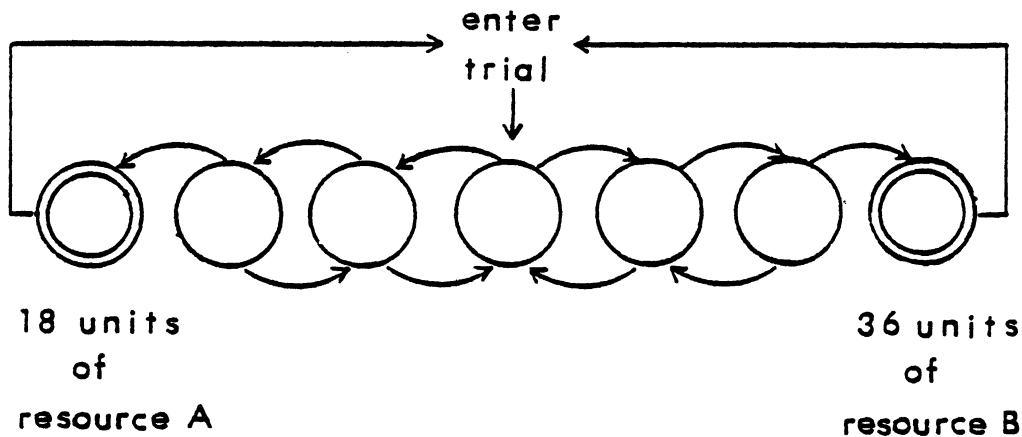


Figure 2.2. The environment for Holland and Reitman's [1978] classifier system. (Adapted with permission of the publisher)

there was of the left resource; therefore, trips to each end must have the corresponding two-to-one ratio in order to simultaneously satisfy both needs. The organism started at the middle node and moved until it had satisfied the most pressing need. It was then returned to the middle node and the cycle was repeated. After a short time the organism performed considerably better than random and eventually learned to consistently make the optimal trip. Given this success, the second experiment was designed to show the ability of the system to use experience. Three new nodes were appended to each end of the original environment and the resources were moved to the new endpoints. An organism that had mastered the smaller environment was placed in this extended one. Having previously experienced the embedded nodes proved to be a tremendous advantage. The experienced

organism learned the environment far sooner than organisms with no prior experience.

Any analysis of the capabilities of this system must be tempered by the observation that it performs in a very impoverished environment. There are no objects, no sense of continuity, distance, or space - properties that are so characteristic of natural environments. This means, in particular, that the system can afford to rely on simple stimulus-response pathways. Moreover, the algorithm for apportioning credit among classifiers responsible for successful behavior is much too coarse. A record is kept of all classifiers active during the search for a resource. The relative contribution of any particular classifier is determined only at the end of the search. In a more realistic environment this would prove to be impractical. Not only would a large number of classifiers be involved in the generation of a single response; there would also be long intervals between the receipt of tangible external rewards. If the system's goals change before a reward is obtained or the system has more than one goal operative at a time, it is more difficult to isolate the contribution of any particular classifier once a behavioral sequence is over. All these concerns point to the need for a more "local" way of evaluating a classifier's activity. Nevertheless, the potential exists in the classifier system framework for more sophisticated models. It is possible, for example, to envision a system in which the response of

one classifier satisfies all or part of the condition for various other classifiers. The system would then be structured much like an associative network. Though such an organization can be theoretically achieved using the genetic algorithm, no system has been constructed to prove that it is practical to implement.

### Summary

Several issues emerge from the above discussion. First there is the importance of choosing an environment with simple but natural characteristics. Mistakes in this area can lead to system designs that are ad hoc and/or irrelevant to the problems faced by natural systems. This underscores the importance of specifying an experimental frame and justifying as much as possible all simplifying assumptions. Moreover, the relationship between the system structure and the structure of the environment should be considered from the beginning. It has been noted that much of the intelligence exhibited by an information processing system resides in the architectural design of that system. The functional characteristics of the design can be understood only if the environment tests it in appropriate ways. This suggests that the most effective simulation strategy is to start with simple models. The environment can then be made more complex in small increments and the system can be elaborated only when a change is necessary to obtain some desired behavior. In this way, a complex design can unfold

gradually and so make its structure comprehensible.

Another important concern that stands out is the question of learning. It is clear that for most natural systems, information acquisition is just as important a job as information handling. When - as is usually the case - that information is uncertain or there is simply too much of it, an information processing system must take the learning and inference issues seriously. Holland [1975] has identified two objectives in processing uncertain information. New information must be examined whenever possible so as not to overlook something that might be important. On the other hand, old knowledge needs to be continually reevaluated and confirmed because it is steeped in uncertainty. Clearly, in order to be successful, a system must pursue both objectives. Obtaining the requisite information, however, is not always easy. Though an organism is swamped with information in a natural environment, salient information is scarce [Kaplan, 1976]. In a very real sense, therefore, an organism must be an active processor of information. It must seek out that which is new and/or relevant. Maintaining a knowledge base that is reliably useful is a never ending task in an uncertain environment; and, learning plays a vital role in the process.

The computational organism model to be developed here will therefore be an adaptive system. It must be emphasized that developing such a model is a very difficult challenge.



People talk fondly of computer programs that will start with some fundamentals and acquire all the knowledge needed by some natural sequence of learning, experiencing the environment in which it must function. Very little effort gets spent studying what it would take to accomplish this, perhaps because there is implicit realization that the task is harder than it might seem.  
[Norman, 1981, p. 284]

The first step toward realizing this goal is to take a closer look at the kinds of information available in a natural environment. Then the challenges facing the adaptive system - and the system's designer - will be clearer.

## CHAPTER III

### A SIMPLE MODEL OF INSTINCTIVE BEHAVIOR

The task at hand is to identify and analyze some of the fundamental issues underlying an organism's commerce with its environment. Given that one of the premises of this research is that natural information processing mechanisms are more functionally determined - reflecting the adaptive demands of the environment - than they are logically organized [Lachman and Lachman, 1979], it makes sense to consider first the relevant properties of natural environments. The properties most pertinent to this study are those that convey information about the identity and location of objects. It must be shown that the kinds of information available in natural environments, together with constraints on how the information can be used, establish functional criteria for information processing. A class of simulated environments will then be devised that embody many of those criteria. Subsequently, a hypothetical organism relying exclusively on innate structures to generate behavior and capable of surviving in the simulated

---

'Logical, that is, in terms of strict adherence to the rules of some known mathematical formalism.

environments will be designed and tested.

It should be pointed out that this approach to organisms, environments, and information is based on the philosophical assumption of *realism*: that there is an objective world with real properties that exist whether or not they are perceived or thought about [Shaw and Bransford, 1977]. The environment, in other words, has a physical nature which places constraints on the patterns of stimulation available to an organism. It makes sense from this point of view to study the properties of an environment and ask how those properties might be supportive of or a hindrance to various information processing activities. This is not to say, however, that there is not

...a necessary mutual interplay between stimulus and organismic properties. An organism cannot engage in pattern recognition, for example, based on a feature analysis unless there are in fact features in the stimulus to be analyzed. On the other hand, there is no need to attempt pure stimulus descriptions in terms that are inappropriate to the processing organism. So the properties of the organism limit the properties of the stimulus to which we pay attention; at the same time, the properties of the stimulus limit what the organism can do with the stimulus. ... In other words, properties of the stimulus provide a limiting condition for the processing organism, but at the same time, the stimulus properties do not completely determine mode of processing at all. [Garner, 1978, p. 101]

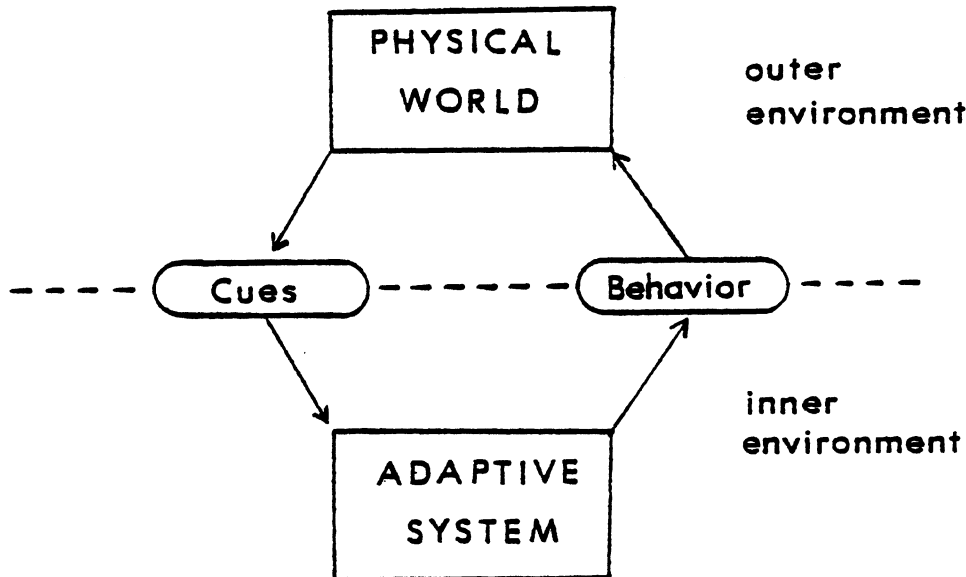
The environment is seen as a source of information and an organism comes to "know" the environment by sampling and processing that information. Knowledge of the environment is indicated by some reliable correspondence between the "psychological" states of the organism and the informational

states of the environment [Shaw and McIntyre, 1974]. This means that the extent to which some set of properties is informative will vary depending on the organism in question.

### The Organism/Environment Interface

It is important to emphasize, therefore, that the organism and environment must be considered in terms of their interactions. For the purposes of this research, intelligent behavior is viewed as an adaptation to the demands of an environment. An organism is accordingly thought of as an adaptive system. Simon [1969] has noted that, from a designer's point of view, an adaptive process is most conveniently thought of in terms of three components: the goals or functions to be realized by successful adaptation; the *inner environment*, or structural details of the adaptive system; and, the *outer environment* or surroundings in which the system behaves. This functional point of view has two potential advantages. First, the inner environment and goals together may specify a functional environment that involves relatively few of the many details of the outer environment. On the other hand, the outer environment and goals may constrain the set of appropriate behaviors to such an extent that the behavior of the adaptive system can be predicted without having extensive knowledge of the inner environment. A good design yields both advantages. It characterizes "... the main properties of the system and its behavior without

elaborating the detail of either the outer or inner environments" [Simon, 1969, p. 9].



**Figure 3.1. The interaction between an adaptive system and its physical surroundings.**

These considerations suggest that it will be most useful to consider natural environments as part of a dynamic relationship as shown in Figure 3.1. An adaptive system interacts with an environment in two ways. First, it continually samples the information available about the current environmental state and state transitions. This information from the environment defines the opportunities for adaptation and behavior available to an adaptive system; specifying, in conjunction with the sensory and motor apparatus of the adaptive system, an environmental niche. What information permeates the interface depends, of course,

on the receptor orientation and behavioral disposition of the adaptive system. Second, the environment can be acted upon by the behavior of the adaptive system. This can result in a change in the set of cues being sampled by the adaptive system and/or a state change in the environment. The relationships depicted in Figure 3.1 emphasize how the interface between an adaptive system and its environment - as determined by the environmental niche and by the functional goals of the system - defines and constrains all interactions.

What characterizes the way the interface mediates the flow of information from a natural environment to an adaptive system? The most basic distinction to be made is between the sources of stimulation and the stimulus cues themselves. Brunswik [1956] and others refer to the former as *distal* objects and the latter as *proximal* stimuli. Distal objects are remote from an adaptive system in the sense that they are defined without reference to any behaving organism. Proximal stimuli are the sensory cues available at the system/environment interface. The distinction is made clearer by the following example:

Physically, this page is an array of small mounds of ink, lying in certain positions on the more highly reflective surface of the paper.... But the sensory input is not the page itself; it is a pattern of light rays, originating in the sun or in some artificial source, that are reflected from the page and happen to reach the eye. Suitably focused by the lens and other ocular apparatus, the rays fall on the sensitive retina, where they can initiate the neural processes that eventually lead to seeing and reading and remembering.  
[Neisser, 1967, p. 3]

The point is that the adaptive system can sense the physical world only in terms of the information available through sensory cues.

### Cues and their Implications

This distinction is perhaps obvious, but it is too often overlooked when considering how to structure a task environment. For example, most of the studies reviewed in Chapter 2 designated entities in the environment as static, non-decomposable symbols. While the cues can indeed be thought of as symbols, the notions of symbolic information prevalent in computer science and information theory differ markedly from the concept of "natural" information that is salient to an active, purposive creature [Shaw and Bransford, 1977]. The cues available in natural environments have meaningful structure. Gibson's [1966] analysis of natural stimuli leads him to conclude "... that the available stimulation surrounding an organism has structure, both simultaneous and successive, and that this structure depends on sources in the outer environment" (p. 267). A similar observation was made by Tinbergen [1951] in his study of instinctive behavior. The stimuli that "release" activity in an innate behavior mechanism as a rule tend to be structural; "that is, it is the arrangement of elements in the visual field, in space and in time, that releases behavior, the elements themselves being the same in the releasing and non-releasing situations" (p. 78). Because the

cues are structured, and their structure is indicative of the source of stimulation, an organism is rarely forced to deal with isolated bits of meaningless data. The relationship between an organism and its environment is mediated by useful information.

Natural environments offer organisms a seemingly inexhaustible supply of such information [Gibson, 1966; Kaplan, 1976]. Even in a single situation, the cues are so diverse and their number is so large that an organism cannot hope to attend to each one individually. To do so would run the risk of being overwhelmed with processing all the cues without ever getting around to making an overt response. Moreover, the stimulation available at any given instant of time is likely to be unique. An organism constrained by a limited storage capacity or a finite processing time clearly cannot afford to treat each cue as a distinct entity; indeed, knowledge of such detail is, from a functional point of view, probably not even important [Bartlett, 1932]. It is therefore unlikely that an organism can rely on mere accrual as a strategy for picking up information. Natural environments demand some degree of discrimination, selectivity, and information reduction.

Another important characteristic of cues and their role in the organism/environment relationship is the equivocality of any particular cue. Brunswik [1956] emphasized that there is seldom a one-to-one, perfectly correlated



relationship between a given cue and a given source of stimulation.<sup>3</sup> In a natural environment the relationship between proximal cue and distal object is fallible and uncertain; in other words, the correlation between cue and object is typically less than one. There also tend to be correlations of less than one among various cues for the same object. This is not to say, however, that the natural environment is chaotic. Kaplan [1978] points out that

Although the environment is uncertain, it is by no means random. Regularities abound, and the organism must identify them. A lion, for example, presents many regularities. A lion has teeth (unless it is a very old lion), a mane (unless it is a female lion), impressive stature (unless it is a cub), and a tail (unless, of course, something happened to it). Lions also roar. Sometimes. But quite apart from the variability of lions, there is the variability in the way one happens to observe them. A side view provides certain information, a back view looks different, and an eyeball-to-eyeball view looks very different indeed. Yet to a potential lion prey, the appropriate action may be the same.

The variability in stimulation from the same objects has still other sources: the background or setting in which it appears can vary; the foliage may obscure portions of the animal; and so on. (p. 30)

Natural environments therefore require that an organism somehow manage to respond to the regularities hidden behind the equivocal nature of cues.

It is of course possible to reduce the uncertainty

---

<sup>3</sup>The Gibsonian position would argue that, in principle, there exists higher order unequivocal information about distal objects. While in some cases such stimulus information is indeed potentially available, it is not clear when or if the information can be used effectively by a perceiver [Hochberg, 1974].

involved in any given situation. A careful examination of all the salient cues, together with precise measurements and calculations, will eventually lead to an unequivocal ascertainment of the underlying regularities. Such a strategy, though, requires at the very least some time for locomotion and manipulation, and an extensive knowledge of the relevant physical laws [Brunswik, 1956]. This is clearly not feasible. Not only is an organism unlikely to have knowledge of all physical laws, it is also unlikely to have the time to do any extensive analysis. Natural environments characteristically have several properties - such as unstable resources, predators, and competitors - that generate a direct adaptive advantage for the speedy processing of cues. Accordingly, the only viable processing strategies are those that are fast, avoid drastic errors, and maintain a reasonable correspondence with the environment [Kaplan, 1978]. The information processing structure and behavior of organisms should therefore be considered in light of their advantages for survival and reproduction in the face of competition and a changing environment. These advantages can be studied directly with the use of optimization or cost-benefit analysis models designed to show an adaptive fit between structure and function [Smith, 1978].

The aforementioned challenges posed by the nature of proximal cues severely constrain the ways information can permeate the environment/adaptive system interface. It has

long been known [James, 1892; Lashley, 1942; Campbell, 1966] that animals and humans meet those challenges by being sensitive to patterns or equivalence classes. The word "pattern" refers to the dimensional and correlational structures that exist in the environment [Garner, 1974]. Patterns provide an economical and reasonably accurate summary of a potentially infinite set of information. They are discernable at several levels, from the informational structure characteristic of a given kind of cue to the simultaneous and successive relationships found in groups of cues. The regularities underlying the information available in natural environments establish a stable framework for information processing. Without such regularities, adaptation and intelligent behavior would be impossible. Bruner [1957, p. 42] notes that recognizing patterns "... represents the simplest form of utilizing inference. It consists of learning the defining properties of a class of functionally equivalent objects and using the presence of these defining properties as a basis of inferring that a new object encountered is or is not an exemplar of the class."

An important principle for the design of an artificial task domain is therefore that the stimulus information have some underlying structure that can be used to generate equivalence classes. Gibson [1966, p. 40] lists three characteristics of natural stimuli that summarize this idea:

- 1) A stimulus has nontrivial properties that characterize its structure in space. It is in

particular not an isolated symbol or mathematical point.

- 2) A stimulus has a structure in time that is lost if the stimulus is isolated in some mathematical instant.
- 3) A stimulus is associated with some components that change, and other components that are invariant.<sup>3</sup>

If a complex task domain provides stimulus information having these properties, then an adaptive system can cope with the complexity in the environment by being sensitive to patterns. Brunswik [1956] used the term *distal focusing* to describe this strategy of generalizing over the variants in the proximal stimulation. The term is appropriate because, by discerning the underlying patterns, an adaptive system focuses beyond the immediate stimuli to establish a stable relationship with the distal source. It should be emphasized that, by being sensitive to patterns, an adaptive system does not completely eliminate the problems associated with complexity and uncertainty. The problems are merely constrained to more manageable proportions.

### Spatial Information

#### The interface between the inner and outer environment

---

<sup>3</sup>Gibson argues that a natural stimulus always has invariant components, or high order properties common to all instances of a particular entity. Several studies of "natural categories", however, indicate it is very unlikely that such invariants exist in the stimulus itself [Bransford, 1979]. Rather, the invariants are residues of the perceptual processes of the observer.

not only affects the quality of information available to an adaptive system; it moreover defines what information is potentially salient. The interface of the proposed adaptive system specifies functional goals involving locomotive behavior; namely, locating resources and avoiding noxious objects. These goals can be realized only if the environment provides spatial information in a way that makes locomotion possible. Tinbergen [1951] makes a useful distinction between two kinds of information required for locomotion: *releasing stimuli* that indicate a particular kind of behavior is appropriate; and, *directing stimuli* that orient the behavior with respect to the relevant spatial characteristics of the environment.

This difference can be made clear by an analogy. The movements of a steamship are dependent on two mechanisms. The propeller pushes it, the rudder steers it. The forward movement is dependent on a releasing stimulus and can go on without further external stimulation as long as the fuel does not run out. The steering, however, is continuously controlled by new stimuli, coming, eventually, from the external environment.... The releasing stimuli that are responsible for the forward motion of the ship are entirely different from the directing stimuli. The sailing order releasing the departure, the open ocean releasing increase of velocity to full speed, or slowing down the speed, &c., provide stimuli controlling the propeller mechanism. Visual and magnetic stimuli (the latter received by the compass, the ship's magnetic receptor) continuously control the rudder mechanism. (p. 83)

The releasing stimuli for the task domain in this study are the cues associated with resources and noxious objects. The nature of such cues has been discussed above. What remains to be considered is the nature of directing stimuli and the

spatial characteristics they indicate.

The most fundamental kind of spatial information required for locomotive behavior is information indicating the location of a source of stimulation. Tinbergen [1951] points out that localization involves ascertaining the direction and/or the distance of the stimulus source.

These two principles of spatial analysis of the environment, viz. localization of direction and of distance, are of great importance for the understanding of the influence of the environment on behavior. First, spatial analysis enables an animal to 'recognize' objects. Further, it enables it to localize objects in relation to other parts of the environment and thus perform oriented movements, that is, movements directed in relation to spatial patterns outside the animal. (p. 25)

The "objects" that are to be localized need not be concrete and specific. Directing stimuli in natural environments include cues such as the direction of gravitational and magnetic forces, the polarization of light, as well as distinctive places and physical objects.

The key issue is what the relevant spatial patterns in the environment are and how they can be assessed by an adaptive system. That is what determines the degree to which an adaptive system must process detailed location information. Fraenkel and Gunn [1961] cite several examples illustrating the variety in spatially oriented locomotion strategies: Consider, for instance, the woodlouse which can only survive in very moist places. By changing its speed of locomotion, moving rapidly in low humidity and slowly or not at all in high humidity, the creature manages to spend most

of its time in the most suitable regions of the environment. The spatial knowledge required for this behavior is minimal. The interface between the inner and outer environment specifies a reliable cue that is directing only in a generic sense. Each individual locomotory reaction is undirected, but the overall behavior is in effect oriented toward moist regions.

A different use of location information is exemplified by creatures that must maintain a general direction of motion that is not straight towards or away from the directing stimulus itself. Ants, for example, use the sun along with other cues to return to the nest. One proposed explanation for their behavior is as follows. By simply maintaining the image of the sun on a fixed ommatidium in one eye, the ant can move at a constant angle with respect to the sun. On returning to the nest, the ant need only keep the sun's image on the "inverse" ommatidium in the other eye. This so-called *light compass reaction* has also been observed in the honey bee in situations where there are no visible landmarks. As long as the trip is quick enough so that the sun has not moved too much, the compass reaction is an effective way for a creature to maintain a general orientation in an uneven terrain.

Accurate estimation of distances not previously traversed or encountered seems to depend almost exclusively on visual stimulation for most organisms [Tinbergen, 1951]. One of the rare known exceptions is the use of sound waves

for "echo location", a capability used by relatively few creatures such as bats and cetaceans<sup>4</sup>. The ability to accurately judge distance is of course important if an organism is to acquire any three-dimensional knowledge of an environment. For simple locomotive behaviors, however, only certain rough estimates of distance seem necessary. For example, an organism must at least be able to detect when it is close enough to food to consume it. It might also be important to know roughly when an object is "in range" or "getting closer" so that other behaviors - such as escape, changing speeds, or paying closer attention - might be triggered when appropriate. Needless to say, the information making these distance judgements possible must be available through the organism/environment interface.

Aside from indicating location, another potential important function of a directing stimulus is to serve as a landmark. A distant landmark that is stationary or that moves in predictable ways can be used for orienting as described above. A higher order function, however, is implied when an object, place, or region serves as a choice point for a change in direction. An organism that uses landmarks as the locus for making a choice must store explicit information about the choices and their associated outcomes. This is a direct use of spatial knowledge about a particular environment. What kinds of entities are most readily used as landmarks? Kaplan [1976] points out that

---

<sup>4</sup>i.e. Whales, dolphins, and porpoises.



distinctiveness is an important issue, and that distinctiveness can arise from at least three factors. First, from sensory qualities that can be easily discriminated from the immediate environment, such as the visual distinctiveness of a mountain surrounded by a flat plain; second, from sensory qualities that can be discriminated using the experience of the organism. "An oak tree, for example, can serve as a landmark to the individual who knows that all the other trees in the area are maples" [Kaplan, 1976, p. 43]; and, third, the existence of some functional relationship with the organism's activities that causes an entity to be frequently encountered. For instance, an otherwise non-distinctive location which offers an unrestricted view of some goal could be a valuable landmark.

An environment can either facilitate or hinder the use of landmarks, depending on how distinctive the individual objects and locations are. An equally important factor, however, is the overall organization of elements in the environment and their interrelationships. Consider, for example, the foraging strategies of various ant colonies and how they are determined by the spatial and temporal distribution of resources [Hölldobler and Lumsden, 1980]. African weaver ants forage in environments that have uniformly distributed, continuously renewing resources. The territory for a colony is defined simply by the amount of area a worker can cover in a relatively short period of time

from the nest. The use of landmarks is not implicated in the worker's search for food. A random exploration in all directions is the cost-effective strategy for the ant colony. Some species of harvester ants, on the other hand, exploit resources that tend to be patchily distributed. Foragers from these colonies travel along well-established trunk routes to known patches of food, and from there perhaps wander off in search of new supplies. These trunk routes are well marked by visual and chemical cues that persist over very long periods of time. The harvester ants, in other words, use landmarks to reliably find their food supply.

If an organism has to learn the relevant "spatiotemporal" patterns of the environment through experience, it is helpful if the patterns are organized coherently so that they can be readily discerned. Lynch [1960] uses the term *legibility* to refer to the ease with which experience with an environment leads to knowledge of its structural patterns. A legible environment is one containing easily identifiable units and well defined relationships among them. Units are grouped together in regions according to their relationships, all in the context of some overall organizing construct. A well structured environment can serve as a frame of reference in which locomotive behavior and spatial knowledge can be organized. "The environment suggests distinctions and relations, and the observer - with great adaptability and in the light of

his own purposes - selects, organizes, and endows with meaning what he sees" [Lynch, p. 6]. The more legible an environment is, the more readily an organism can learn to find its way through it. A chaotic or "information poor" environment can make navigation extremely difficult or even impossible. Efficiency is not the only issue here. Survival could be at stake as well. In humans, for example, the need to experience order in the environment is closely tied to feelings of security and well-being. Lynch notes that, consequently, experts who have learned to navigate unaided in more or less "featureless" domains such as the open sea still experience strain and anxiety during every trip.

A legible environment can also help an adaptive system obtain certain generic information useful for the control of locomotion. The control of motor behavior requires, among other things, an assessment of the possibilities afforded by each alternative action [Gibson, 1979]. An obstacle affords a collision, an "opening" affords unobstructed movement, a resource affords a positive experience, and danger or noxious elements afford a negative experience. These factors, once perceived, can be used by an adaptive system to help decide on the most appropriate course of action.

#### Implications for an Artificial Environment

The above discussion is only a brief review of the many issues relevant to understanding how information about the

identity and location of objects can be made available. It would be extraordinarily difficult, and computationally expensive, to design and implement a simulated environment having all of these properties. Consequently, some attempt must be made to identify some small subset of basic issues most pertinent to this research. Four aspects of an organism's commerce with a natural environment summarize what is of interest here. First, the organism/environment interface makes available to the organism only a limited sample of all the information potentially available in the environment. Second, there is a meaningful and functional distinction between proximal cues and distal objects. Third, the cues available from any given object are likely to be highly diverse and variable. Finally, there are structural regularities underlying the object/cue relationships that have potential functional significance to the organism.

Even given this reduction in scope, there are still several parameters involved in designing an environment with realistic informational properties. Consider, for example, some alternatives relating to object/cue relationships. The relevant cues "defining" an object can have varying degrees of accessibility. All the cues might be accessible all the time; or, perhaps, only one of any number of different subsets of those cues might be accessible due to the vantage point of the observer, the inherent variability in the object class, etc. In some dynamic situations, the cues

might even be dispersed over time. There is also potential diversity in the ways cues can be correlated with objects. A given cue might be perfectly and exclusively correlated with a given object. Detection of the cue in that case is a reliable indicator of the presence of the object. Alternatively, a cue might be correlated with several objects to varying extents, or even with other cues. Determining object identity under these circumstances is a much more challenging process. Values along these dimensions obviously impact the degree to which an environment is considered complex or uncertain.


Given the broad range of choices regarding environmental parameters, where is one to begin in the design of an artificial environment? The review in Chapter 2 of previous work points out a consistent failure to specify functional differences between cue and object in the environment. Only the Holland and Reitman [1978] simulation forced the adaptive system to acquire knowledge exclusively from sensory cues; however, their environment did not contain any information about objects or space. The problem that must be faced by the designer of an artificial environment is how to make sensory information available in a way that suggests the desired object and spatial properties. Since the primary requirement for guiding a locomotive task is information about the direction and distance of objects, the localization issue will be dealt with first.

The first decision that must be made in constructing an appropriate environment concerns the number of spatial dimensions. It is hard to imagine any kind of spatial behavior that would be interesting in a one-dimensional environment. On the other hand, designing a three-dimensional environment requires detailed attention to the layout of surfaces, volumes, and so on. These factors are important spatially, but they present additional complications to the more basic notion of locus required here. A simple two-dimensional environment can have many spatially interesting properties, such as area and occlusivity<sup>1</sup>, that vary depending on the vantage point [Benedikt, 1978]; and, at the same time, in two dimensions it is easy to directly manipulate direction and distance parameters. Accordingly, the environments to be considered will be two-dimensional.

As a further simplification, the loci in an environment will be indicated by discrete points on a square 20 by 20 grid (see Figure 3.2). In order to approximate the continuity of a more realistic space, a grid is given hexagonal coordinates so that each point has six neighbors; and, the boundaries are "joined" to make a torus so that the environment is a surface without any edges. An object is specified spatially by giving its locus and all coordinates from which its stimulus signals are potentially available.

---

<sup>1</sup>Occlusivity is a term used by Gibson [1966] to refer to the extent to which surfaces visible from a given vantage point seem to cover each other.



**Figure 3.2.** A typical distribution of signal intensities in an environment. Each object is located at the center of its distribution, and intensity falls off 30% per unit distance. The dots indicate loci where no signals are available. At the top are the distributions for two objects. Below, two similar objects are placed closer together so that their distributions merge.

Since radiant energy - in particular the intensity of light rays - provides the primary source of orienting information for many biological organisms [Fraenkel and Gunn, 1961], an intensity parameter is associated with each stimulus signal. The intensity of a signal is a scalar quantity that is highest at the object locus and falls off

with distance. This parameter is meant to approximate the measurable intensity due to a source of light or heat in the physical world, or the concentration level generated by a chemical stimulus. Signal intensity decreases uniformly with distance in all directions. In this sense, the simplifying assumption is made that all objects are sources of stimulus energy in a uniform medium. This avoids complications such as intensities generated from reflective surfaces, etc. An intensity gradient is thus made available for an adaptive system to determine the direction of an object. If an adaptive system has some knowledge of what the full strength signal intensity is, it can infer the distance of the object as well.

Each kind of signal is available via a characteristic "channel". A channel can be thought of in terms of certain physical properties of the medium through which signals from objects are made available. Two signals are said to be on different channels if they can both be available at the same locus at the same time without interfering with each other. An organism with an appropriate interface can then have access to either signal independent of the other. If two or more signals are available on the same channel at the same locus, they are merged into one signal. This means that the signal and intensity detectable by an adaptive system at that locus is some function of the various alternatives. Alternative intensities are simply added together. The nature of stimulus signals and how they combine with each



other will be discussed later when object identity becomes an issue. As noted above, the most relevant information for localization is intensity. The state of any point in an environment is therefore fully specified by indicating what kind of object, if any, is located there; and, by giving a vector designating the signal and intensity available on each channel.

### Simulation of a Simple Model

Does this class of environments adequately specify the desired spatial information? Clearly, any adaptive system equipped with sensory apparatus that detects intensity gradients will have no trouble locating objects. The answer is not so obvious, however, if the adaptive system is not allowed to directly perceive or compute the gradients; that is, if the system must act solely on the basis of intensities currently available. Recall that one design criterion for an artificial environment is that it enable an adaptive system to take a relatively limited sample of the available sensory information and figure out what its implications are. In this context, that means it should be possible to use the given signal intensities as they are. An adaptive system should be able to follow the gradient without ever having to compute it or detect it. In particular, it should be possible to design simple motor routines that a system could use to produce this behavior.


To see if this kind of environment has the desired

properties, a simple organism model is proposed. The overall organizing construct for this and all subsequent models is summarized by Simon [1969] as follows:

The condition of any goal-seeking system is that it is connected to the outside environment through two kinds of channels: the afferent, or sensory, channels, through which it receives information about the environment; and the efferent, or motor, channels, through which it acts on the environment. The system must have some means of storing in its memory information about states of the world - afferent, or sensory, information - and information about actions - efferent, or motor, information. Ability to attain goals depends on building up associations, which may be simple or very complex, between changes in states of the world and particular actions that will (reliably or not) bring these changes about. (p. 66)

The signal intensities are the only part of the environmental state to be used for orienting by the proposed organism. It is reasonable to suppose, at the simplest level, that built-in or "hard-wired" information is available about which motor components are appropriate for each "channel" in the environment. A typical arrangement is shown in Figure 3.3. An organism has innate structures allowing information from the environment to directly initiate the appropriate motor behavior. Tinbergen [1951] called such structures *innate releasing mechanisms*. While a system using only such mechanisms is not at all an adaptive or goal-seeking system, it does provide an adequate framework for testing proposed motor routines.

What kinds of innate motor complexes should be considered? There are several hierarchical levels to choose from: activation of particular muscle fibers; overall



**Figure 3.3. An innate releasing mechanism, the simplest arrangement for generating appropriate behaviors, in which a stimulus signal directly elicits the correct response.**

activity in a given muscle; coordinated activity in several muscle complexes to move a limb or joint; coordinated movements of several limbs; stereotyped, fixed pattern responses; and goal-directed, flexible instinctive behavior [Tinbergen, 1951]. The lower levels involve sophisticated and flexible coordination of specific motor channels in response to specific input configurations. Exactly how this occurs is still a subject of considerable controversy [Bindra, 1978] and, for the purposes of this research, those issues are best avoided. Indeed, it can be argued that at the lower levels it is not even possible to make reliable predictions about the relationship between input and output [Brunswik, 1956]. At the more generic levels of overt response - either stereotyped or flexible instinctive behavior - studies of animal behavior [Tinbergen, 1951] indicate that the major issues involve how one response is selected over alternative responses. Handling motor organization from this perspective means the focus in designing a hypothetical organism can remain at the level of a general control architecture.

Given this generic approach to motor responses, Bindra [1978] identifies two kinds of overt actions characteristic of animal behavior: transactional actions and instrumental responses.

Transactional actions, such as eating, drinking, sniffing, grooming, copulating, biting, and struggling to get free of a predator, are performed when the animal is in contact with an environmental object, including parts of its own body. Instrumental responses, such as walking, climbing, lever-pressing, and head-turning, bring the animal closer to a particular stimulus (or otherwise make the stimulus available) or take it away from a particular stimulus (or otherwise make the stimulus disappear). (p. 46)

This basic classification - consummatory or defensive actions in response to contact with an object and instrumental responses affecting potential commerce with an object - will guide the choice of motor packages for the organism being designed here.

Accordingly, consider a primitive system that has a few very basic motor actions: TURN, which changes the direction of motion 60° to the right or left; MOVE, which transports the system from the current locus to the one directly ahead; ESCAPE, which removes the system from contact with an object; and, CONSUME, which enables the system to "nourish" itself when in contact with some resource. It is easy to see how these primitive actions might be combined to produce somewhat more sophisticated motor packages. An exploration routine EXPLORE, for example, might consist of a sequence of


---

'In a hexagonal grid, the angles specifying the direction of neighboring points are all multiples of 60°.


MOVE's interspersed with a few random TURN's. The organism has its detectors in front, and they can pick up only the intensities at the loci directly ahead, to the right, and to the left (see Figure 3.4a). The system cannot detect signals emitted by objects behind it. This restricted "retina" gives the system a well-defined orientation, reduces the potential for being overloaded with information at any locus, and is an approximation of the sensory interfaces common to many biological systems. Assume, in addition, that the organism can detect when it is in "contact" with - that is, at the same locus as - some object.

Now consider a particular version of this hypothetical organism that is to use the signal intensities to approach objects in a one-channel environment. All objects are considered resources that are to be consumed. The existence of many loci where no signals are available is a simplification based on the assumption that, in a natural environment, salient information is relatively scarce [Kaplan, 1976]. In the absence of salient information, exploratory or seeking behavior is the most appropriate strategy to follow until adequate stimulation is found [Tinbergen, 1951]. There is a further simplification in that the objects are more or less uniformly distributed, the environment is unchanging, and there are no competitors for the resource. All these factors make random exploration an adequate foraging strategy.


Figure 3.4




The adaptive system can only pick up signals from the indicated loci




TURN to center the strongest intensity



TURNING away from one side means TURNING toward the other



TURN away from the most stimulated side



**Figure 3.5.** The environment used for organism simulations 1, 2, and 3.

The system has two innate releasing mechanisms: one that is triggered by contact with an object and activates the CONSUME response; the other is triggered by the presence of the stimulus signal from an object and activates some as yet unspecified APPROACH motor package. When no signals are detected, the EXPLORE routine is the default response. The priority scheme for how these generic motor routines compete for control of system behavior is - from highest to lowest - CONSUME, APPROACH, and EXPLORE. As a simplifying assumption, resources are considered stable and renewable in

the extreme. The organism is satiated after each activation of CONSUME and the availability of the resource is effectively undiminished.

What remains to be specified is the organization of the APPROACH motor package. The most straightforward kind of mechanism orients with respect to the strongest intensity currently detected. Computer simulations of this simple strategy are summarized below:

#### ORGANISM SIMULATION 1

##### Environment

One channel with "resource" objects only. Six objects are distributed so as to present a representative set of intensity configurations (see Figure 3.5).

##### Required Behavior

When a signal is detected, the organism should orient toward the object and move in that direction until contact is made. On contact with an object, CONSUME is activated. When no signals are detected, EXPLORE is activated. EXPLORE causes the system to randomly TURN left or right, then MOVE  $n$  times, where  $n$  is a random number between 1 and 6. This sequence is interrupted, of course, when a signal is detected.

##### Mechanisms

- 1) The restricted retina described above prevents the system from seeing behind it. This, along with a satiable need for the "resource", enables the system to move away from an object after contact and experience other parts of the environment.
- 2) The APPROACH motor routine: If the strongest intensity is not centered on the retina, TURN until it is (see Figure 3.4b), then MOVE.

##### Results

The simulation was run until the organism had encountered all the given stimulus configurations from several angles. In every case, the system successfully approached the object.



The simulated behavior shows that this APPROACH motor routine, applied at each locus, enables the organism to follow the gradient and find the object. A set of representative approach trajectories is shown in Figure 3.6. Using the strongest intensity as a directional cue is, under these circumstances, an acceptable strategy.'

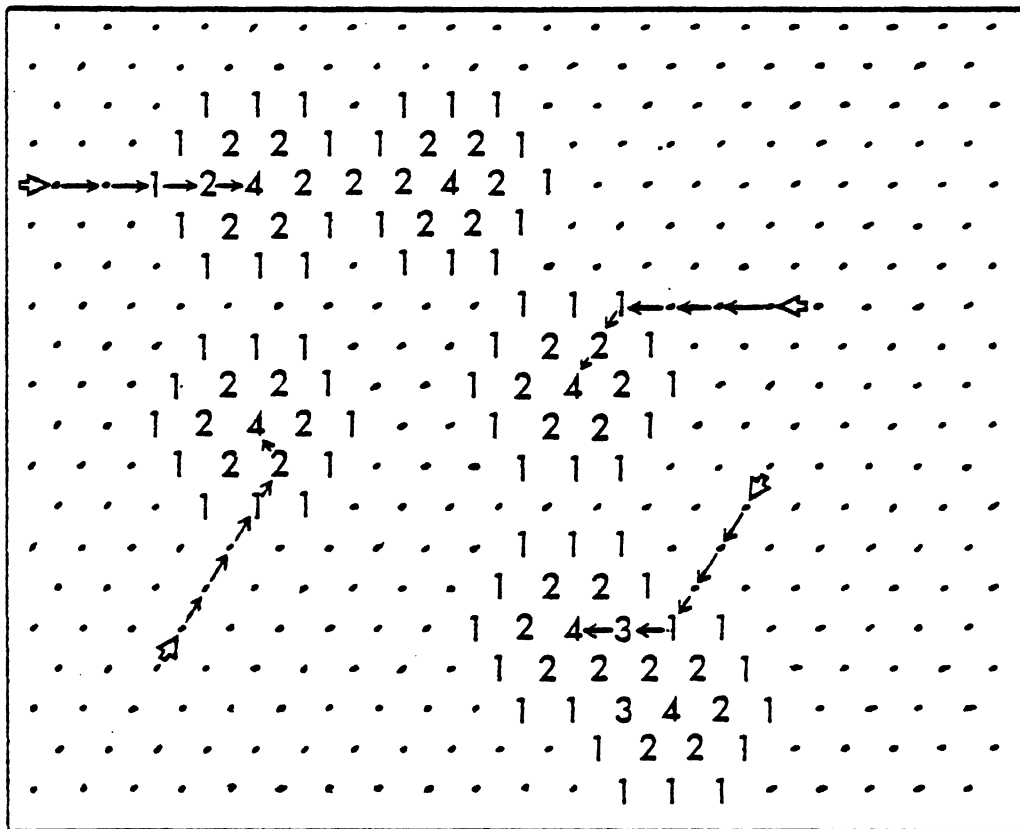


Figure 3.6. Some representative approach trajectories.

---

'Eventually there must be some limits on how strong an intensity can be and still remain useful for orienting. In a more realistic situation, intensities that are "too strong" might indicate danger or might be damaging or painful.

The behavior of this primitive system can be classified according to how each movement is related to the source of stimulation, in a manner analogous to the classifications used for biological organisms [Fraenkel and Gunn, 1961]. The system moves directly toward the source of stimulation, and therefore the behavior is a positive "taxis" or orientation. When confronted with two equally intense and equally attractive sources, each one stimulating a different half of the retina, the system orients toward one and disregards the other. The behavior is thus a positive "telo-taxis", where the prefix "telo" means goal or objective. It is hypothesized that biological organisms use some kind of central inhibitory mechanism to ignore one stimulus and orient to the other during telo-taxis. In place of such a mechanism, the organism described here merely chooses between the stimuli at random.

Now consider a one-channel environment containing objects that the organism is to avoid contact with. Again the system will use only two innate releasing mechanisms. This time, however, contact with an object triggers the ESCAPE response and detection of a stimulus signal triggers an AVOID motor package. How should this new motor package be organized? It would be convenient if the organism could manage to always move directly away from the source of stimulation. Such a negative telo-taxis, analogous to the strategy used above, would be a realistic mechanism,

provided the organism can "look" directly back.\* However, the organism as described cannot detect signals directly behind it. Since most biological organisms operate under the same limitation [Fraenkel and Gunn, 1961], it seems advisable to consider a different orienting strategy. An obvious alternative mechanism is simply to turn away from the strongest intensity currently detected. Simulations of this mechanism are summarized below:

## ORGANISM SIMULATION 2

### Environment

Same as in Figure 3.5, but the objects are now considered noxious.

### Required Behavior

When a signal is detected, the organism should orient away from the object.' When no signal is detected, the system explores. Any inadvertent contact with an object triggers the ESCAPE response. ESCAPE causes the organism to choose a new direction at random, then MOVE.

### Mechanism

An AVOID motor routine: TURN away from the strongest intensity, then MOVE.

### Results

The behavior was simulated until all the stimulus configurations had been encountered from several angles. The organism had difficulty only with the bilaterally symmetric stimulus configuration described previously (see Figure 3.4c). In that situation, the motor routine can get stuck TURNing back and forth - choosing one side to turn away from, then turning away from the

\*Another way to achieve negative telo-taxis would be to enable the system to move backwards. Many simple organisms do not have this capability, however, and neither does the hypothetical one being considered here.

'Since noxious objects are not mobile, a "freezing" reaction is inappropriate.

other side only to leave the organism back where it started.

This motor routine, like the approach routine, uses the direction of strongest intensity as the basis for an orienting reaction. The strategy is confounded, however, by the presence of equally intense stimulation from both sides.

A mechanism that leaves the organism in a behavioral deadlock is, to say the least, inconvenient. Some more functional solution must be found. The temptation is to remedy this problem by devising a strategy that uses the direction of weakest intensity directly, orienting toward weak intensities rather than away from strong ones.

Unfortunately, there are no studies of animal behavior that justify using such a mechanism [Fraenkel and Gunn, 1961].

Given equally intense stimulation from both sides, there is no way for an adaptive system to determine - using only the current retinal intensities<sup>10</sup> - when turning away from one source does not imply turning toward the other. The only option remaining is to go straight ahead, making sure the system never orients toward either object and, therefore, eventually avoids both of them.

The factor missing from the current strategy is an assessment of the equality, or lack of equality, of the

---

<sup>10</sup>Recall that the mechanism can only specify which simple motor actions are evoked by which intensity configurations. At this level, it cannot rely on anything suspiciously "cognitive" like remembering and reevaluating the configuration detected at some other location.

total stimulation across both halves of the retina.' This consideration suggests a somewhat different AVOID mechanism; namely, to turn away from the most stimulated half of the retina. The following is a summary of the results of simulations using this mechanism:

### ORGANISM SIMULATION 3

#### Environment

Same as in Figure 3.5, but the objects are now considered noxious.

#### Required Behavior

When a signal is detected, the organism should orient away from the object. When no signal is detected, the system explores. Any inadvertent contact with an object triggers the ESCAPE response.

#### Mechanism

The AVOID motor routine: If the stimulation on both halves of the retina is not equal, TURN away from the most stimulated half; then, MOVE.

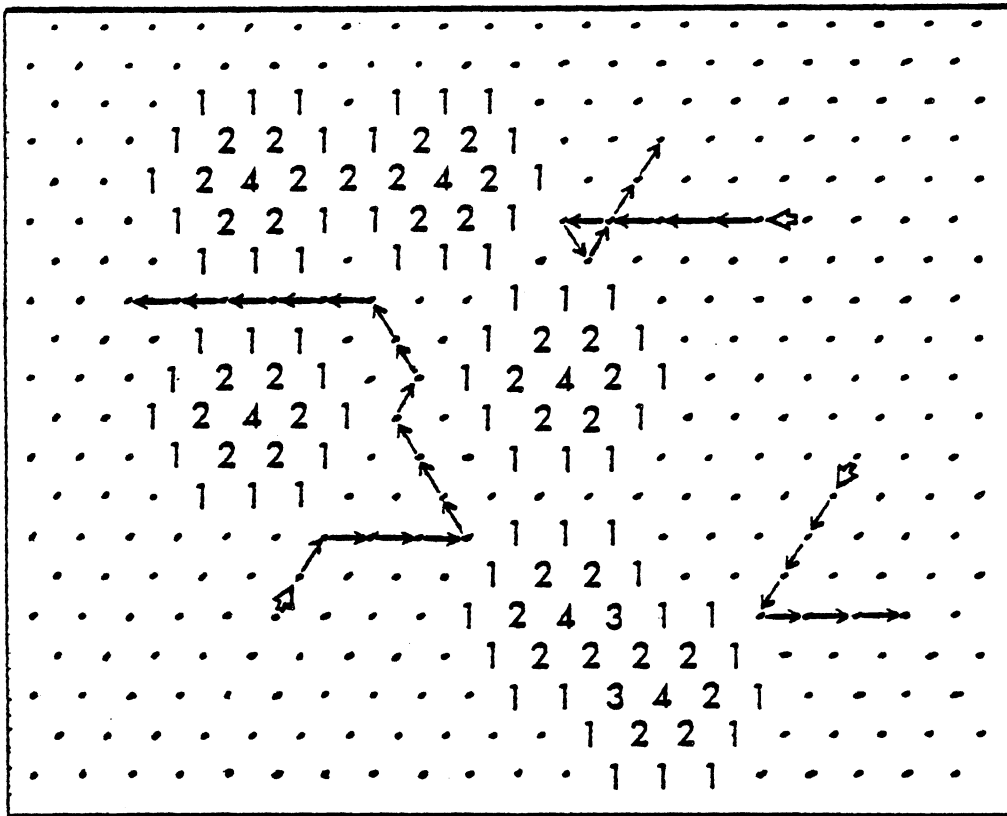
#### Results

The behavior was simulated until all the stimulus configurations had been encountered from several angles. The organism behaves appropriately even in the presence of the bilaterally symmetric stimulus configuration (see Figure 3.4d).

In the two-stimulus configuration this modified AVOID routine always guides the organism between and then away from both sources. The behavior is classified as a negative "tropo-taxis" since the orientation is based on a resultant of the stimulation levels on both sides. Figure 3.7 illustrates some typical trajectories.

---

'The "center" of the retina is considered part of both halves. A more biologically realistic system would be bilaterally symmetric and have two retinas.



**Figure 3.7. Some typical avoidance trajectories.**

At this point a brief discussion about the two-stimulus configuration is appropriate. Two equally intense stimuli, bilaterally and symmetrically located with respect to an organism, are not likely to occur in a natural environment. An organism is much more likely to be inundated at each moment with several stimuli of varying magnitudes from all directions. The fact that this particular situation arises so often in the simulated environment points out the extent to which stimulus configurations have been simplified. Moreover, even when the situation is artificially created

under laboratory conditions, it is not always the case that an organism behaving tropo-tactically will stay on the middle path between both sources [Fraenkel and Gunn, 1961].

Deviation from the middle path may be simply accidental, due to a somewhat erratic mode of locomotion. On the other hand, there may be asymmetries in the animal that give a bias to its locomotion. If one eye is in a different state of adaptation from the other, if the muscles of one side are fatigued, or if there are slight structural asymmetries, the maintenance of the middle path is not to be expected. (p. 152)

In other words, the bothersome "dilemma" faced by our hypothetical organism is an artifact of the "idealized" environment and the "perfect" set of detectors and effectors being used. Nevertheless, in the circumstances under which tropo-tactic behavior can be reliably elicited, animals do indeed move directly between both stimuli. The solution adopted here is therefore in line with the way simple animals behave under equally artificial conditions.

On the basis of the above simulations, it can now be concluded that the given class of environments makes spatial information available in a useful way. Simple mechanisms using that information have been shown to generate meaningful, directed orientation behaviors. Moreover, the use of balance of stimulation to control the action of symmetrical effectors is a very widespread orienting mechanism in animal behavior [Fraenkel and Gunn, 1961]. It is encouraging that these artificial environments afford the successful use of a mechanism that is consistent with what is known about biological systems.

### A Basic Goal-Directed System

The environments considered so far have been overly simplified in one other important respect. Each one has contained only one kind of object and, in each case, contact with an object directly releases the appropriate behavior. The fact that all objects considered so far have the ability to directly elicit responses is not an unreasonable simplification. Instinctive mechanisms rely heavily on the environment to signal the appropriate context for a potential response [Tinbergen, 1951]. Innate releasing mechanisms are in many ways cornerstones for instinctive behavior and for the acquisition of higher order units of behavior. Indeed, Tinbergen argues that "... it is not generally understood that learning and many other higher processes are secondary modifications of innate mechanisms, and that therefore a study of learning processes has to be preceded by a study of the innate foundations of behavior" (p. 6). It is appropriate, therefore, to begin our design of a hypothetical organism by first endowing it with simple inflexible behaviors. Of course, any goal-seeking organism will be capable of much more than this. By relying exclusively on innate releasing mechanisms, our hypothetical organism must wait for the environment to provide an action mandate. In a more complex environment, this organism would most likely have a long and dangerous wait. An organism capable of taking a more active role in the process of generating behavior would have a significant adaptive



advantage. What factors in the environment lead to such different demands on the organism? How can the organism be redesigned to meet those demands?

One way the simulated environment could be more demanding would be if both kinds of releasing stimuli, appetitive and aversive, were present at the same time. An organism with limited output capabilities would be forced to respond to one of the stimuli and ignore the other. It is hard to imagine circumstances in which it would be advantageous to always make the same choice regardless of the intensity of the stimuli or the physical requirements of the organism. Flexibility with regard to these factors is observed even in the behavior of creatures at low phylogenetic levels [Milner, 1970]. A further complication would be to introduce neutral stimuli into the environment and thereby make overall stimulation more complex and releasing stimuli perhaps more difficult to discern. An organism that can somehow make use of the periods of neutral and/or noisy stimulation to increase the likelihood that appetitive objects will be encountered and reduce the chances of encountering aversive objects will have a tremendous advantage over its competitors. Similarly, if the environment contains more than one type of appetitive object, it is to an organism's advantage to have some kind of "agenda" that enables it to gather everything it needs as efficiently as possible. This agenda must of course be flexible - capable of adjusting to the moment to moment

fluctuations in stimulus conditions. All these potential complications to the environment imply a basic principle about the internal design of a suitably adapted organism:

...the motor output comprising a response is separated from current sensory inflow by the intervention of some central system ..., and it is this system, not sensory-motor associations, that determines what the response will be. It is assumed that such a central system somehow flexibly adjusts the motor output to the ever-changing stimulus and organismic conditions. [Bindra, 1978, p. 44]

The required sophistication of such a central system will of course depend on the adaptive demands of the environment.

It should be pointed out that there are risks involved when an organism uses internal processes to select a response. Using an internal mechanism must not be such an involved process that the outcome - that is, the response - comes too late or not at all. On the other hand, the process must involve enough details of the current sensory array and organismic state to insure that the response will be an appropriate one. As the internal processes get more sophisticated, the necessary balance between central and external determinants of behavior becomes more difficult to achieve. This is one area in which natural systems have been more successful than their artificial counterparts. Thomas [1977], for example, notes that the pragmatic - as opposed to strictly logical - nature of human inference procedures is an advantage. "The fact that people are locally driven and subject to constant, parallel, and multidimensional stimulus input enables them to operate

under a much larger set of inference rules than would be optimal for a deterministic, sequential, and top-down goal-driven machine" (p. 12).

The most basic kind of internal factors involved in modulating response selection are called "motivational" factors [Tinbergen, 1951]. These factors depend on metabolic conditions such as hormone concentrations, food or water deprivation levels, etc., and their effect is to determine the *central motive state* of the organism. While there is some disagreement over exactly how this central motive state is generated and coordinated with action [Bindra, 1978], there is general agreement that the motivational state changes the relative effectiveness of stimuli to elicit behavior. In particular, the various metabolic conditions seem to serve as "... 'gates' or modulators in the paths between the receptors and the motor system. Each gate is selective for the particular impulses generated in the sensory systems by the relevant goal stimuli, so that the animal will approach food when it is hungry, a mate when it is sexually mature, and so on" [Milner, 1970, p. 399]. This rather straightforward control mechanism allows an organism to vary its behavior in appropriate ways.

It is important to emphasize the difference in function between the motivational factors and releasing stimuli. A motivational factor selects for some behavioral goal by priming or otherwise increasing the readiness of appropriate motor complexes [Tinbergen, 1951]. Conflicting action

mandates compete with each other for this control, probably through some kind of mutual inhibition. Once a set of alternative motor routines has been primed, the releasing stimulus activates - via an innate releasing mechanism - the one that best suits the current situation. The priming is necessary for the releasing action to be effective. In this way, motor control is not vested exclusively in sensory input. It is this kind of interaction between control information and the appropriate context that is the minimum processing requirement of any goal directed system [Scott, 1979]. It should be noted that motor complexes related to certain defensive or aversive reactions must be effectively "ungated" so that the releasing stimulus can reliably elicit behavior. It would be maladaptive, for example, if the detection of a predator did not elicit an innate inclination to flee. An organism with a central motive state is therefore best characterized as a system having both rigid, stimulus controlled behaviors and flexible, motivationally controlled behaviors.

These considerations suggest that motivational control can be imposed on the set of innate releasing mechanisms as shown in Figure 3.8. Pathways under control of the central motive state require facilitation in order for the releasing stimulus to have the capacity to elicit behavior. Facilitation designates one or more of the paths as alternative ways of achieving some behavioral goal. Environmental stimulation then activates the pathway most

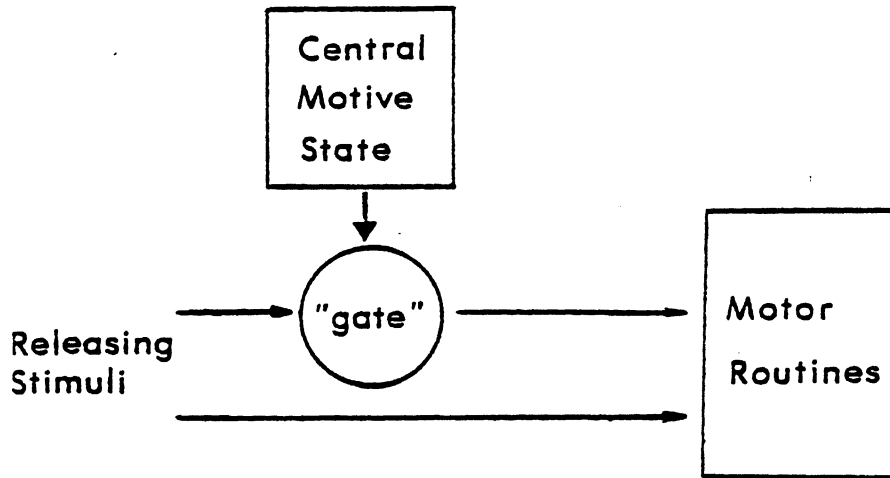


Figure 3.8. A simple implementation of motivational control. Some pathways require facilitation from the central motive state to be effective. Others are not subject to this control mechanism. All activated pathways must compete for control of the motor system.

appropriate to the current context. Pathways not subject to this control operate as before except that for conflicting action mandates there is a competition - based on strength - among all activated pathways for control of the motor system. It must be emphasized that, though the mechanisms here are still innate, they allow for flexible internal control. This is an important step in the elaboration of simple behaviors to generate more "intelligent" ones.

Lashley [1949] notes that "the mechanisms of instinctive and intelligent behavior ... seem fundamentally the same....

Higher levels of intelligence are based on a greater variety of types of organization, but this does not mean that they are any less dependent upon genetic factors" (p. 31).




Figure 3.9. The two-channel environment containing both resources [squares] and noxious objects [circles]. Each object has an intensity aura like those shown in Figure 3.5.

In order to demonstrate the effectiveness of this control architecture, a new environment is proposed. This environment has two channels, one for noxious stimuli and the other for appetitive or resource stimuli. Eight objects, four of each type, are arranged so that there is ample opportunity for making choices between the two stimuli (see Figure 3.9). Note that signals from resources and noxious objects don't interfere with each other because they are on separate channels. Our primitive organism is

modified to include the aforementioned motivational control structure. It is proper, at this point, to refer to the organism as a goal-seeking system. The system retains the innate releasing mechanism for ESCAPE. Now, however, the effectiveness of the CONSUME, APPROACH, and AVOID pathways are under "motivational" control. It is assumed that the organism must CONSUME enough of the resource to keep some internal deprivation level within tolerable limits. The deprivation level is simply a counter that is incremented every time step. When the counter is above a fixed threshold, the organism is "motivated" to seek resources. Everytime a needed resource is contacted, CONSUME resets the counter to zero. Contact with a noxious object is costly in that the counter is incremented by a substantial amount. If the deprivation level gets too high, the organism "dies". It is important to note that in this environment death is a very real possibility. With the maximum deprivation level set at 100, for example, a primitive organism using only EXPLORE to stumble across resource objects usually lives only 191 time steps.<sup>12</sup> The ability to make use of the aura of stimulus signals is therefore a crucial adaptive advantage.

When both kinds of signals are detected by the organism, the control of behavior is resolved as follows. If the deprivation counter is above the fixed threshold, the

---

<sup>12</sup>This was empirically determined by simulating 10 such organisms.

APPROACH pathway is facilitated in proportion to the difference between the counter value and the threshold; otherwise, the AVOID pathway is facilitated in proportion to the threshold value. In this way, the threshold corresponds to an innate, constant "motivation" associated with aversive stimuli. The pathway facilitated can be considered the winner of a competition between the appetitive and aversive influences. Final control is determined by computing a strength for each path, which is the product of the facilitation - if any - and the total intensity detected on that channel. The pathway with the highest strength controls which motor routine is activated. The strength associated with a pathway therefore represents the combined influence of internal and external factors. Note that CONSUME implicitly has higher strength than APPROACH and ESCAPE has higher strength than AVOID due to the special significance of contact with an object. It is not possible in this environment for the organism to be in contact with more than one object at a time. A more complete priority scheme for the output pathways is therefore not necessary.

The goal-seeking system just described was placed in the two channel environment with the implicit goal of keeping the deprivation level within tolerable limits and surviving. As a simplification, it is assumed that the effectiveness of the motor packages does not diminish with an increase in the deprivation level. The computer simulation is described below:



## ORGANISM SIMULATION 4

Environment

The two channel environment depicted in Figure 3.9.

Required Behavior

On contact with a needed resource, CONSUME it. This resets the deprivation level to zero. On contact with an aversive object, ESCAPE. This increments the deprivation level by 10. For other stimulus configurations, APPROACH or AVOID as indicated by the central motive state. When no signals are detected, EXPLORE.

Mechanism

The CONSUME, APPROACH and AVOID pathways are under motivational control. The threshold for motivating CONSUME and APPROACH is a deprivation level of 10. AVOID is always motivated at the constant level of 10. The organism dies if the deprivation level reaches 100.

Results

The organism demonstrated its ability to survive in this environment as the deprivation level was kept below 80 during a 433,428 time step interval of observation. The average deprivation level was 10.36.

The success of this organism in a non-trivial environment demonstrates that the motivational mechanism, as implemented, is an effective way to exercise flexible control over behavior. In fact, the control architecture of this goal-seeking system contains all the basic components found in Tinbergen's [1951] comprehensive model of the way instinctive behavior is coordinated.'<sup>3</sup>

The behavior of our organism can be described more systematically within the framework of Tinbergen's model.

---

<sup>3</sup>The principles of organization are considered here without subscribing to Tinbergen's assumptions about the accumulation and "draining away" of various impulses.

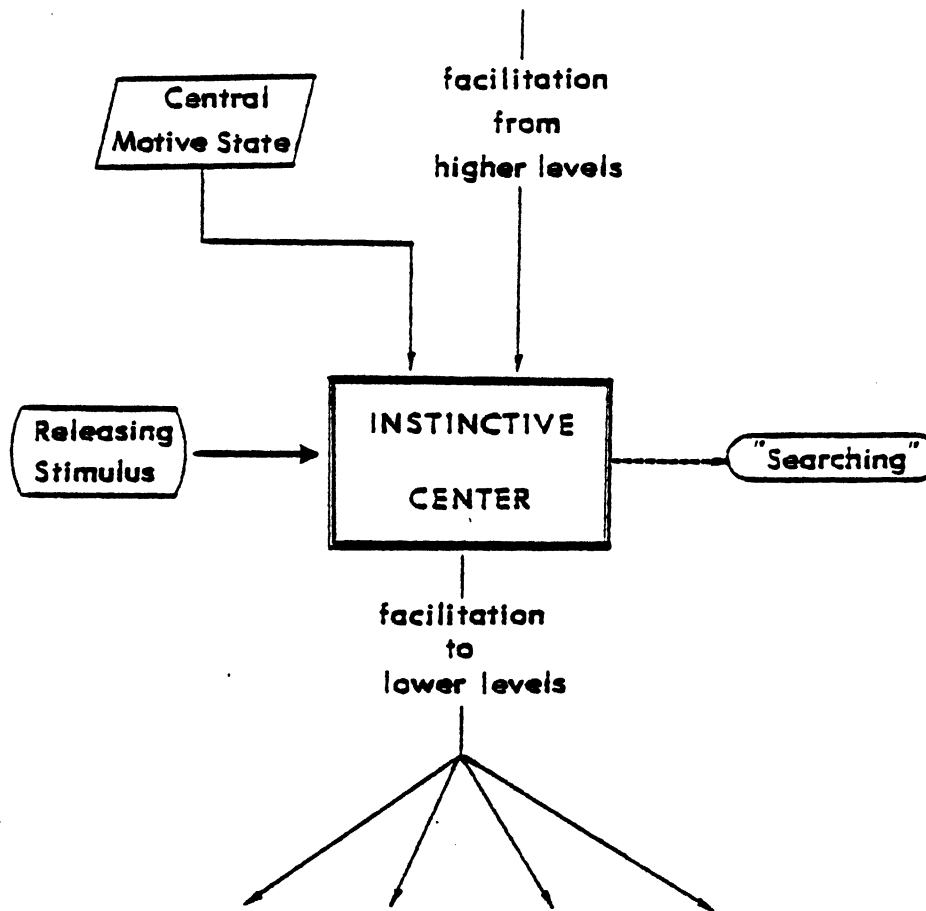


Figure 3.10. An instinctive center, the basic element in Tinbergen's [1951] model of the organization of instinctive behavior. Centers are linked together in a hierarchy. A typical center receives facilitation from the central motive state and from centers at higher levels. These factors combine to implement the "gate" effect on the releasing stimulus as shown in Figure 3.8.

According to this model, behavioral control is hierarchically organized. The lowest level components are the so-called "consummatory actions", simple responses like actual eating, escape, etc. that are characterized by stereotyped motor responses. These actions are either

released directly by external stimulation or they are under the control of higher "instinctive centers". Each such center itself receives control signals from the central motive state and, perhaps, from other superordinate centers as well (see Figure 3.10). Activation of an instinctive center - via the combined effect of a facilitating control signal and a releasing stimulus - leads to facilitation of all directly subordinate centers and consummatory acts. If none of these subordinate components becomes active, the center initiates a searching or exploratory behavior which strives to attain some behavioral goal. For example, the behavior might be directed toward finding a stimulus that will release one of the lower level centers.

Coordination in the system is achieved by the interaction of two selection criteria. Within a hierarchical level, active components are chosen on the basis of intensity. There are competitive - that is, inhibitory - interactions within any given level so that only the component with the strongest total mandate will become active and thereby initiate behavior. Between levels in the hierarchy, active components are selected on the basis of a simple priority scheme: output is controlled by the lowest active center in the hierarchy. In this way, the exploratory behavior associated with an active center is the default response when control is not assumed by one of the facilitated subordinate centers. Each hierarchy represents the organization of a major instinct. The highest center in

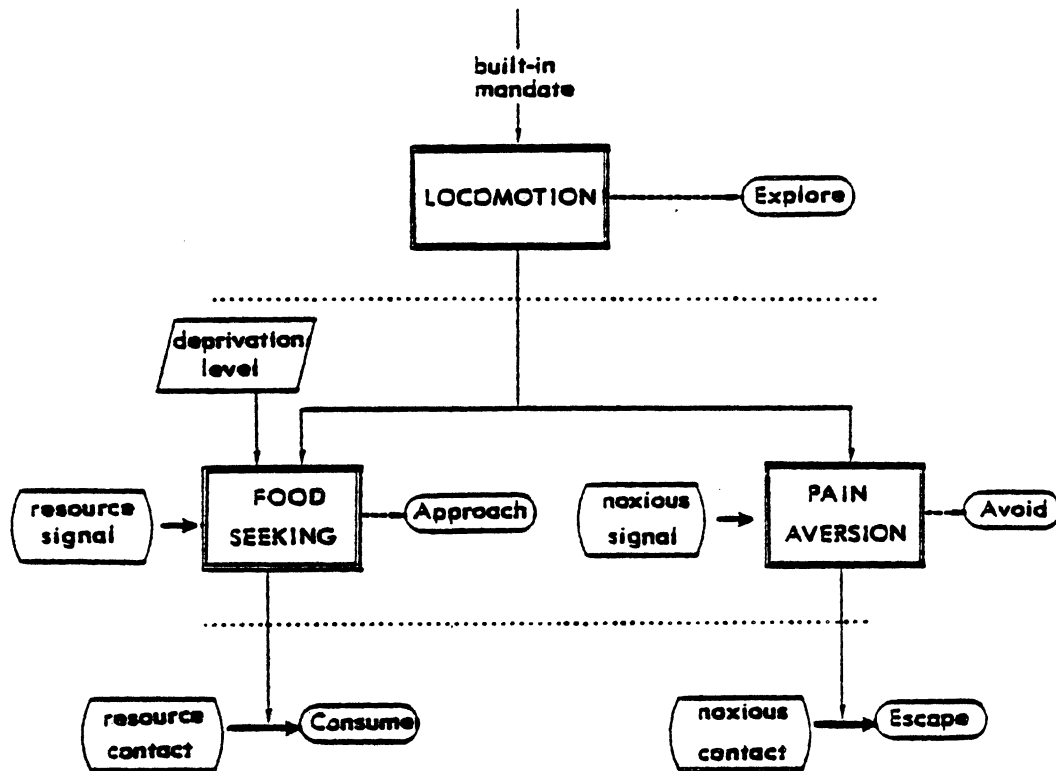


Figure 3.11. The control structure of the goal-seeking organism, shown as an "instinct" in the manner of Figure 3.10.

the hierarchy represents the overall goal. This center may or may not require a releasing stimulus to be activated.

Figure 3.11 indicates how the "locomotion instinct" of our goal-seeking organism fits into this framework. The system is organized into three levels. At the highest level is a center responsible for overall control of locomotion. In our organism this center is always active. Whenever none of the lower level components is active, the default motor response is EXPLORE. The middle level contains two centers subject to motivational control. Only one of these centers

can be active at any given time. The food-seeking center, when active, causes the organism to APPROACH a resource until contact is made and consumption is possible. The pain-aversion center, which always has enough facilitation to be released by a noxious signal, causes the organism to AVOID noxious objects. At the bottom level in the hierarchy are the two "consummatory actions" CONSUME and ESCAPE. CONSUME requires facilitation from the food-seeking center and contact with a resource to become active. ESCAPE is released automatically by contact with a noxious object. Given this explicit designation of the way various motor routines are coordinated, it is easier to see how the organism might be redesigned to perform more sophisticated tasks in more complicated environments. Moreover, we can ask well posed questions about how such a control structure might arise from experience with the environment (see, for example, Scott[1979]); and, hopefully, about how other forms of internal control of behavior might supplement and work with this structure.

## CHAPTER IV

### SOME CONSEQUENCES OF UNCERTAINTY

The goal-seeking organism model developed in the previous chapter has, from a functional point of view, two important information processing structures: the innate releasing mechanism and the motivational control system. The innate releasing mechanisms afford the generation of appropriate behaviors in response to functionally significant stimulation from the environment. When the organism/environment interface provides for direct and unambiguous selection of the relevant sensory/motor pathway, these mechanisms are - in and of themselves - a sufficient means of generating behavior. As the complexity of the interface increases, however, mechanisms only sensitive to external stimulation are no longer adequate for efficient functioning. One example of this increased complexity is the simultaneous presence of appetitive and aversive stimuli at the same locus in the environment. An organism in such an environment is faced with a series of choices, each of which impacts its prospects for survival. In order to meet this challenge, an organism must have criteria for deciding which stimulus is most important in a given situation; and,

it must have the structural apparatus to selectively process and respond to the chosen stimulus. Our goal-seeking model was modified to address these concerns by including an internal selection mechanism based on innate motivational factors. The motivational control system - using facilitation to nominate alternatives and competition to resolve conflicts - augments the innate releasing mechanisms to provide for flexible, adaptive responses; and, therefore, allows the organism to function effectively in more complex circumstances.

Complexity, however, is not the only issue that impacts the functional structure of organisms in a natural environment. As was argued in the last chapter, uncertainty is an important factor as well. The simulated environments and hypothetical organisms discussed so far have taken for granted the issue of object identity. Each sensory/motor pathway in the organism has been innately "tuned" to receive only those signals from the relevant channel in the environment, thus providing unambiguous information about the identity of the stimulus. Because the organism in this sense responds only to the presence or absence of stimulation on a channel, the structure of the stimulus signals has been left unspecified. Functionally, each signal has been an arbitrary, non-decomposable symbol. We have already argued that the stimulation available in natural environments has meaningful structure; and, that this structure is indicative of the distal stimulus source.

When the relationship between objects and their proximal stimuli is uncertain, information processing strategies relying exclusively on the moment to moment details of isolated stimuli are no longer adequate. An organism must discern and respond to the regularities or patterns underlying the proximal stimulation.

Accordingly, the simulated environment and goal-seeking system developed so far must be re-designed to account for uncertainty. For the environment, this means introducing variability and pattern structure into the object definitions; and, making stimulus signals available in a way that suggests the underlying regularities. For the goal-seeking system, this means that sensory/motor pathways must no longer all be hard-wired into their relevant environmental channel. Instead, the system must somehow be made sensitive to patterns. Information strategies and mechanisms must be found that discover which stimulus patterns are indicative of an object. The resulting equivalence classes must then eventually be associated with the appropriate motor pathways.

#### Specifying Stimulus Patterns

Given that the simulated environment is to be modified to include stimulus pattern information, how should those patterns be structured? To answer this question, it is helpful to examine the correlational structures in the real world. Gibson's [1966] analysis of natural stimuli,



referred to in Chapter 3, points out that there is considerable spatiotemporal structure in any natural stimulus. As for the pattern structure across various stimuli, the most fundamental concept is the notion of a *category* or equivalence class of objects. "A category exists whenever two or more distinguishable objects or events are treated equivalently" [Mervis and Rosch, 1981, p. 89]. Humans structure their object categories so that, in general, members of one equivalence class are more "similar" to one another than members of alternative equivalence classes [Medin and Schaffer, 1981]. There is considerable flexibility concerning the way the term "similar" can be used to define a category. At one extreme, membership can be rigidly determined by the presence or absence of certain critical features in the stimulus. The features in this case are said to be "criterial" for class membership. A more flexible criterion would check for several alternative, perhaps overlapping sets of features, any one of which would be sufficient for class membership. Such features are "characteristic" of the class as opposed to criterial. At the other extreme, one can imagine equivalence classes in which the members have no common features at all [Wickelgren, 1969]. Membership in this case is based on some higher level notion of similarity derived from the knowledge state of the observer making the classification.

There has recently been a considerable amount of research and theory examining the structural details of the

equivalence classes formed by humans and the corresponding structural basis for those classes in the environment. The equivalence classes are usually referred to as *natural categories* (see Mervis and Rosch [1981] for a review). Several useful characterizations of real-world correlational structures and the categories that delineate them have emerged. Real-world attributes or features do not occur in arbitrary combinations.

... The world is structured because real-world attributes do not occur independently of each other. Creatures with feathers are more likely also to have wings than creatures with fur, and objects with the visual appearance of chairs are more likely to have functional sit-on-ability than objects with the appearance of cats. That is, combinations of attributes of real objects do not occur uniformly. Some pairs, triples, or n-tuples are quite probable, appearing in combination sometimes with one, sometimes another attribute; others are rare; others logically cannot or empirically do not occur. [Rosch *et al.*, 1976, p. 383]

There is increasing evidence [Mervis and Rosch, 1981] that most natural categories do not have rigid necessary and sufficient membership criteria. The boundaries of natural categories tend to be fuzzy or ill-defined. Some members of a category may be more "typical" of the category than others because the members vary in the number of characteristic attributes they possess. The issue of category membership is therefore best described by the phrase "family resemblance", indicating criteria possibly more general than any feature-based similarity.

An organism that categorizes stimuli in the real world must find a balance between two conflicting constraints. On

the one hand, it is advantageous for the organism to delineate as many categories as possible to reduce the potential errors in any future discriminations between categories. On the other hand, as it was argued in Chapter 3, the primary rationale for using equivalence classes is to "... reduce the infinite differences among stimuli to behaviorally and cognitively usable proportions" [Rosch *et al.*, 1976, p. 384]. As Kaplan [1978] has pointed out, there is a tradeoff between accuracy on one hand and speed and economy on the other. For an organism in a natural environment, unerring accuracy must be sacrificed to achieve increased speed and economy.

In spite of the flexible structure of natural categories, organisms must somehow manage to make consistent and economical categorizations. How can the structural information about a category be usefully summarized? Rosch *et al.* [1976] argue that class inclusion is an important relation among categories with respect to their informational structure. The more inclusive a category is, the higher its relative level of abstraction. Within each such hierarchy of natural categories there seems to be a "basic level" of abstraction. This is the most inclusive level at which an organism can categorize and efficiently discriminate among real-world correlational structures. 'In general, the basic level of abstraction ... is the level

---

'Which level in a hierarchy is "basic" will vary depending on the organism in question.

at which categories carry the most information, possess the highest cue validity, and are, thus, the most differentiated from one another" [Rosch *et al.*, 1976, p. 383]. For example, the concept "chair" is a basic level category for humans because there are predictable clusters of attributes and functions common to all or most category members: chairs have legs, a seat, arms, a back, they can be sat upon, etc. Superordinate categories like "furniture" have fewer attributes shared by members of the category. Subordinate categories such as "kitchen chair" contain many attributes which overlap with other categories at the same level (e.g. "desk chair"); these categories are therefore less efficient for making discriminations. The basic level is in this sense the most inclusive level at which the informational content of attribute clusters is maximized. Rosch *et al.* have shown that humans learn basic level categories before categories at other levels; that at the basic level people tend to use similar motor actions to interact with category members; and, that the basic level categories reflect the fundamental classifications made during perception.

Given that the boundaries of a category are not clear cut, even at the basic level, how is it that categories come to be differentiated? One way

... to achieve separateness and clarity of actually continuous categories is by conceiving of each category in terms of its clear cases rather than its boundaries.... Categories can be viewed in terms of their clear cases if the perceiver places emphasis on the correlational structure of

perceived attributes such that the categories are represented by their most structured portions. [Rosch, 1978, p. 35-36]

At the basic level, members of a category are sufficiently similar that the entire category can be characterized by a *prototype* or ideal cluster of features. Because of the heightened similarity at this level, the probabilistic cue validity [Brunswik, 1956] of the category - a measure of the extent to which attributes are reliable indicators of category membership - is maximized. Prototypes are an exaggeration of category structure in the sense that they represent only the most characteristic attributes clusters of the category; and, just as basic level categories maximize the cue validity of attributes across hierarchical levels, prototypes maximize cue validities within a category level. All these considerations point to basic level categories and prototypes as useful and efficient summaries of the correlational structures in the real world.

It is not unrealistic, therefore, to impose a prototype-based structure on the signals available in the simulated environment. Each object category in the environment will be designated by some defining or characteristic set of feature values. The goal-seeking system must infer from experience with individual signals or *exemplars* which features are relevant to that category and which are irrelevant. The resultant categorical information

then provides a basis for classifying novel signals.' The next chapter will discuss how the goal-seeking system can manage to make the required inferences. What must obviously be considered first is how to generate the feature values for each exemplar.

There are several ways in which an object category can be generated on the basis of a prototype. The basic idea is to generate category exemplars by varying some attribute values of the prototype. In the simplest case, when all members of the category must share a set of common attributes, the variations among exemplars comes from substitutions within the irrelevant attribute dimensions of the prototype. A more flexible, family resemblance-type category can be generated by independently modifying the relevant attributes of the prototype. These modifications or distortions are done in a way that makes the prototype the statistical "center" of the category; that is, the prototype attribute values can be the average of the values across all exemplars, the most frequently occurring values, or some other perceptual designation of "ideal" values. At

---

<sup>2</sup>Medin and Schaffer [1978] point out that many prototype-like effects can be accounted for without resorting to prototypes - using, instead, only information about specific exemplars. This point of view does not account for experiments showing that, days after a pattern recognition task, subjects' memory for prototypes is stronger than for any exemplar [Posner and Keele, 1970].

yet a more general level<sup>3</sup>, the prototype might be "structural" in the sense that it is both the relevant attribute values and their correlations that determine category membership; that is, certain values must occur simultaneously in an exemplar. Modifications to relevant attributes in this case cannot be made independently. They must preserve the correlational structures that are characteristic of the category.<sup>4</sup>

Hayes-Roth [1973,1974] provides a theoretical analysis of these various alternatives in terms of the difficulty of the abstract classification problems they present. One approach to generating pattern classes is called the "volume" approach [Hayes-Roth,1974]. The basic intuition behind this approach is that exemplars for a category are contained within some well-defined "volume" or region in an M-dimensional feature space. Pattern generators based on this approach typically produce exemplar features that vary stochastically, within limits, from prototypical values. Values along each dimension are usually determined independently, although sometimes they are chosen to preserve pairwise linear correlations between dimensions.

---

<sup>3</sup>Although prototypes are often thought of exclusively in terms of some highly specific exemplar (e.g. Reed [1972]), Palmer [1978] points out that there is a continuum of prototype-like approaches based on the degree to which within-category variation is represented by the prototype as opposed to within-category invariants.

<sup>4</sup>There is evidence [Whitehead,1977] that this feature-correlation information is in fact used by human subjects in pattern recognition tasks.

The classification problem associated with this approach involves locating a test item at a point in the M-dimensional space based on its values along each feature dimension. If the point lies within a region centered around some prototype, the test item is classified as a member of that class. Given the assumption that all categories are at the basic level so that cue validities are maximized, each point lies in only one volume in the space. The classification problem is therefore trivial - merely determine which volume, if any, a point belongs to. The simplifying assumptions of this approach, however, make it useful only for a very restricted set of problems.

An alternative way of generating patterns is called the "schematic" approach. This approach acknowledges the fuzzy boundaries of natural categories by using combinations of feature values - or *schema* - to define membership in a category. More specifically, a schema is simply a set of feature values. In an M-dimensional feature space, a schema of k features corresponds to a k-dimensional hyperplane. The schematic approach assumes, in the simplest case, that each pattern class can be defined by a single structural prototype or *characteristic*. Each such characteristic is a schema designating a set of feature values required for class membership. Unspecified values are assumed to be irrelevant to the given pattern.

It follows from such a view that many of the ... attribute values may be irrelevant for some patterns and that, in general, different patterns may be best predicted by different subsets of



attributes. The ... pattern recognizer must determine which attribute values are coadapted, that is, which combinations are necessary and sufficient for determining membership in each pattern class. [Hayes-Roth, 1973, p. 343]

The obvious generalization of using just one characteristic to define a pattern class is to permit several characteristics to define a class disjunctively. Even more flexible membership criteria can be attained by allowing a characteristic to have less than perfect correlation with a pattern class. Pattern generators based on the schematic approach generate exemplars by assigning the mandatory combinations given by one or more of the pattern characteristics and producing irrelevant feature values probabilistically. In this way, each exemplar of a class manifests at least one of the pattern class characteristics. The classification problem can be very difficult under the schematic approach, since any given item can *match* or instantiate the characteristics of several alternative pattern classes.

On the basis of the previous discussion of natural categories, the schematic approach to generating patterns is the most appropriate one for the simulated environment. It allows classification problems having varying degrees of difficulty to be posed that challenge the organism at any one of several levels of environmental complexity and uncertainty. Hayes-Roth [1976] suggests a simple, yet powerful representation framework for a schematic pattern generator. "Each feature can be associated with a single

value in a bit vector - the value being one if the feature is present and zero if not ..." (p. 319). This is basically just a standardized property list description of an object. However, the bit string encoding means that the category level pattern information can be recovered by simple bit operations: the bitwise logical product of two bit strings yields the set of features common to both strings; and, if pattern characteristics are represented by bit string templates of relevant features and essential feature values<sup>1</sup>, the match between an exemplar string and a pattern characteristic can be determined by simple masking operations. The use of bit string encodings for information in the simulated environment therefore affords simple and efficient abstraction and classification processes for any proposed organism.

At first sight it might appear that the bit string representation is too weak. For example, property lists have a pre-determined fixed length which, in many cases, can be a disadvantage [Minsky, 1963]. It is not unreasonable, however, to assume that an organism has a limited set of sensory detectors and hence that the fixed length signal is appropriate. The bit string representation cannot be used to express arbitrary, general relationships among properties. However, this is a concern only from the point of view that all the information about an object must be

---

<sup>1</sup>Either the presence or absence of a feature might be designated as essential.

economically captured in a single bit string; for example, by trying to designate relationships among features as simply another kind of feature [Palmer, 1978]. Such an enterprise is not likely to be successful given the flexible structure of natural categories. Several characterizations of each type of information might be required and they could need different kinds of processing. For example, there is a growing body of evidence indicating that organisms use two largely independent mechanisms to process contour information and spatial information [Kaplan, 1970]. It is therefore reasonable to assume that the many facets of properties and relations in object categories can be adequately designated by sets of binary strings, each set specifying a particular aspect of the overall representation of an object.

Of equal concern is the facility with which graded comparisons of objects designated by bit strings can be made.

The attractiveness of bit operations to effect learning (abstraction) and pattern recognition in the framework of feature list descriptions is seriously diminished by the fact that such matching operations are fundamentally all-or-none. Each bit in a feature vector represents an attribute that is or is not present and provides no basis for fuzzy comparisons of two objects. In this context, the term fuzzy refers to a graded measure of the degree to which the values of the same attributes of two objects are similar. The capacity to retain the information that two objects are fuzzily equal is important in many learning problems involving continuous, ordinal, or noisy data. [Hayes-Roth, 1976, p. 319]

Fortunately, there are ways of organizing features so that

bit operations can be used to make graded comparisons. One such organization is based on the notion of "receptive field" in visual perception and is called a "radial generalization manifold" [Hayes-Roth, 1976]. The basic idea is that for a given feature dimension, some finite subset of values  $x_1, \dots, x_n$  is chosen so that any point in the maximum possible range of all values lies within a radius  $r$  of one of the  $x_i$ . The parameter  $r$  is called the "radius of generalization" and designates "... the maximum difference between the values of the same attribute which can be viewed as fuzzily equal" [Hayes-Roth, 1976, p. 321]. The "receptive field" of a value  $x_i$  is the set of values within the radius of generalization. It is assumed that adjacent fields overlap. The distance between adjacent values  $x_i$  and  $x_{i+1}$  determines the amount of discrimination possible between two patterns along this dimension. The bit string representation of this feature is simply a string of length  $n$  where the one or zero at position  $i$  indicates whether or not the feature value belongs to the  $i^{\text{th}}$  receptive field. Given this encoding, fuzzily determined matching and abstraction operations are possible with the simple bit operations described above. For example, suppose the range of values for a given feature is the continuous interval  $[0, 10]$ . If the four values 2, 4, 6, and 8 are

---

'Hayes-Roth [1976] also gives a method for representing disjunctive sets of features in a single bit string. For the purposes of this research, however, disjunction is adequately represented by having one bit string for each disjunctive component.

chosen as the centers for receptive fields of radius  $r = 2$ , then any value in the interval can be represented by a bit string of length four. The value 2.9, for instance, is represented by the string 1100 and the value 5.2 by the string 0110. A simple bitwise logical product reveals these two values to be fuzzily equal since they are both contained in the second receptive field.

The use of bit strings as stimulus signals therefore does not impose any unjustifiable limitations on the complexity of objects that can be defined in the simulated environment. Moreover, from the standpoint of the learning algorithm to be presented in the next chapter, the bit string representation is in some sense an optimal choice (see Holland [1975]). Patterns in the simulated environment will therefore be defined over bit strings. The most straightforward way to define a pattern class or object - following the schematic approach - is to specify a set of characteristics associated with the class. Each characteristic will be a string in the alphabet  $\{1,0,*\}$  where the \* is a place holder for an irrelevant feature and each 1 or 0 designates a relevant feature value.<sup>7</sup> A characteristic is a schematic template for generating binary strings in the sense that the 1 and 0 indicate mandatory values and the \* indicates values to be generated at random. Thus the characteristic  $1*0*$  generates the four strings

---

<sup>7</sup>The interpretation of a characteristic could also be in terms of receptive fields for various features dimensions as discussed previously.

1000, 1001, 1100, and 1101. When more than one characteristic is associated with a pattern class, each one is given some non-zero probability of being applied to generate a pattern exemplar. As a simplification, each pattern class will be assigned exclusively to one channel in the environment. This means that when two or more signals are available on the same channel at the same locus, they are in some sense redundant. One signal can therefore be chosen at random to be available to the goal-seeking system without any significant loss of information.

#### Identification of Distal Objects

The new class of simulated environments now makes a functional distinction between proximal cues and distal objects. Each stimulus signal, as an isolated entity, is an unreliable basis for generating behavior because of the variability in the way signals correspond to objects. How does this change impact the design of our hypothetical organism? The instinctive version of the organism had all of its motor pathways hard-wired into the relevant environmental channels. In order to make sure that the organism has to rely on the structure of the stimulus signals, these hard-wired pathways are eliminated. Moreover, suppose that information about which external channel a signal belongs to is lost when the signal is transmitted across the environment/organism interface; that is, the interface preserves - or reconstructs - the

distinctive nature of each signal but, internally, all signals belong to the same sensory modality. For example, the appetitive and aversive objects in the environment might be sources of chemical stimulation. The stimulus signals can be on different external channels in the sense that the concentration gradient generated by one type does not interfere with the gradient generated by the other; yet, internal to the organism, all the signals are processed within the same modality. This added restriction is to make sure that the only information available to the organism about object identity is the underlying structure of the proximal stimulation.

One rather straightforward, but perhaps subtle, implication of these conditions is that the mere detection of a signal - or sensation - is no longer a sufficient information processing mechanism. Somewhere between the sensation and the overt response there must be a process that ignores the variants in proximal stimulation and discerns the distal source. An organism must have, in other words, some kind of perceptual machinery. Psychologists have established several general characterizations of the perceptual process and its functional significance to an organism.

Perhaps one of the most useful and durable descriptions of the perceptual process is William James's: "Perception is of probable and definite things". By "probable" he meant that we tend to perceive what is likely, what is familiar, even when the stimulus is in fact not familiar. By "definite" he meant that we tend to perceive clearly, even when the stimulus is vague, blurred,

or otherwise ambiguous....

A percept that is definite is a necessary condition for speedy response. And being probable makes the percept reasonable; it reduces the need for accuracy, even if it is not a direct substitute for it. Essentially what the organism is doing, in classical functionalist terms, is making a "best guess". [Kaplan, 1978, p. 31-32]

Implicit in this characterization is the importance of prior experience and learning. The "guess" the organism makes about a stimulus configuration is derived from many concrete experiences with detailed stimuli. Learning processes distill these experiences into an internal residue or abstraction that summarizes which stimulus properties usually occur together. This abstraction has a structural basis [Hebb, 1949] and, because it more or less corresponds to the experiences it is derived from, it is an *internal representation* of those particular stimuli [Kaplan, 1973]. Since the residues of varied experiences with an object are the structural regularities underlying the proximal stimulation, an internal representation is a way of achieving a perceptually stable environment [Hilgard, 1978] - one that is in correspondence with the distal stimulus source. The fact that the perceptual process leads to "definite" outcomes points to a fundamental bias to perceive clearly [Woodworth, 1947]. A clear percept has two properties that Hebb [1949] calls "unity" and "identity". Unity refers to the extent to which the perceived object is distinct or segregated from the background or concomitant stimulation. A perceived object has identity "... when it falls at once into certain categories and not into others



... [and] is capable of being readily associated with other objects or with some action" (p. 26). The orientation of perception toward potential action is particularly important for an organism as it interacts with its environment [Arbib, 1972].

These rather general characterizations of the perceptual process are useful constraints to keep in mind as we formulate a perceptual component for our hypothetical organism. The other source of constraint is, of course, the prototype-based, schematic structure of object categories in the simulated environment. It is important to emphasize that the specification of object categories in terms of prototypes constrains but does not determine the mode of internal representation, processing, and learning to be chosen [Rosch, 1978]. Consider, for example, how prototypes are assumed to account for perceptual classification.

There are categorical representations stored in memory.... However these representations are constructed, the input pattern is represented along the same dimensions as the prototypes. A measure of similarity is computed between the input pattern and each categorical prototype. The similarity is assumed to be highly resolved or even continuous. A decision strategy is used to assign the pattern to a category on the basis of degree of similarity. The most common classification rule is of the best-fit variety such that each pattern is classified into one-and-only-one category. [Palmer, 1978, p. 289]

It is easy to see how several representation schemes - generalized templates, binary features, multidimensional features, or complex structural descriptions - might implement this rather general process, each using its own

particular similarity metric and decision criteria. Given a suitable sensory interface, our hypothetical organism might use any of these alternatives as the structural basis for its perceptual machinery.

A representation scheme is available, however, that is particularly appropriate for the binary strings in the simulated environment; and, for the given locomotive tasks, does not require the organism to do any additional preprocessing of the signals.<sup>4</sup> The representations are called *classifiers* and were originally proposed by Holland [1976]. In its simplest form, a classifier is a straightforward kind of "production rule" [Davis and King, 1976]; that is, it is an ordered pair of symbol strings, the left string denoting the conditions under which the rule is applicable and the right string denoting what actions or consequences the rule engenders when it is evoked. For classifiers, each of these strings is called a *taxon*. The left string is the *input taxon*, a fixed length string designating the set of all binary strings the classifier is sensitive to. The right string is the *message taxon*, a string of the same length that determines what kind of *message* or binary signal the classifier is capable of generating. The only consequence of evoking a classifier is that a message is produced.

---

<sup>4</sup>The signals available in the simulated environment have already been preprocessed in the sense that each signal is a standardized feature description of the object as detectable by the organism.

More specifically, the input taxon of a classifier specifies some subset of all binary strings of a given length. A classifier is "sensitive to" signals belonging to this subset in the sense that it can be evoked whenever one of these signals is present. Each taxon of length  $k$  is a string  $\langle s_1, \dots, s_k \rangle$  in the alphabet  $\{0, 1, \#\}$ . This string denotes a subset of binary strings in the following sense. A binary string  $\langle b_1, \dots, b_k \rangle$  belongs to the subset just in case  $b_i = s_i$  whenever  $s_i = 0$  or  $1$ . In other words, each symbol  $s_i$  designates a value that must be present at position  $i$  for all binary strings belonging to the subset. When  $s_i = \#$ , position  $i$  is irrelevant for determining membership in the subset and the binary string can have either a  $0$  or  $1$  at that position. Strings belonging to the subset are said to *match* the given taxon. Thus, for example, the taxon  $\#11\dots 11$  designates the subset  $\{011\dots 11, 111\dots 111\}$  and the taxon  $\#0\#\dots\#\#$  is matched by any binary string with a  $0$  in the second position. This notation for designating subsets does not allow arbitrary subsets to be represented by a taxon.' For instance, there is no single taxon that designates exactly the subset  $\{000\dots 00, 111\dots 11\}$ . It is the case, however, that any arbitrary subset can be specified as the union of subsets denotable by taxa. In this sense, the taxa constitute a "basis" for representing the set of all possible subsets.

---

'Indeed, given any large set there is no compact and useful way of specifying arbitrary subsets without simply listing the members.

To date, the ability to produce messages has not been implemented in any of the systems designed to use classifiers [Holland and Reitman, 1978]. Nevertheless, this capability has a tremendous - albeit untapped - potential as a way of realizing some sophisticated information processing mechanisms. Message taxa are strings defined over the same alphabet as the input taxa and have the same length. A message taxon designates the binary string to be generated by a classifier as follows. Given the taxon  $\langle s_1, \dots, s_k \rangle$ , a binary string  $\langle b_1, \dots, b_k \rangle$  is produced such that  $b_i = s_i$  whenever  $s_i = 0$  or  $1$ . When  $s_i = \#$ , the value of  $b_i$  is the same as the value in position  $i$  of the binary string that evoked the classifier. In other words, the meaning of the symbol  $\#$  in a message taxon is that a particular value is "passed along" from the incoming signal to the outgoing message. For example, consider a classifier with input taxon  $1\#\#\dots\#\#$  and message taxon  $0\#\#0\dots00$ . The signal  $111\dots11$  matches the input taxon and, therefore, when it is present the classifier can be evoked. The message generated will be  $0110\dots00$ , passing along the information contained in the second and third positions of the signal. In this sense, a classifier can be said to "process" the signal that evokes it. It defines a transformation or mapping over the set of binary strings that match the input taxon. This transformation might have any of several interpretations: certain values in the string are replaced, indicating that some additional information is available or that a step has

been completed in some complex string computation; some subset of values is transmitted to be processed elsewhere, the modified values indicating where or how; and so on. It does not take much imagination to realize that the syntactic simplicity of a classifier is deceptive with respect to its potential as a component in an information processing system.

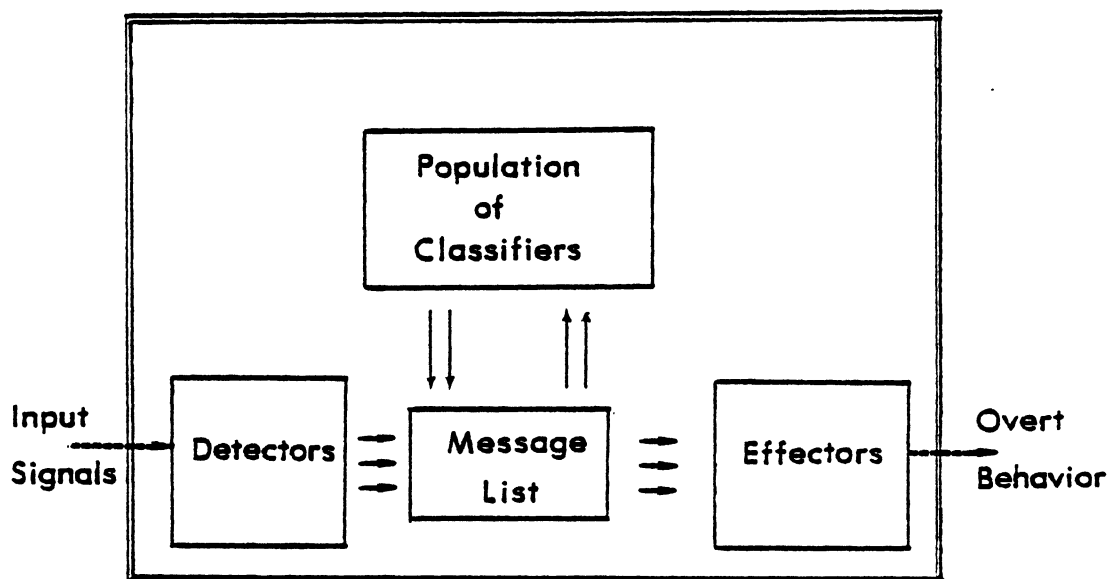


Figure 4.1. The structure of a simple classifier system.

Classifiers are implemented within the framework of a *classifier system*. In its most basic form, a classifier system has four components (see Figure 4.1): a set of input detectors; a *message list* or working memory; a set or *population* of classifiers; and, a set of effectors. In more detail,

- 1) The input detectors translate the information available from the system's task environment into fixed length binary strings. For example, each detector can be thought of as a device that analyses the current environmental configuration for the presence of some feature or property. The output of a detector is then a binary string indicating the presence or absence of the corresponding feature, a value computed along some feature dimension, a binary encoding of some structural information, etc. The signals generated for subsequent processing by a population of classifiers are formed by concatenating the output from some subset of detectors into a single binary string.
- 2) Signals derived from the detector outputs are placed on the message list. The contents of this list at any given time determine which classifiers and/or effectors will be evoked. In this way, all operations performed by a classifier system are "data-driven" and therefore potentially "... sensitive to any change in the entire environment and potentially reactive to such changes with the scope of a single execution cycle" [Davis and King, 1976, p. 304].
- 3) The population of classifiers is a fixed number of ordered pairs of taxa as described above. When a signal appears on the message list that matches a classifier's input taxon, that classifier has the

potential to become *active*. Active classifiers influence the overall operation of the system by placing their message on the message list.

- 4) The effectors can be thought of as classifiers whose message taxa specify actions the system can perform with respect to its task environment. When an appropriate message appears on the message list, an effector can be activated and the system can therefore make an overt response.

The basic execution cycle for these systems is straightforward. At the beginning of each cycle, the message list is emptied and a new list is constructed by including all signals derived from the input detectors. All active classifiers generate a message that is appended to the list, after which these classifiers cease to be active. Now each classifier in the population scans the message list and determines how much of a mandate it has to become active. This determination involves computing a *match score* for each message which indicates the degree to which each string matches the specifications of the input taxon.<sup>10</sup> On the basis of these computations, one or more classifiers are probabilistically selected to become active. At the end of a cycle, the system generates any overt behavior specified by active effectors.

This general characterization of classifier systems

---

<sup>10</sup>Strings belonging to the subset designated by an input taxon match the taxon perfectly.

clearly encompasses a large class of particular implementations. For example, we shall see in the next chapter that there are several ways to compute a match score. It is also possible to imagine situations in which variations on these basic themes might be appropriate: classifiers could be defined to have more than one input taxon to integrate information from several sources; classifiers might generate more than one message when activated; classifiers might be activated by the absence rather than the presence of a message on the message list; classifiers might remain active across several execution cycles; messages could have variable durations as well; and so on. The point to be made is that classifier systems are a powerful computational tool and representation scheme. It is a simple exercise to construct a classifier system that is computationally complete. Indeed, "classifiers, in combination, readily implement arbitrary production systems, providing the designer with a natural way of organizing the system initially" [Holland, 1980, p. 252].

We are now in a position to see why classifier systems are such an appropriate choice as the structural framework for our hypothetical organism. The correspondence between the syntax of an input taxon and the designation of schematic templates for signals in the simulated environment is obvious. One is a string in the alphabet  $\{0,1,\#\}$  and the other is a string in the alphabet  $\{0,1,*\}$ . Classifiers can therefore serve directly as internal representations of



environmental patterns and the objects they indicate. Because classifiers are capable of sending and receiving messages, they are capable of being "associated" with each other and/or with system effectors. An "association" means that the message generated by one classifier matches the input taxon of some other classifier. In this sense, an active classifier can be a percept that has "identity" insofar as it pertains to associations with other structures and with potential action. Furthermore, the generation, transmission and processing of messages has been cited [Hebb, 1972] as one of the key issues related to behavior in higher organisms.

... The complex communications network of the higher animal has developed so that messages run to and fro within it as well as into and out of it. Such internal activity, infinitely more complex than these words can suggest, is mind; and possession of this internal complexity is what distinguishes higher from lower animal, making the behavior of the higher animal less directly under the control of sensory input. (p. 79)

Classifiers system provide a simple but powerful framework in which complex message passing activities can be systematically analysed.

A classifier system can be "pre-programmed" to recognize certain patterns, form specific associations, etc. Moreover, classifier systems were specifically designed to be compatible with certain general learning heuristics [Holland, 1976]. This means classifiers can be modified, based on experience with the environment, to be sensitive to relevant sensory patterns. Since more than one classifier

at a time can be active, classifier systems can easily process object information designated by sets of incoming signals. Finally, note the similarity between the execution cycle of a classifier system and Palmer's [1978] description, cited earlier, of the perceptual classification of input patterns based on prototypes. If we assume that the message generated by an active classifier indicates a categorization decision, then the execution cycle is a way of implementing the desired perceptual process. All these factors, taken together, imply that the classifier system framework will provide our hypothetical organism with a more than adequate set of information processing tools.

#### Modifying Stimulus-Response Probabilities

Identifying the entities in the simulated environment is not the only issue raised by the introduction of variability and pattern structure into the stimulus signals. Of equal importance is determining which behaviors are appropriate in response to the presence of those entities. The hypothetical organism has to this point been equipped with hard-wired pathways leading from stimulus to response. This innate organization was successful because the design anticipated all possible behavioral contingencies. In effect, the organism was provided with a set of specific "instructions" that was sufficient for adaptive functioning. As Pulliam and Dunford [1980] point out, however, "in a complex, unpredictable environment ... there can be no set

of specific instructions that successfully anticipates all eventualities" (p. 4). An organism must have some discretion, within limits, to generate "instructions" appropriate for the actual circumstances it is confronted with. In other words, an organism must have the "... capacity to *learn*, where learning is understood to be nothing more than a change (in the environmentally appropriate direction) in stimulus-response probability relations" [Dennett, 1978, p. 76].

The increased flexibility engendered by this learning capability has two fundamental implications for the design of our hypothetical organism. First, the organism must start with some reservoir of modifiable, potentially useful pathways or "soft-wiring". Given the assumption that it is impossible to anticipate or explicitly provide for all eventualities, this set of soft-wired connections obviously cannot be exhaustive. Indeed, until the perceptual machinery develops some internal representations, it is not clear where specifically the soft-wiring must be available. From a design standpoint, therefore, it is important that the initial soft-wiring be chosen in a way that "covers" the space of possibilities.

Since it is impossible to determine, at the outset, which changes experience will require, the designer should consider the full range ... of systems that could result from various combinations of changes.... It is vital that the changes allowed for ... be rich enough to give a reasonable chance of correcting faults in the initial design. In other words, the system should not only learn, but some guarantee should be given that it can adapt to a wide range of

situations. [Holland and Reitman, 1978, p. 314-315]

This means that the initial soft-wiring must be robust in the sense that it can be selectively modified to instantiate any of a large number of stimulus-response probability relations.

The plasticity of soft-wiring is an advantage only if the organism is "... able to distinguish good results of plasticity from bad, and preserve the good" [Dennett, 1978, p. 75]. Consequently, a second implication of plasticity for the design of our organism is that criteria must be available that indicate whether or not a given pathway is adaptively salient; and, a mechanism must be provided for selecting and preserving the soft-wiring deemed adaptive. The organism needs built-in definitions of "good" and "bad" in order to select the pathways most consistent with its overall goals of functioning and surviving. Definitions tied explicitly to these long term goals would clearly not be very useful. An organism "... would have to try an option long enough to see if it survived and reproduced, by which time it would be too late to try another option" [Pulliam and Dunford, 1980, p. 6]. Instead, the definitions must allow for more short term, ongoing evaluations of experience. In the selection of an appropriate hard-wired response, we saw that the simplest mechanism relied on the environment to evoke the correct alternative. Similarly, definitions of good and bad are, in the simplest case, derived directly from the presence of certain stimuli in the environment whose

significance to the organism is innately determined. These stimuli are usually called *reinforcers*. Examples of such stimuli

... include sights, sounds, tastes, odors, temperatures, and cutaneous textures that are provided by such biologically important objects, events and situations as food, water, a sexual partner, a nest, the call of a distressed offspring, the shape of a predator, and injurious levels of heat or cold. [Bindra, 1978, p. 45]

By detecting these stimuli, an organism obtains a useful evaluation of the current situation.

An organism can make use of these evaluations by being innately disposed to seek out good experiences and avoid bad ones; that is, if the soft-wiring is modified so as to increase the probability that good experiences will recur and decrease the probability that bad experiences will recur. In this sense, good experiences are "rewarding" and bad experiences are "punishing". This mechanism for selecting adaptive soft-wiring illustrates the commonly observed principle of learning known as the Law of Effect: "rewards or successes further the learning of the rewarded behavior, whereas punishments or failures reduce the tendency to repeat the behavior leading to punishment, failure, or annoyance" [Hilgard and Bower, 1975, p. 34]. In a realistic environment, stimuli can be rewarding (or punishing) to varying degrees or with varying probabilities [Pulliam and Dunford, 1980]. Part of this is due to variability in the stimulation itself. The rest can be attributed to the presence or absence of the proper

motivating factors in the organism. As a simplification, we assume that, for a given level of internal motivation, the reinforcing effects across stimulus categories are all equally strong and equally reliable.' ' Moreover, we assume that the organism's evaluation of its experience is binary - either positive or negative. These assumptions reduce the complexity of the evaluative information and preserve the essential categorizations, thereby minimizing the processing burden on the organism.

The learning mechanism to this point includes a set of stimuli the organism is innately predisposed to interpret as good or bad; and, a simple learning rule that strengthens soft-wired behavior leading to encounters with good stimuli or the cessation of a bad stimulus. An organism with only these capabilities can generate a set of adaptively salient connections from stimulus representations to responses. It is often not enough, however, merely to have in place an appropriately strengthened set of pathways. Using these pathways, especially when a stimulus configuration nominates more than one alternative, might require some assessment of what the consequences will be. Pulliam and Dunford [1980] illustrate this with a simple example. Consider a hypothetical lizard with a motivationally controlled, innate disposition to eat ants. This lizard can function adaptively as long as the ants in its environment are all

---

' 'As a stimulus category becomes familiar and well learned, a more realistic organism might habituate to the stimulus and thereby neutralize the reinforcing effects.

edible and nourishing. Suppose, however, that black ants are edible but red ants are toxic. In order to avoid toxins the lizard must store information about the outcome of eating each kind of ant and refer to that information before following its impulse to eat the next ant. William James [1892] makes a similar point in his explanation of the way instinctive impulses are controlled.

Some expectation of consequences must in every case like this be aroused; and this expectation, according as it is that of something desired or something disliked, must necessarily either reinforce or inhibit the mere impulse. (p. 262)

These considerations suggest that it is advantageous for an organism to be able to store the evaluation - or affect - derived from a reinforcing stimulus; and, subsequently use the evaluation to influence which behaviors are selected by the motivational control mechanisms.

There is evidence that many organisms do indeed code and store these evaluations and that activation of these codes can serve a motivating function. Olds [1969] reviews several studies which demonstrate that mutually inhibitory reward and aversive centers are located in the hypothalamic- limbic regions of the central nervous system. This is part of the "old brain" that evolved before the neocortex and is a significantly large fraction of the total brain tissue in many lower animals [Zajonc, 1980]. Coupled with Olds [1969] observation that the anatomical distribution of reward effects is the same in man and at least nine other kinds of animals, this suggests that the reinforcement mechanisms

appeared very early in phylogeny. As for the coding function of these affect regions, Olds [1969] points to research suggesting that remembered behavior patterns previously associated with rewards can induce a temporary neural "memory" or activity in the hippocampus. This temporary activity serves a motivating function in the sense that it is correlated with the nature of the effective goal stimulus and resembles the activity induced by an actual reward. Furthermore, Olds [1969] speculates that "a reward stimulus might make a motivational inscription on both the afferent and the efferent sides of the memory elements" (p. 129). In other words, the available evidence indicates that affective codes are in fact stored and used by many organisms.

It is interesting to note that some of the affective regions are directly related to the control of consummatory actions and the central motive state. Electrical stimulation of various sites in these regions can elicit specific motor responses [Flynn, 1978] as well as complete, well coordinated generic behavior patterns in which the specific actions vary depending on the situation [Tinbergen, 1951; Bindra, 1978]. This suggests that the affective regions are also the locus for much of the "gating" functions of the central motive state discussed in the previous chapter. Having the evaluative code closely associated with behavioral control has obvious adaptive advantages. Activation of the code can lead directly to an



appropriate response, in spite of an inconclusive perceptual analysis. Zajonc [1980] offers a good example illustrating how this is adaptive.

A rabbit confronted by a snake has no time to consider all the perceivable attributes of the snake in the hope that he might be able to infer from them the likelihood of the snake's attack, the timing of the attack, or its direction. The rabbit cannot stop to contemplate the length of the snake's fangs or the geometry of its markings. If the rabbit is to escape, the action must be undertaken long before the completion of even a simple cognitive process - before, in fact, the rabbit has fully established and verified that a nearby movement might reveal a snake in all its coiled glory. (p.156)

The affective evaluations referred to here are defined very much like primitive sensory qualities; that is, they can be detected quickly and automatically. Zajonc argues that such evaluations are the dominant factor in the behavior of lower organisms and that the affective system is a basic and powerful determinant of behavior in higher organisms.

#### A Revised Design

To summarize, our hypothetical organism must be changed in several ways before it can cope with the uncertainty of the simulated environment. It has to have perceptual machinery in order to be sensitive to the patterns in the proximal stimulation. The organism also must have the means to modify its repertoire of stimulus-response pathways. This involves, among other things, the designation of certain stimuli as reinforcers having special significance in terms of evaluating which pathways need to be modified

and how they should be modified. Moreover, the organism should have the capability to store the evaluations or affective codes derived from reinforcing stimuli. These codes are then available to serve two roles: as a criterion for choosing among alternative behaviors that augments the function of the central motive state; and, as control signals that have a motivating function in and of themselves.

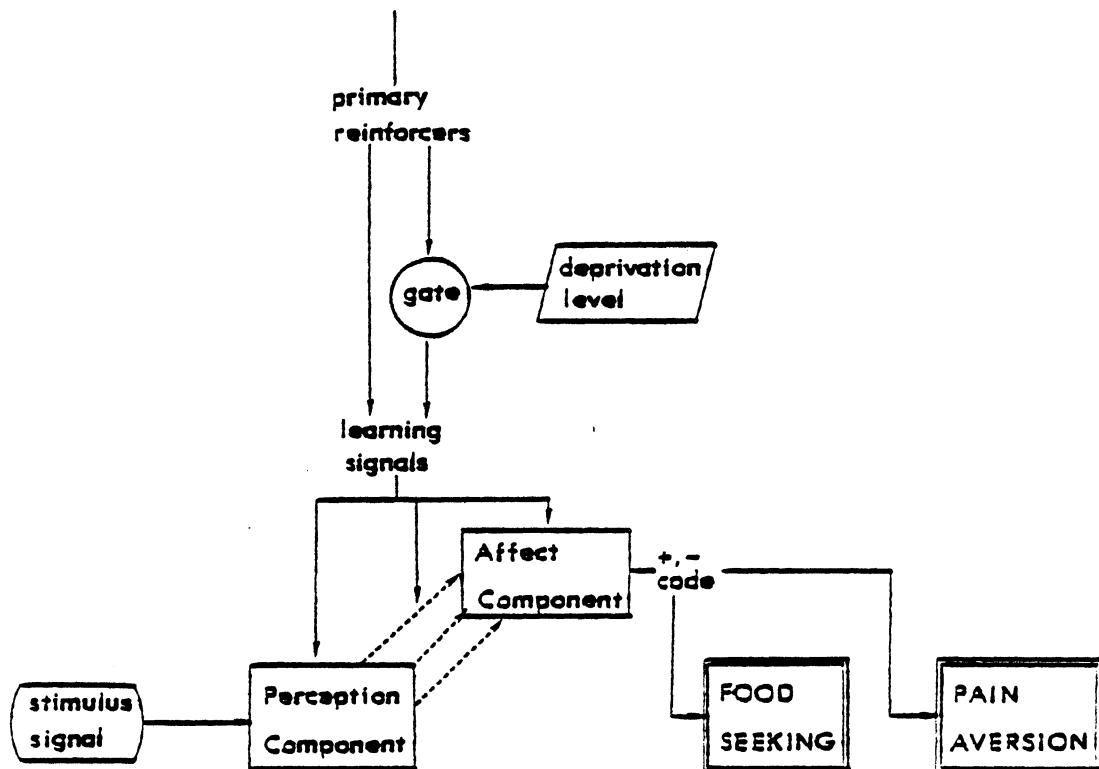


Figure 4.2. The integration of perception, affect, and learning into the goal-seeking system.

We can integrate all these concerns into our overall system as shown in Figure 4.2. There are two new components: a perceptual component that generates and stores representations of the stimulus patterns in the environment; and, an affect component that stores representations for the evaluations obtained from reinforcing stimuli. The soft-wiring between the two components allows the system to build associations between the representations of a stimulus pattern and the appropriate affective code. The perception/affect pathways are a replacement for the hard-wired inputs to the food seeking and pain aversion centers in the previous model. It is assumed that these two centers monitor the activity in the affect component. The overall code associated with that activity - either "good" or "bad" - determines which center is most likely to control behavior.

Under the assumption that in primitive organisms the affective assessment is the dominant behavioral factor, no direct connections between the perceptual representations and the motor routines are modeled here. Instead, it is assumed that the system's repertoire of innate releasing mechanisms contains all the necessary direct sensory/motor connections. Stimulus categories come to exercise selective control of behavior through association with an appropriate affective code. This preserves the basic role of releasing stimuli while at the same time permits the flexibility mandated by the uncertainty of the simulated environment.

It is worth emphasizing that the role of affect as a mediator between perception and action - though not absolutely necessary for our simple creature - is a very important elaboration of the system. It allows response pathways to be evaluated before the system is committed to an overt response [Kaplan, 1973]. In this sense, affect is a useful "pre-adaptation" for more sophisticated systems relying on internal processes to nominate potential courses of action.

Contact with either a resource object or a noxious object is the only reinforcing stimulus in the simulated environment. Obviously, contact with a resource is "good" and contact with a noxious object is "bad". If the organism's deprivation level is below threshold, contact with a resource object has no reinforcing effect. The result of reinforcement is learning in three areas:

- 1) The concomitant stimulus signals are assumed to be important. Therefore, the set of internal representations is modified to improve the system's ability to categorize the signals in the future.
- 2) The evaluation of good or bad is used to generate an affective code for the current situation. While there are only two generic kinds of codes, each situation elicits its own characteristic affective "pattern".
- 3) The soft-wiring connecting the current perceptual representation with the current affective representation is strengthened. This allows the

affective code to be retrieved upon subsequent occurrences of this kind of situation.

Note that, because modification of internal representations also goes on at a slow background rate without external reinforcement, the system has what amounts to an innate motivation to match all the signals it encounters. This assures that the most frequently encountered patterns will eventually get represented and that the soft-wiring is anchored to a stable and complete set of perceptual representations.

It is instructive to ask how sophisticated a learning system can be explained using only the basic principles and mechanisms discussed so far. Organisms relying solely on reinforcement from the environment are limited in terms of what they can learn. This is because "... they can learn only by actual behavioral trial and error in the environment. A useful bit of soft-wiring cannot get selected until it has had an opportunity to provide some reinforcing feedback from the environment" [Dennett, 1978, p. 76]. Trial and error learning is not only time consuming and limited. In a realistic environment it can also be dangerous. An organism would have a significant advantage over its competitors if it could somehow evaluate a piece of soft-wiring without resorting to overt behavior. A small step in this direction is the ability to make tentative responses to get the needed information - what Tolman [1948] called "vicarious" trial and error behavior. The behavior

is "vicarious" in the sense that an organism faced with a choice pauses to sample and compare the various alternatives. The response pathway chosen is the one that is deemed most desirable as a result of this overt comparison. This behavior implies the use of stored affective codes that can be retrieved in the presence of stimuli that are associated with them.

A more powerful strategy is to allow the codes to be used in the absence of supportive external stimulation. In other words, to provide for some kind of internal reinforcement that selects adaptive soft-wiring.

Ultimately of course it is environmental effects that are the measure of adaptivity and the mainspring of learning, but the environment can delegate its selective function to something in the organism ..., and if this occurs, a more intelligent, flexible, organism is the result. [Dennett, 1978, p. 78]

In its most primitive form, this internal selection could be a mechanism for "secondary" reinforcement where a stimulus associated with reinforcement from the environment acquires reinforcing properties of its own. The importance of this simple elaboration of the basic principles has long been recognized by psychologists [Hilgard and Bower, 1975].

An even more sophisticated mechanism arises when the internal reinforcement is not derived solely from external reinforcement. For example, the extent to which a given stimulus matches the existing representations in the system is an indication of how *novel* the stimulus is. Novel stimuli usually cause an increased alertness in an organism.

This increased alertness must accompany any change in the environmental conditions, any appearance of an unexpected (and sometimes, even an expected) change in those conditions. It must take the form of mobilization of the organism to meet possible surprises, and it is this aspect which lies at the basis of the special type of activity which Pavlov called the *orienting reflex* and which, although not necessarily connected with the primary biological forms of instinctive processes (food-getting, sexual, and so on), is an important basis of *investigative activity*. [Luria, 1973, p. 55]

One consequence of this activity, of course, is that learning is much more likely. In this sense, novelty is an internal criterion for identifying what is important. Other such criteria might include accurate predictions, "... well-ordered *preferences*, sound *plans of action*, in short all the favorite tools of the cognitive psychologist" [Dennett, 1978, p. 80]. The point to be made is that the basic principles established so far provide a framework in which many sophisticated processes can be described and, eventually, implemented.

#### Implementation as a Classifier System

Now that the goal-seeking system has been redesigned and a representational framework has been chosen, it is time to specify in more detail the way the system will be implemented. In simplest terms, the organism model developed in the last chapter will be augmented with two populations of classifiers: one to serve as the organism's perceptual component, the other to serve as the affective component. Two separate populations are used rather than one larger one so that the learning algorithms can benefit

from having classifiers already separated into gross functional "niches". Signals from the environment activate perceptual classifiers which in turn transmit messages to the population of affect classifiers. The resulting activity in the affect population, together with the central motive state and motor control hierarchy, determines what kind of overt response the organism makes.<sup>12</sup> Each population of classifiers is continually modified based on the organism's experiences in the environment. These modifications are designed to improve the system's ability to recognize objects in the environment and respond to them appropriately. If the learning heuristics are successful, the organism will acquire the same ability to function and survive that was built-in to the instinctive model given in the previous chapter. Implementing these changes requires that some of the basic definitions regarding classifier systems be broadened and that appropriate learning heuristics be developed and tested.

The Message List. One classifier system component that must be revised is the message list. To present the system with an undifferentiated list of stimulus signals is to assume that the organism is neutral with respect to the information contained in each signal and that all signals should be equally effective in terms of eliciting activity. This is a common, though often unstated, assumption in many

---

<sup>12</sup>Note that two sets of messages are processed before the system generates an overt response.



information processing models [Lachman and Lachman, 1979]. Natural organisms have an evolutionary history which predisposes them to be sensitive to certain adaptively salient stimuli. "Some concepts, percepts, and relationships should be easier to process than others because they have a longer history of adaptive salience" [Lachman and Lachman, 1979, p. 144]. Even though our hypothetical organism has no evolutionary history, its task environment does subject it to particular adaptive demands. The organisms's structure should have the capacity to be sensitive to those demands by selectively processing functionally important signals. Selectively processing one signal over another is one aspect of the phenomenon usually referred to as *attention*. Attention is a very complex issue [Norman, 1976] and no attempt will be made here to discuss its many facets in any detail. It suffices to say that the extent to which a given signal is likely to influence subsequent information processing depends, among other things, on both the physical and informational properties of the signal. Selection based on the information content of a signal is already evident in our organism's innate releasing mechanisms. The only physical property of signals in the environment that has functional significance is the signal intensity. The higher the relative intensity within a stimulus aura, the closer the organism is to a given object. In this sense, intensity is a rough comparative measure of how salient a signal is. Accordingly, the message list will

be modified to include both signals and their intensities. The relative intensities of signals on the list is a factor in the relative ability of each signal to elicit activity.

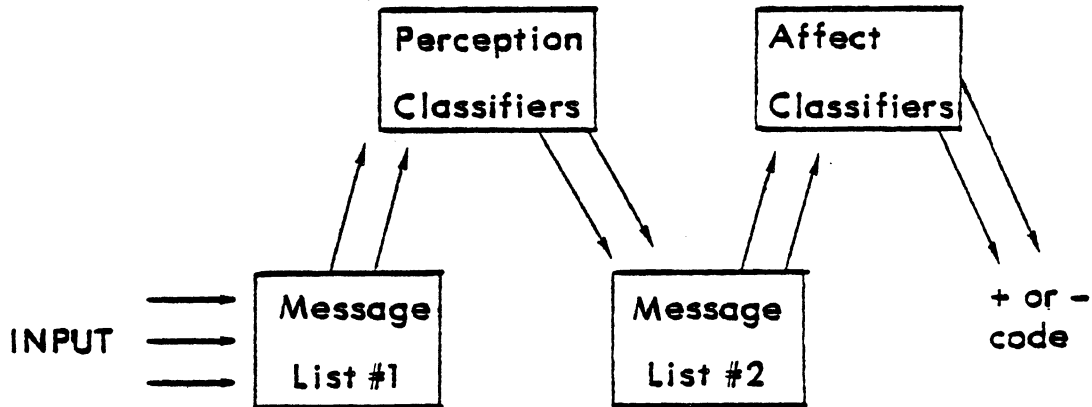


Figure 4.3. Implementation of the model using two message lists.

Because the system has two populations of classifiers, it is convenient to implement it with two message lists as shown in Figure 4.3. Signals from the environment are placed on the first list and determine the activity in the perceptual classifiers. The messages generated by this activity are then placed on the second list to be processed by the affect classifiers. It must be emphasized that the use of two lists is merely a convenient simplification. Holland [1980] points out that the definition of a classifier can be broadened to include a binary prefix as part of every input taxon and every message. This prefix can function much like an "address" for a classifier in the

sense that the classifier can only receive messages having an identical prefix. For our purposes only two prefixes would be required: one designating all classifiers in the perception population and the other designating classifiers in the affect population. If messages from the environment are given the perception prefix and messages generated by the perceptual classifiers all have the affect prefix, our organism can function with one message list - and, if desired, one combined population - in the traditional manner shown in Figure 4.1. However, because it is easier to analyse and debug a set of classifiers and messages that are physically separated according to their function, the organism will be implemented using two message lists.

Even though prefixes will not be used here, it should be noted that prefixes and "addresses" can be used to implement some very sophisticated information processing mechanisms. The simple use of prefixes proposed above can turn the message list into a global database analogous to the *blackboard* used in the Hearsay-II speech understanding system [Erman *et al.*, 1980]. Each population of classifiers can be thought of as a set of independent, specialized processors or *knowledge sources*. We can define a loose hierarchical structure over a population by extending the notion of an input taxon prefix to include a more specific kind of address. As before, all classifiers in a population have the same binary prefix. Now, however, each classifier also has its own particular address prefix which is a string

in the alphabet {0,1,#}. One classifier is said to be subordinate to another in the hierarchy if the set of strings matching its address is a subset of the set of strings matching the other classifier's address. For example, the address #101#0 is at a lower level than the address ##01##. Addresses at the lower levels have more 0's and 1's and are more specific in the sense that they match relatively fewer strings. Address at higher levels are correspondingly more general with respect to the set of strings they match. This addressing scheme is flexible enough to implement many kinds of control architectures for communication and cooperation among distributed processors. Such architectures have proven to be very useful in several practical applications [Erman *et al.*, 1980].

Moreover, this extended use of prefixes could allow our hypothetical organism to treat a set of signals as a structured array instead of as a mere list. Specifically, each input signal might be spatially tagged with certain "retinotopic" or "tonotopic" coordinates from the organism's sensory system. Messages at the output interface might be organized spatially as well, perhaps encoding environmental coordinates needed to direct a movement. Each classifier can now be thought of as having a meaningful relative "location" in the system: one classifier is close to another to the extent that the messages they process come from the same "neighborhood" in sensory coordinates. The "spatial" layout of activity at the sensory interface then has an

orderly relationship to the location of activity in the population of classifiers. In this way, classifiers sensitive to messages from the sensory interface are organized to preserve rough spatial relationships like proximity among incoming signals. Arbib [1972] has pointed out how this kind of "somatopic" organization of elements is a characteristic and computationally important principle of the way the brain processes information.

Classifier Strength. Implementing the revised organism model as a classifier system requires a broadening of the basic definition of a classifier. The input taxon of a classifier designates some pattern of messages that is presumably important to the system. The message taxon designates a set of responses appropriate to that pattern. Because the environment is uncertain, the system must have some way of discerning which patterns - among all the possible patterns that could be extracted from a set of messages - are the most probable and the most salient; and, which combinations of input and message taxa are the most useful in terms of generating behavior. In order to help provide that information, each classifier has a *strength* parameter associated with it. The strength of a classifier is a number estimating how well the classifier characterizes the situations in which it is a candidate to become active. Classifiers with high strength are more likely to become active than classifiers with low strength; and, when the system generates and tests new classifiers, classifiers with

high strength are less likely to be replaced.'<sup>1</sup>

A classifier represents a good characterization of a situation to the extent that it specifies an appropriate pattern and generates useful messages. It is important that a classifier's input taxon - the organism's internal representation of a pattern class - corresponds as closely as possible to the categorical prototype from which the signals are generated. In this sense, a taxon can be too general or too specific by designating a superset (too many #'s) or a subset (too few #'s) of the actual category. For example, if an object prototype is defined to be 1\*00\*\*...\* then the perceptual classifier 1####...# is too general a representation and the classifier 1000##...# is too specific. In order to bias the system away from classifiers that are too general, strengths are continually adjusted so that they are proportional to, and therefore estimate, a classifier's expected match score. In this way, specific classifiers accrue more strength than general ones and, relatively speaking, general classifiers are at a disadvantage. As we shall see later, there are other mechanisms at work in the system selecting against classifiers that are too specific. The result is a system biased toward classifiers at an appropriate level of generality based on the system's experience.

A classifier is said to generate useful messages if its

---

<sup>1</sup>Recall that a population of classifiers has a fixed size. When a new classifier is inserted, an old one must be deleted.

messages are instrumental in activating subsequent classifiers and/or enabling the system to obtain resources or avoid noxious stimuli. When a message is put on the message list, classifiers that match it may or may not become active. Active classifiers are chosen probabilistically so there is some uncertainty about which messages on the list will ultimately be successful in eliciting activity. Since messages are at the heart of the system's processing capabilities, classifiers whose messages reliably elicit activity are indispensable if the system is to generate consistent behavior. Such classifiers are good "predictors" in the sense that they indicate the most likely next "locus of activity" in the system. Obviously, it is also important that the sequence of activity generated by the system have meaningful consequences in terms of overt behavior. For the given environment, sequences of behavior leading to external reinforcement are the most significant. The most useful messages, therefore, are those that are reliably associated with a sequence of activity resulting in reinforcement. In order to formulate a classifier's "prediction" about the effects of its message, each classifier receives feedback about whether or not its message elicits activity and, ultimately, reinforcement. Strengths are adjusted to estimate the prediction of expected feedback and therefore bias the system toward classifiers that generate effective messages.

The strength parameter can be thought of as a kind of

long term structural code, activity being a more short term code. Over time, strength should be modified so that the most functionally valuable representations in the system have the highest strength. There are several simple heuristics available [Minsky and Papert, 1969] for adjusting this parameter to estimate the expected match score and feedback. The version used here postulates that, for each message on the message list, classifiers competing to become active lose a fixed fraction of their accumulated strength. In return, each classifier receives an amount of strength proportional to its current match score and feedback. More specifically,

Let  $s(t)$  be the amount of strength in a classifier on the  $t^{\text{th}}$  time it is evoked.

Let  $e$  be the fraction of strength lost in the process of competing to classify this message.

Let  $m(t)$  be the match score of the classifier.

Let  $M$  be the highest possible match score.

Let  $S$  be the relatively constant amount of strength made available to be recovered by a classifier.

Then the change in strength for a classifier is given by

$$s(t+1) = s(t) - e*s(t) + (m(t)/M)*S$$

It is easy to show (see Minsky and Papert [1969]) that, given an appropriate choice for  $e$ , the amount of strength in a classifier will be  $(m/M)*S$  where  $m$  is the expected match score.

This simple scheme can be extended to estimate feedback by making the strength recovered proportional to both the match score and the feedback. The most effective way to combine the two factors will be described in conjunction with the algorithm for generating new classifiers.

A sophisticated feedback mechanism is not necessary for our organism. Feedback is only required for the messages



generated by perceptual classifiers'<sup>4</sup> and all the message-passing "sequences" in the system are of length one. There are three pieces of information needed to evaluate the effect of one of these messages: First, some measure of how strong the "associative connection" is between the perceptual classifier and those affect classifiers responding to the message; second, an indication of whether or not the associated affect classifiers become active; and, finally, an assessment of how closely those affect classifiers represent the affective code indicated by the current reinforcement signal. The average match score of the message with the active classifiers it evoked - assuming the score is zero if it evokes no activity - provides the first two items of information. The third factor is summarized by the average "tag score" (see discussion of tags below) of the active classifiers. The product of these two scores is therefore sufficient feedback about the effectiveness of a message. Because of the extent to which our hypothetical organism depends on reinforcement from the environment, feedback is available only when a reinforcing stimulus signal is present. The expected feedback for a classifier is estimated by its strength as indicated above. Note, however, that a separate parameter is required to keep track of the expected feedback. Modifying strength directly will not work since the relatively infrequent presence of

---

<sup>4</sup>We shall see later how the affect classifiers have hard-wired connections into the motor control hierarchy.

feedback information would be mere noise in the background of the steady flow of input taxon match scores. This problem can be avoided by using a separate feedback parameter to modify strength and updating that parameter whenever feedback is available. This means that the system will be biased toward classifiers that produce effective messages.

Even though the feedback requirements of our hypothetical organism at this stage in its design are minimal and straightforward, it is important to indicate how more powerful feedback mechanisms could be implemented within the classifier system framework. In general, classifier systems are capable of generating arbitrarily long sequences of message-passing activity before making an overt response. Moreover, during the course of any behavioral sequence in a complex task, many classifiers will be activated before the system can expect to receive reinforcement. A classic difficulty under such circumstances is the so-called "credit-assignment" problem [Minsky, 1963]. Simply stated, the problem is how to apportion credit (or blame) for some behavioral outcome when a system uses many elements over several time steps to produce that behavior. When the number of classifiers involved is large, and/or there is a long delay between the activation of a classifier and the receipt of a reward, it is not practical to keep track of what each classifier did and when. Not only would this require large amounts of

storage, it would also require a sophisticated critical analysis of the entire sequence and the contribution of each classifier. The component of the system performing this analysis - the "critic" - would more than likely have to be very complex, an anomaly in a system designed to achieve computational power and simplicity at the same time. The only realistic alternative for a classifier system, or any "self-organizing" system, is to define reinforcement for sequences in terms local to the interacting elements involved and to discover and make use of reliable subsequences [Minsky, 1963].

Making predictions can be an effective way to identify the local decisions in a behavioral sequence that are responsible for each ultimate success or failure [Samuel, 1959]. The feedback parameter defined above designates a classifier's prediction about the effects of its message. This parameter can be redefined to include information about expected reinforcement from the environment as follows. Suppose that every classifier active when the system is reinforced has its feedback parameter increased by an amount proportional to the size of the reward. Suppose further that every time a classifier transmits a message its feedback parameter is adjusted based on two factors: the average match score of the message with the classifiers it evokes (if any); and, the average feedback estimate of those newly activated classifiers. This arrangement allows information about reinforcement to

be passed back to all classifiers involved in a message-passing sequence leading to a reward. Classifiers belonging to such a sequence will have a higher feedback estimate - and, therefore, higher strength - than classifiers that do not belong. Subsequences consistently leading to reinforcement will thus have high average strength while those only occasionally resulting in reinforcement will be correspondingly weaker.

Tags. Several aspects of the revised organism model implicitly assume that the system has some way of internally selecting one representation over another. Consider, for example, the way motivationally controlled stimulus-response pathways are facilitated. In the instinctive model the facilitory connections between a given motivational factor and the appropriate pathway were determined in advance. Identifying the functional significance of a pathway was never an issue. The soft-wiring in the revised model changes all that. There are no important representations or pathways in the naive organism for the motivational factors to connect to. As the organism experiences the environment and develops representations, the appropriate facilitory connections must be generated as well. This cannot happen unless the functional significance of a representation can be identified.<sup>13</sup>

It is important to note that this identification must

---

<sup>13</sup>Note that the input taxon does not do this. The organism has no *a priori* knowledge about which patterns are important or why.

be done at the level of the perceptual representation. To illustrate why, suppose that some motivational factor is at work selecting for behavior appropriate to stimulus category A. If the organism is confronted with a stimulus configuration involving two kinds of objects - categories A and B - the dominant perceptual activity will be, all other things being equal, determined by the object generating the most intense set of signals.' ' If that activity happens to be in response to category B, there is a contradiction between the "attention" mandated by the environment and the behavior mandated by the central motive state. Behavior directed at object A requires the activation of the representation for object A. The currently active representation is what the system is attending to and that is what must determine subsequent action [James, 1892]. If the organism is to use internal control in a way that is sensitive to the environment, the facilitory influences must be exerted at the perceptual level.

Another example of internal selection is the way affect classifiers come to designate affective codes. The organism is innately predisposed to interpret reinforcement signals as good or bad. If these evaluations are to be stored for later use, each affect representation must be labeled with an identifier that indicates its valence; and, the current activity must be biased toward representations having the

---

'Recall that the relative intensity of a signal is an important factor in determining how much activity will arise in response to that signal.

"right" code. In order to meet these concerns, the definition of a classifier is broadened once again to include a *tag*. A tag is simply a binary string that serves to identify a classifier to other parts of the system. In the case of affect classifiers, for example, the tag is the affective code. The instinctive centers for food seeking and pain aversion monitor the affect population for activation of the appropriate code. The occurrence of such an active code has a built-in releasing effect on the center in question. The need was expressed earlier for a *tag score* to measure how closely a given affect classifier designates the code indicated by a reinforcing signal. Since codes and tags are binary strings, a simple count of the bit positions at which they differ is an adequate way to measure their proximity.

More generally, tags can be thought of as coordinates locating each classifier in the perceptual-conceptual "space" of the organism. The dimensions of this space are innately determined and the location of a classifier in the space has functional significance. For instance, affect classifiers in one "region" of the space code for good events while those in another region code for bad events. Since classifiers in the same region have similar functions, facilitatory effects due to motivational factors or reinforcement can be thought of as built-in phenomena targeted to regions instead of specific representations. Facilitation biases activity in favor of one kind of

classifier versus another. This gives our organism the important capacity to have and use "dispositions" or momentary biases on processing [Minsky, 1979].

Revised Execution Cycle. Having specified all the changes required for the basic structures in our classifier system, we are ready to see how these changes are reflected in the system's execution cycle. On each time step, the system processes messages as follows:

- 1) Both message lists are emptied and all active classifiers are deactivated.
- 2) The *excitation* level of every classifier is set to zero. Excitation is a number indicating how much evidence there is that a classifier should be activated.
- 3) Signals from the environment are placed on the first message list.
- 4) The set of active perceptual classifiers is determined (see below).
- 5) Each active classifier places exactly one message on the second message list. If a classifier matched more than one signal, its message is generated by processing the signal with the highest intensity and match score. The intensity of the message generated by the most excited classifier is set to some fixed value  $k$ . The other messages are assigned intensities that preserve the relative excitation levels of the classifiers that generated them. So, for example,

classifiers with excitation levels of 1000, 700, and 500 will generate messages having intensities  $k$ ,  $0.7*k$ , and  $0.5*k$  respectively.

- 6) The set of active affect classifiers is determined (see below).
- 7) If reinforcement is available, all active classifiers have their feedback estimates revised as explained previously.
- 8) The affective code corresponding to the activity in the affect classifiers is used by the hierarchical control mechanism as it determines an overt response.

The following steps are involved in deciding which classifiers in a population are activated by the contents of the message list.

- 1) For each message on the list,
  - a) Determine the  $N$  classifiers with the highest match score. These are the classifiers in the system considered most *relevant* to the message.
  - b) Increment the excitation level of each relevant classifier by an amount equal to the product of three (or four) factors: strength, match score, message intensity, and incoming facilitation (if there is any). Sensitivity to facilitation is directly related to a classifier's tag score.
  - c) Update the strength of each relevant classifier as discussed previously.
- 2) Determine the  $M$  most excited classifiers and assign



each one a probability of becoming active based on relative excitation values. If  $E(i)$  is the excitation level of classifier  $i$  and  $E$  is the total excitation in the set of  $M$  classifiers then the probability that classifier  $i$  will be activated is simply  $E(i)/E$ .

- 3) Choose  $J$  active classifiers without replacement using this probability distribution.

There are two new concepts introduced here. First, there is the excitation level associated with each classifier.

Excitation can be thought of as a number indicating the mandate a classifier has to become active. Each message can engender an increase in excitation that is the simple product of the factors cited earlier as determinants of activity in the system. These factors include considerations of prior structure (strength and match score) and the current processing environment (message intensity and facilitation). The more messages on the list a classifier is relevant to, the more excited it gets.

Relevance is the second new concept. The relevant classifiers represent the set of options available to the system at any given time for processing a message. The limited time and resource constraints imposed by realistic environments dictate that this set not be complete and exhaustive. It has been argued previously that information about a pattern class is often best designated by a set of messages rather than a single, perhaps more complex message. Because of this, and the fact that the message list is

likely to contain information related to more than one pattern class, each message is processed by the most specific classifiers available. In this way, a population can be thought of as a group of specialists in competition for processing resources - that is, activity - and capable of operating in parallel. A similar use of specificity to determine which rule to activate has been made in other production systems (e.g. Anderson and Kline [1979]), but without the added flexibility of having more than one production activated, none of which has to match all of the messages.

There are two ways in which classifiers compete with each other to become active. For any given message, classifiers compete to get in the relevant set and accumulate excitation. This competition is in most cases won deterministically by the most specialized classifiers.' At the next stage, classifiers compete probabilistically for activation based on their total excitation. This stochastic process enables the system to function more effectively in an uncertain environment. In effect, each classifier is treated like a tentative hypothesis about the distal source of a message. The reliability and limitations of such a hypothesis can be discovered only if it is tested in various contexts and compared with various alternatives.

Summary. It is useful to review some of the important

---

'A complete discussion of the issues involved in choosing a set of relevant classifiers is given in the next chapter.

properties of the system as it has been described so far. It is clearly a system with a distributed memory. The system's knowledge about any given situation or category is represented by a set of classifiers. This representation captures much of the variability and correlational structure of a category in terms of the various combinations of taxa, tags, and strength within the group. Storage and retrieval of information can be accomplished either on a content basis - using messages and input taxa - or on the basis of a particular function or goal - using tags. Of course, both factors could be operative at the same time. The inherent redundancy in this kind of knowledge representation makes it resistant to the loss or destruction of an occasional classifier. This simplifies the problems that might arise if some "critical" element in the system is modified or replaced. There are no critical classifiers. The system can function adequately using any one of several alternative subsets of classifiers.

The system is also a model of parallel, distributed processing. More than one classifier can be active at a time. Because of the way excitation is computed, behavior in the system is determined by the locations where several messages "converge" to induce a high probability of activity. In this sense, classifiers can increase their effectiveness by cooperating with other classifiers - each one sending its message to the same part of the system. Competition is another important aspect of the way

classifiers interact. The competition mechanism enables the system to generate and test a set of hypotheses derived from experience. Holland [1980] points out that competition has another important function.

The competition mechanism provides more than a set of hypotheses ... based on experience. It allows the [learning] algorithm to inject new classifiers into the system with a minimum of disruption. The new classifiers must win competitions to be activated and tested. Typically, this will happen only when the system is inadequately handling some [situation] so that the scores of the classifiers competing in this context are relatively low. Then the new classifier has a chance of winning the competition. If, upon winning, it provides an improvement it will become one of the preferred classifiers in that context. (p.266)

This means, in particular, that the system changes gracefully - it can be modified without incurring disastrous consequences for behavior. Similarly, performance changes gracefully even in the face of potentially disruptive events such as the loss or distortion of a message. All of these properties work together to give the system a useful "common sense" knowledge representation of its environment [Kuipers, 1979].

Tags give the system another interesting capability. If we assume that tags have the same length as messages, then the tags from active classifiers in one part of the system can serve as messages for classifiers in another part of the system. This means that the system has the ability to *recognize and act on the basis of its own internal patterns of activity*. This is a very sophisticated ability. For example, it allows the system to integrate the activity

generated by patterns and objects into representations of higher order entities like scenes. It also allows hierarchical or heterarchical processing to be accomplished easily, automatically, and parsimoniously. Top-down and bottom-up processing can be implemented as well. Tags used in the manner described here make classifier systems a very powerful computational tool.

Finally, the statistical processing leading to patterns of activity as described here is analogous to the way the brain has often been characterized as an information processor [Hebb, 1949; Lashley, 1951; von Neumann, 1958; John, 1967]. There are certain parallels that can be drawn between classifiers and structural representations - or "cell assemblies" [Hebb, 1949] - hypothesized for the brain. Both are generic pattern sensitive elements derived from experience with the environment. They are activated and processed in a statistical manner in a network of rich associations. Both make extensive use of strength, proximity, facilitation, and competition in their respective computations. There are also, however, some significant differences. Cell assemblies have many partially active states that allow them to process patterns extended over time; and, their activity can be influenced by large amounts of barely discernable or "fringe" activity. What a classifier accomplishes with tags a cell assembly achieves more elegantly with its physical structure. Moreover, while a classifier system requires explicit computations for



matching and relevance, cell assemblies are activated automatically by structures that allow for switching and convergence of activity. The implications of these similarities and differences are not clear. The analogy is worth pursuing, however, as a subject for further research.