

## Making predictions in an uncertain world: Environmental structure and cognitive maps

ERIC CHOWN

*Bowdoin College*

This article examines the relationship between environmental and cognitive structure. One of the key tasks for any agent interacting in the real world is the management of uncertainty; because of this the cognitive structures which interact with real environments, such as would be used in navigation, must effectively cope with the uncertainty inherent in a constantly changing world. Despite this uncertainty, however, real environments usually afford structure that can be effectively exploited by organisms. The article examines environmental characteristics and structures that enable humans to survive and thrive in a wide range of real environments. The relationship between these characteristics and structures, uncertainty, and cognitive structure is explored in the context of PLAN, a proposed model of human cognitive mapping, and R-PLAN, a version of PLAN that has been instantiated on an actual mobile robot. An examination of these models helps to provide insight into environmental characteristics which impact human performance on tasks which require interaction with the world.

**keywords:** cognitive maps, navigation, associative networks, gateways

The world is a complex and changing place and yet humans are able to quickly adapt to a wide array of new environments. This does not stem from an innate ability to exactly memorize every facet of a new environment; on the contrary, it is much more likely that humans are able to efficiently extract and organize the minimum information they are likely to need. This article examines the mechanisms for this efficiency and when these mechanisms do and do not work. The premise of the article is that humans, through the process of evolution, have learned to exploit structures common to many environments. The discussion in the article focuses upon environments of the sort involved in navigation, but the scope is more general, even including abstract environments and other organisms. The article relies heavily upon the example of human cognitive maps, but this is because they exemplify so many representational principles and because so much of cognition is spatially based. A central topic is how an organism, when faced with a completely new set of information, can effectively organize that information in a useful fashion. Even before then, however, the organism must first be able to recognize that it is faced with new information, rather than simply a new view of old information.

An organism which only interacts in a single niche can evolve specialized structures to maximize its efficiency within that niche; however, as organisms adapt to more and more environments, some of this efficiency must be sacrificed for general purpose structures. Success in a new niche still requires efficiency, and indeed humans are able to quickly identify new niches and the salient components within them. The knowledge structures involved in human navigation, called cognitive maps, are good examples of representations which are both general purpose and efficient. Humans are able to effectively navigate in environments ranging from oceans to forests; in most cases quickly learning an individual niche, even a seemingly novel one. The thesis of this article is that the basis for both the efficiency and the generality of cognitive maps results from the human ability to recognize and exploit common structures which occur in widely disparate niches. Further, these structures are not limited to the navigational domain, but are found in abstract environments as well. This has several key implications: 1) The reliance upon generic structures means that niches rich in such structures will be easy to learn, and that niches lacking in it, conversely, will be more difficult to learn; 2) Designed environments, whether they be buildings or computer programs, can be tailored to take advantage of this; and 3) Learn-

ing systems, such as mobile robots or artificial organisms, can be devised to take advantage of these structures as well.

The guiding principal for this work is that an adaptive organism's internal representation of a niche must reflect the salient features of its experiences within the niche. Since saliency cannot be rationally determined a priori, there must be general heuristics which afford strong clues as to what aspects of an environment should be learned and how the information should be organized. Further, uncertainty is always a problem in learning new environments because the world is constantly changing and because information is often incomplete or noisy. These issues are central to any organism dealing with a new niche, but are particularly relevant to organisms which interact in as wide a range of environments as do humans. The claim in this paper is that humans do not come up with a new strategy for learning in every new environment that they encounter, rather that there are basic principles and structures which are useful across all of the environments that are likely to be encountered. The basis for discussion of these principles are a pair of models called PLAN (Chown, Kaplan & Kortenkamp, 1995), and R-PLAN (Kortenkamp, Chown & Kaplan, 1997). PLAN (Prototypes, Locations, and Associative Networks) is a model of human cognitive mapping and R-PLAN (Robot-PLAN) is an instantiation of PLAN on an actual mobile robot. This article could be considered the third of a sequence in which 1) a model of human cognitive mapping was presented, 2) that model was extended for use by a mobile robot, and 3) it is shown that the structures of the model are further generalizable, applying even to abstract environments. Since this article relies upon the models presented in the previous papers, the essential properties of those models will be reviewed over the course of the article.

This research is based upon the premise that operating in real environments has guided and constrained the evolution of human mental representations. If the claims of this article are true - that the ability of humans to learn new environments is reliant upon generic properties and structures of environments - then there are implications in terms of human functioning. In fact there is already strong evidence for this; the literature linking different environmental configurations to health and performance is substantial (for extensive discussions and reviews see (Kaplan & Peterson, 1993; Kaplan and Kaplan, 1988). The primary thesis of this literature is that some environments (generally natural environments) enhance functioning because they resonate with human cognitive structure. For example, many studies have shown that natural environments reduce attentional demands on humans; further, experiencing

these environments can have a kind of restorative impact cognitively, leading to improved performance (Kaplan, 1993; Tennessen & Cimprich (1995)) and even health (Cimprich, 1993). Closely related to this work is the idea that people are more at ease in some environments than in others (Weisman, 1981; Kaplan & Kaplan, 1989). In many respects this article is a bridge to such work by examining the relationship between environmental and mental structure.

It has been argued by numerous researchers that spatial structures and spatial reasoning are fundamental to much of cognition (Lakoff, 1987), and spatial cognitive structure, or cognitive maps, serve as the basis for most of the discussion in this article. One of the purposes of this article is to show how the same principles exploited in spatial representation are generally applicable. Apparently disparate environments can be organized cognitively using a spatial framework. More and more computer interfaces, for example, are designed using spatial principles; the main tools for using the Internet are even called "Navigator" and "Explorer."

The PLAN model will serve as the basis for the examination of spatial representational structure. We have proposed several distinct spatial structures, each of which exploits a different environmental characteristic. Numerous researchers in developmental and spatial cognition have divided spatial acquisition into several distinct phases (Shemyakin, 1962; Piaget & Inhelder, 1967; Siegel & White, 1975), each corresponding to one of the structures in PLAN. The first phase concerns the most basic element of environmental structure, that environments tend to be collections of distinct objects. In cognitive mapping these objects are generally called "landmarks," and in PLAN they are represented by prototypes. In the next level of the developmental hierarchy, structures called "route maps" are acquired. Route maps reflect another structural property of environments: that the objects which make up an environment are experienced sequentially. In real environments this typically occurs as an organism moves through the environment, but it can also occur as the organism looks from one object to the next, or it can occur as a simple result of the changes that naturally occur with the passage of time. In PLAN, route maps consist of associative networks of prototypes. The final phase of development concerns the acquisition of "survey maps," representations which appear to be more directly visual in character and which capture a larger spectrum of spatial relationships than are contained in the route maps. Because survey maps are visual in character, they require that there are particular locations in environments, called "gateways," which are especially suitable to anchor the representations. These three elements - objects, sequence, and gateways - form the core of the

representational theory presented here. Of the three, objects and sequence have long been studied, while the gateway construct is relatively new.

PLAN is similar to a number of other computational models of cognitive mapping, notably TOUR (Kuipers, 1978) and Navigator (Gopal, Klatzky & Smith, 1989), with regard to the underlying theory of route maps. Much of the analysis of those parts of the representation applies to those systems as well (for a more thorough comparison of these and other systems see (Chown, et al, 1995)). Differences include the use of scenes to create survey maps, a feature that other cognitive mapping representations lack, and in the fact that R-PLAN has been instantiated on an actual mobile robot. In addition the gateway construct represents a significant step towards understanding the relationship of the environment to its representation because it provides the basis for hierarchical cognitive structure.

The article begins with an examination of two of the major paradigms for dealing with environments, the behavior-based and knowledge-based approaches, and proposes that a kind of synthesis of the two is both desirable and possible. The major purpose of this section is to introduce some of the issues involved in trying to cope with real world environments. In the two subsequent sections the essential components of the PLAN model are presented in light of the proposed synthesis. In these sections the emphasis will be on the strong relationship between environmental and cognitive structure and the resultant structure's ability to cope with uncertainty. The critical factor involved in balancing these approaches is the proposition that there is enough structure inherent in most environments to make internal models worthwhile even though there is enough uncertainty to make precise models impractical. Finally, a more general framework is discussed, which travels beyond the scope of navigation. Ironically, although it shares the same basic structure as PLAN, R-PLAN is central to this discussion because it is an example of how the general principles work effectively even though a robot's abilities are substantially different from those of a human.

## STRATEGIES FOR INTELLIGENT ENVIRONMENTAL INTERACTION

The complexity of the real world makes intelligent interaction a difficult proposition. At any moment in time an agent interacting in a given niche is faced with the problem of what to do next. Such an agent has two main sources of information in order to make decisions: 1) a model of the world based upon prior experience, and 2) immediate sensory information (a third source

is also possible, stemming from evolution, e.g. the tendency to move to high ground when in danger). While these information sources need not be mutually exclusive they are sometimes treated as such.

### The knowledge-based approach

The classical, or "knowledge-based," approach to interacting with the world can be characterized as trying to build as complete and accurate a model of the world as possible and then using that model as the basis for all plans and actions. The emphasis on complete world models is both a strength and a weakness of the knowledge-based approach.

If a complete and correct internal world model can be assumed then the possessor of such a model is free to plan and reason about its environment to any desired degree. Of course this begs the question of the cost and complexity of such reasoning. The potential value of internal models has been stated elegantly by Craik:

"If the organism carries a 'small-scale model' of external reality and of its own possible actions within its head, it is able to try out various alternatives, conclude which is the best of them, react to future situations before they arise, utilize the knowledge of past events in dealing with the present and future, and in every way to react in a much fuller, safer, and more competent manner to the emergencies which face it," (Craik, 1943, p. 61).

This is a powerful statement, reflecting several important aspects of models. First, models are 'small-scale.' Too much information can sometimes be as paralyzing as too little. A useful model contains *relevant* information. Second, models are predictive. If an organism can predict a negative outcome in its head, it can avoid taking a dangerous action. Finally a model is synthetic; past experiences are blended together in such a way that the usage of the model generalizes beyond specific experiences.

Because of the potential power of models, it is not surprising that so much research in cognitive science has focused upon this aspect of cognition. Further, because understanding object recognition and perception has proven to be so difficult, it is not surprising that AI has generally focused on knowledge representation. Systems with limited perceptual capabilities must necessarily rely upon internal models because they cannot efficiently extract needed information from the environment. A robot with noisy sensors but an accurate map of an environment might be able to navigate effec-

tively relying almost completely upon its internal model. Since the environment is not providing information, the internal model must necessarily be as complete and accurate as possible.

However, there are many problems with this philosophy. Among them are the problem of how such knowledge structures are developed (Bates and Elman, 1992), the inherent brittleness of such systems (Holland, 1986), and perhaps most telling, the fact that even if such perfect knowledge is achievable, putting it to use is an intractable problem (Chapman, 1987). All of these lines of research illustrate that real environments demand systems that work, even if not optimally. Bates' and Elman's work is a reminder that organisms must be able to effectively cope in a new niche long before they have had enough experience to build a complete model of the niche. Further, learning is not simply a matter of adding new information, but often consists of building entirely new structures on top of existing ones. This is closely related to issues Holland has explored; he has searched for knowledge structures which apply to ranges of situations and which can be used as building blocks for other structures. Chapman's work on planning shows that optimality in complex environments is not feasible given the inherent complexity of such environments.

Another problem with the search for complete and precise representations is that perfect knowledge is not possible for systems that deal in extended time. The world is inherently uncertain; what is true now may not be true in a moment. Further, information normally available may at any moment in time be degraded, impoverished, or obscured as environmental conditions change; for example as fog descends on a valley. In recent years, knowledge-based AI has begun to combat these problems in a variety of ways, but generally by placing more emphasis on learning. As AI systems become more learning based, they necessarily make fewer assumptions about perfect knowledge; for example, a great deal of work in reinforcement learning deals with systems where the effects of operators are stochastic (Kaelbling, Littman & Moore, 1996). On the other hand, such systems do not plan in the usual sense. Instead of developing a long term strategy they merely choose actions one at a time. One of the challenges for such systems is to adapt planning methods to worlds where the transition from one state to the next cannot always be predicted with complete certainty. For humans, planning does not seem to be an all or nothing proposition; plans do not need to be precise, but they may extend into the future. There is a balance between the uncertainty involved and the knowledge which might be useful.

In summary, placing too great an emphasis on building precise and complete models is unlikely to yield sys-

tems which perform effectively in real environments, on real tasks. Knowledge structures which emphasize principles such as universal computation, information maximization, and other mathematical artifacts may be useful in the abstract but are not necessarily practical. After millions of years of evolution it is reasonable to assume that internal representations are optimized for the specific types of environments, or niches, that an organism is likely to find itself in. A thesis of this article is that an organism which can exploit a wide variety of niches, as humans do, cannot specialize in each niche, rather it must recognize and exploit patterns which occur across niches. This presupposes that there are common properties in seemingly disparate niches. Such optimizations will also necessarily bring tradeoffs, meaning that there will be niches in which the organism will not fare well.

### **The Behavior-based Approach**

The "behavior-based" approach in its purest form ignores world models altogether. The most often quoted reason for this is that "the world is its own best model." One of the major proponents of this position, Rod Brooks, argues that the problem of uncertainty is so great that it cannot be overcome through the use of models (Brooks, 1991). Behavior-based systems do not rely upon high level world models. Instead, they count on the environment to provide any necessary information. While this position is often presented as being somewhat radical and new it actually has a great deal in common with behaviorism and the theories of J.J. Gibson. Gibson, like the behavior-based group, postulated that there was no representation of the world in the head (Gibson, 1966). Instead, different environmental configurations caused the brain to resonate in different ways. Learning about an environment, therefore, would not be akin to building a model, but more along the lines of tuning internal responses.

The advantages of such systems are obvious. Because they are not tied to internal models and are essentially completely reactive, they do not need to access or process representational structures and are, therefore, faster than systems that do. While they still must deal with problems of noisy sensors, uncertainty is not as much of a problem because they only react to things that are directly perceivable and, therefore, do not rely upon beliefs and predictions.

For the most part the behavior-based approach has been limited to mobile robots, where such systems have achieved a great deal of success (Brooks, 1986; Connell, 1992; Madadevan & Connel, 1992; Horswill, 1993; Bonasso, Kortenkamp & Miller, 1995). Despite such success, however, it is not clear how far this approach

can be pushed. Returning to Craik's quote, it would appear that a system without internal models is at a serious disadvantage versus one which does model its environment.

### Striking a balance

Storing all possible information is not a viable option for a naturally intelligent system. Given a relatively small brain in proportion to the enormously complex world around us, it is essential to economize. As Clark (1989) points out in describing his "007 Principle," "Know only as much as you need to know to get the job done," (p. 64). On the other hand, doing without stored information altogether would be a serious error. The benefits of some sort of coherently stored information, of something approximating an internal model of the environment, are many. Chief among them is the ability to predict environmental events before they occur by using information from past experience to go beyond what is present in the sensory array (Bruner, 1957). Further, internal representations can reduce perceptual computation; recognizing a familiar place, for example, means that it is not necessary to look around to know what can be found.

There would seem to be great advantage in finding a middle ground, a way of functioning based on having "as much as you need to know" and for the rest depending on the information presented by the environment. In theoretical terms this means finding an integration of these two apparently antithetical approaches. These positions represent extremes, of course, and recent research trends in research on adaptive systems and putting agents such as robots in real environments have necessarily led to an increased understanding of the synergy necessary between the environment and representations. A closer look at Clark's principle suggests some interesting directions for achieving such a synergy. Clark's principle as stated suggests two interpretations. "As much as you need to know," at one level implies that if information is readily available from the environment, there is no need to store it. But it can also mean that one need not usually store details; that generic structures are an economical means of handling most situations. This would appear to be an accurate characterization of cognitive maps as described by Kaplan (1973):

"The cognitive map is a construct that has been proposed to explain how individuals know their environment. It assumes that people store information about their environment in a simplified form and in relation to other information they already have. It fur-

ther assumes that this information is coded in a structure which people carry around in their heads, and this structure corresponds, at least to a reasonable degree, to the environment it represents. It is as if an individual carried around a map or model of the environment in his head. The map is far from a cartographer's map, however. It is schematic, sketchy, incomplete, distorted, and otherwise simplified and idiosyncratic. It is, after all, a product of experience, not of precise measurement."

The next two sections describe the representational components of the PLAN model, focusing on the basic tension between the ability to predict and the inherent uncertainty of the real world, and how both of these are tempered by experience. The next section is concerned with the most basic element of any niche, the objects which make it up, and how those elements are structured in a mental representation. In the subsequent section, it is shown that regular properties that appear in a wide variety of niches can be exploited to ensure the effectiveness of these representations. Just as an organism must be able to discriminate between uncertain and predictable events in a niche, it is also advantageous if it can exploit general patterns that are found across niches.

## REPRESENTATIONAL ELEMENTS - OBJECTS AND SEQUENCE

### Objects - landmarks

Objects form the core of most modern theories of cognition, and significant portions of psychology and computer science, e.g., categorization and machine vision, are directly concerned with the issue of how objects are learned. In the cognitive mapping literature the relevant objects are called landmarks. Landmarks function in navigation as a kind of marker or index which can be used to determine where one is in an environment. As such, for an object to be a useful landmark, it must be stable within the environment and it must be uniquely identifiable within its niche. An object which moves from one spot to another is not useful as a tool for identifying location, nor is an object which is just like all of the other objects. Thus, as one of the basic elements of environmental structure, building a representation based upon landmarks requires that niches are made up of uniquely identifiable and stable objects. This is the first, and most basic, example of how internal

representations rely upon recurring environmental patterns. It is possible to imagine alternative representational schemes which do not require objects; for example a navigation system could be based purely upon an absolute coordinate system. The implications of a system based upon landmarks will be discussed later in this section.

Not all of the objects found in an environment are stable, of course. Such dynamic objects can be quite important to understanding the environment, but representationally they are separate. Dynamic objects interact with the environment, but are not directly a part of it. In real environments a dynamic object might be something like a car, which is used to interact with the environment. In an abstract environment an example would be one's opponent in a chess game. The representational distinction will be discussed subsequently in the section dealing with links between objects.

Although a landmark's location must be stable, it will not be stable with respect to a given observer since people generally interact with environments by moving through them. One consequence of this fact is that recognizing that a landmark seen on one trip is the same as one seen on a different trip is far from a trivial task (Kortenkamp, et al. 1997). In this respect and many others, landmarks are no different than any other object. One of the most difficult aspects of dealing with objects is the problem of recognizing them across multiple perspectives. A requirement of any object representation scheme is that multiple experiences with a single object or landmark, often at different orientations and at different distances, must be recognized as being with a single object. For this reason these representations cannot be highly specific, but must reflect a variety of experiences. Having more generic representations such as these carries other advantages as well.

Basing a knowledge system on generality has great potential pay-off. Fewer representations can cover more ground, building in economy. Uncertainty management is enhanced, since degraded and incomplete stimulus patterns are more readily handled by a system that focuses on neither specificity nor detail. And speed of retrieval will benefit from the far smaller storage requirements of a system that maps many instances of the same thing into a single representation.

Generic representations are, of course, far from novel entities in cognitive science. Over the years they have been framed in a variety of ways and referred to, for example, as categories, schemata, or prototypes. PLAN uses the latter construct for reasons that are not essential to the present discussion. PLAN actually contains two types of units. The first type corresponds to what is usually thought of as an object (or landmark). These units are taken to be cell assemblies, recurrent struc-

tures that have the capacity to hold activity, to persist in time after the cessation of the input (Hebb, 1949; 1980; Kaplan, Sonntag & Chown, 1991). The second type of unit corresponds to a spatial representation of a scene which in turn might contain links to objects of the first type; this type of unit will be discussed in much more detail in the next section. R-PLAN uses a third type of representation, based upon vertical lines. In this case the prototype construct appears to be powerful enough that it works even upon such a primitive implementation.

One of the difficulties in using this type of structure in a real system, such as a mobile robot, is that the available information about the structure of such entities is mainly descriptive. For example, landmarks are often described as being perceptually unique, but there are no available algorithms to make such a determination. The PLAN model calls for generic representations based upon cell assemblies and though models of cell assemblies are available (Kaplan et al., 1991) they necessarily make the abstraction that they are provided with information already processed by the visual system. In short, there is still a long way to go before such representations can be automatically implemented. On the other hand, experience with the R-PLAN model has shown that even an extremely limited version of a generic representation can provide a robot with a great deal of perceptual power. The landmarks created in R-PLAN are created by extracting vertical lines from visual images and creating a Bayes-net to probabilistically determine whether a given input set of lines corresponds to stored representations. Even though this representation uses only a fraction of what might be considered useful visual information, it is enough to enable the robot to navigate interior environments.

### Sequence - networks

In the early developmental stages of children (and for adults in novel environments), cognitive maps are essentially a topologically structured collection of landmarks (Piaget and Inhelder, 1967; Siegel and White, 1975), primarily capturing the spatial relationships proximity and order, but lacking other spatial properties such as distance and direction. Although this structure stores only a fraction of the spatial information available, potential journeys can easily be planned as going from one landmark to the next. The topological structure is an especially effective way of capturing the inherent sequentiality in navigation. Further, this sequentiality extends beyond the domain of navigation. Perhaps the two primary modes of experience which natural systems have are locomotion through space and observing events from a single vantage point as they unfold in

time (even a single, stationary, object is generally experienced sequentially through eye movements). In both cases the experiential stream is similar, namely sequences of things one after another. Since the world is not random, such sequences will often be patterned, providing clues to what Brunswik (1956) referred to as the “causal structure of the universe.” Since storing such sequences permits prediction, it is information of considerable value (Macphail 1987). Thus, rather than storing information in lists or other arbitrary structures, an efficient method of storing objects is in terms of their sequential relationships to each other. As more and more sequences are experienced, they will begin to overlap, resulting in a network structure. In PLAN (and other cognitive mapping models) the objects in these networks are landmarks. Not only does this network structure reflect how the knowledge was acquired, it also reflects how it is likely to be put to use. Smolensky has called systems of this type *semantically transparent* (Smolensky, 1988) because they do not require additional computation either to store or to interpret.

While a network structure is informationally sparse, lacking information on direction and distance in the spatial domain, among other things, it is nonetheless functional. For example, if an environment is represented as a network of landmarks, the problem of getting from one place to another is transformed into the problem of finding a connecting path in the network between the two places. For such a problem any number of search techniques, such as spreading activation, breadth first, etc. can be used. R-PLAN, for example, uses spreading activation search to determine a sequence of landmarks which will lead from the robot’s initial location to its goal (this portion of PLAN and R-PLAN is actually based upon Levenick’s NAPS system (1991)). Spreading activation search consists of activating the nodes in the network which correspond to the start and goal nodes and propagating activity through the network. Nodes which are active will in turn activate nodes to which they are linked. The result is that activity will spread out from the two initial nodes until the two waves of activity coalesce. Once the waves have intersected a potential path has been found (for more details see (Levenick, 1991)). Contrast this with a typical cartographic map which contains much more information but can be difficult to use; finding a plan with a cartographic map relies upon abstract reasoning, while planning with a network can be done automatically. Further, the plans created with a node and link structure rely only on being able to get from one landmark to a nearby one. Studies comparing the plans created by such systems have found to high degree of similarity to plans formed by humans in the same environments (Byrne, 1979; O’Neill, 1990; 1991). Such plans include the

built in assumption that the environment will provide much of the necessary information at execution time. In this case the information needed is the heading towards the next landmark in the sequence. By contrast, a plan with the more information rich cartographic map might rely upon exact spatial heading or distances.

There is one other key advantage of using a sparse, semantically transparent, network structure for knowledge representation and planning. With such a representation a given plan consists of a sequence of smaller plans. In the navigation domain each of these plans consists of going from one place to another place nearby. Detailed knowledge, in the form of precise headings and distances, is not required in the execution of these small plans since it can be directly acquired from the environment. In this way such plans are virtually immune to errors provided that the perceptual system can accurately identify the necessary landmarks. A system using a cartographic-type spatial map, on the other hand, requires not only an extremely accurate representation, but also requires precise execution necessitating accurate headings and distances, a dangerous situation when over a long distance even a small error in an initial heading can lead to an extremely large error in a final destination. More detailed discussions of these issues can be found in (Brooks, 1985), among other places.

### Links

One of the key properties of landmarks is that their absolute location is stable. Unfortunately, environments change, and a tree that was a useful landmark one day may be gone the next. Therefore, to effectively cope with a niche over time, internal representations, such as cognitive maps, must be able to adapt to changing environmental structure.

PLAN is a connectionist knowledge structure. The nodes are cell assemblies (which actually are representations considerably more complex than the nodes found in many connectionist systems), and the links between them are variably weighted through a compensatory learning rule based upon Hebb’s original learning rule. Hebb’s learning rule essentially postulated that two neurons which were active simultaneously would become more strongly connected. In PLAN this rule applies at the level of nodes. When two landmarks are experienced near each other in time the connection between their representations is strengthened. As a sequence is repeated the relevant links continue to be strengthened, reflecting the increasing certainty that the close spatial proximity is a stable relationship. Furthermore, the learning rule is compensatory - as some links out of a node are strengthened, others are automatically weakened. This means an object that is only temporarily in an en-

environment will only become weakly connected to it, and after it is no longer in the environment the connections will eventually tend to drop out. Such a system places a high premium on stability and experience.

The compensatory learning rule has a number of consequences with regard to how environments are learned. First, such a rule is recency based; new learning tends to automatically expunge old learning. A heuristic version of such a rule would postulate that the best predictor of the state of an environment is the most recently experienced state of the environment. Of course this is tempered by the long term structure present; a well learned environmental feature will tend to remain in a representation for some time after it is no longer present, much as amputees often experience “phantom limbs” for extended periods of time after losing an appendage. One interesting question in designing learning systems is how to tailor a compensatory learning rule to the kinds of environments it will be used in. Oddly enough, the rate at which elements should drop out of a representation does not seem directly related to the relative stability of the environment. For example, in highly unstable environments, one might prefer a very responsive learning rule in which old elements dropout quickly. On the other hand, in environments which are extremely stable a responsive learning rule might also be useful because once an element is gone it will be unlikely to reappear. Indeed the reappear factor would seem to be central to this issue. If environmental elements tend to appear, disappear, and reappear at irregular intervals (for example a car in a driveway) then a less responsive learning rule would appear to be more logical than in an environment where once an element is gone it is probably gone for good.

There are other advantages to variably weighted links. In a spreading activation planning system, such as is used by PLAN and R-PLAN, the increasing ability to propagate activity through high-strength links means that the shortest path between two points will not always be the one selected. Instead the paths selected will tend to favor those sequences that have been experienced the most times. Such a policy not only tends to reduce uncertainty, but, assuming that routes that have been taken numerous times are safer than new routes, also places a premium on safety over optimality.

In summary, the simple connectionist learning rule of strengthening links with repeated experience is a natural method of coping with uncertainty. The links represent a sequential relationship and the strength of the links comes to represent the certainty of the association.

### The compensatory nature of input and structure

In most knowledge-based systems there is a substantial apparatus, apart from the knowledge structure per se, whose function is to organize, compare, and otherwise manage the knowledge in the system. An architecture consisting of networks of cell assemblies affords a more economical solution to this problem. Rather than having the knowledge managed by other agents, in such an architecture the elements of the knowledge structure themselves determine the course of processing.

Cell assemblies are semiautonomous (Hebb, 1963) activatable representations of reality that do not rely uniquely on the sensory interface to become or to stay active. Cell assemblies have an autonomous character because a system of cell assemblies processes by propagating activity within its structure rather than by moving data from storage into a CPU for processing. At the same time such autonomy must be tempered. One way is by preserving the connectionist notion of nominating percepts through competition, while enhancing such competition by placing it in a far more symbolic context. Another is through the use of internal control mechanisms which preserve the sequential nature of the structure by imbuing a natural time course of activation on any given symbol.

The time course of activity of a cell assembly provides an example of the potential power of a semiautonomous process. A given cell assembly will tend to resonate to particular types of environmental stimuli. If a particular stimulus is strong enough, neurons within the cell assembly will become active. At this point the structure of the cell assembly provides natural feedback mechanism to enhance and sustain the activity (subject to competition from other cell assemblies). As the cell assembly remains active it expends resources (such as transmitter substances) and eventually fatigues and begins to shut down. These processes impose a natural time course of activity, which in humans is thought to be about five seconds. This time course impacts numerous aspects of cognition including the number of elements which can be held in working memory at one time. Further, this time course can be altered by other cognitive processes, enhancing the flexibility of the system especially with regard to learning. Much more detail on these issues can be found in (Kaplan, et al. 1991; Sonntag, 1991; Chown, 1994).

As in other connectionist representations, environmental input and internal structure have a compensatory relationship in a cell assembly structure. Incomplete knowledge can be supplemented by information picked up from the environment. Impoverished environmental information can be filled in by previously



acquired knowledge. Such an arrangement would seem consistent with the conjecture that a modest sensory system and a modest internal model building capability, properly integrated, are likely to manage uncertainty far more effectively than even a highly developed version of either functioning alone.

### Networks and environments

The network architecture is based upon two important properties of the environments that are likely to be encountered: 1) they are made up of identifiable objects, 2) the objects are experienced in recurring temporal patterns. There are numerous examples in nature of organisms whose interactions with the world are not based upon these properties; for example, it is common for organisms to react to very specific, hard-wired stimuli. The payoff of a network architecture is that it generalizes to many potential niches. To be effective these networks require that a given niche consists of uniquely identifiable, positionally stable objects which occur relatively close to each other. A topological structure such as a network is based upon proximity. This can most clearly be seen in the action of the learning rule which is based upon simultaneous activation. What this means is that in environments in which objects are sparse, such as oceans or deserts, this type of representation will be relatively ineffective. This is not to say that humans are unable to navigate these environments effectively. Gladwin provides a beautiful counterexample in his discussion of Puluwat navigators (1970). What is fascinating about these navigators is that, even though the Pacific Ocean would appear to be relatively devoid of landmarks, their functioning highlights some of the essential properties of the network representation. For example, the Puluwat navigators make effective use of “seamarks,” breaking long voyages up into series of shorter ones. However, while some of these seamarks are simple, such as reefs, others require sophisticated pattern recognition, such as the ability to distinguish patterns of interference set up in waves by nearby islands. Further, even while the Puluwatans have developed an efficient method of dead reckoning based upon sun and star positions, they rely upon it as little as possible. For example, the navigators have three known routes to the island of Satawal, an island about 130 miles distant, each varying in the amount of dead reckoning required. One is essentially a straight course to the island, the next involves a stop at the island of Pikelot 100 miles along the way. The third run, the favored one, is described by Gladwin as follows

“Most navigators making this passage actually depart not under the star for Pikelot at

all, but rather under the setting of Cassiopeia, the star course for the southern tip of Gray Feather Bank. After a run of only thirty miles they are over its edge. Thereafter they can follow along over it or near it to the next reef, and so on all the way to Condor Reef. Even if they stray off course they know where they are: if they see reef under them on both sides they know they have veered too far to the north and should turn south, while if the water is deep and clear below they know they are too far to the south. No wonder men dare leave for Pikelot even when drunk!” (p. 163)

The passage elegantly shows how certain, stable, environmental features can compensate for uncertainty, in this case arising from factors such as changing current and wind. What is interesting about this example is that even though the environment apparently does not support the underlying knowledge structure very well, because of the lack of obvious landmarks, the Puluwatans still appear to use the same strategies for navigation as landlocked navigators. As Gladwin points out, the Puluwatans do not compensate for the relatively poor environment by using alternate representation strategies, rather they appear to have an increased reliance on pattern recognition. In essence, therefore, the Puluwatans retain a relatively sparse, simple, representation with the tradeoff that using it accurately requires more pattern recognition. Such a reliance has costs of course, the foremost being the extended periods of time that must be spent learning the relevant patterns. This problem is not explicitly addressed in R-PLAN since the environments used are information rich. Indeed indoor environments often have the opposite problem the Puluwat navigators faced since many hallways, for example, contain the same features repeated over and over. One reason why robots can perform in such environments is that pattern recognition is relatively less important and consequently a good internal model becomes more important.

### VISUAL STRUCTURE - GATEWAYS

A network of nodes is an economical method of storing information in a way that effectively manages uncertainty. However, while there is a degree of uncertainty in any environment, in many niches there is also a certain degree of stability. Object-based representation, while economical, is sparse with regard to the actual quantity of usable information. Useful information which is left out can normally be picked up through the

sensory interface, but such an arrangement can tie up processing and slow down execution. For example, if cognitive maps consisted purely of a network of landmarks, there would be no stored information indicating relative directions from one landmark to the next. A journey using just such a structure would consist of going to a landmark, looking for the next one, going to it, etc. The structure stores “what” the next landmark is, but not “where” it is. While the human visual system makes the “where” search process relatively fast, some sort of stored directional information would reduce processing and could work even in the absence of perceptual information, such as at night, or during a dense fog.

The human visual system is split into two distinct parts often called the “what” and “where” systems (Ungerleider & Mishkin, 1982; Otto, Gradguillaume, Boutkhal, et al., 1992). The “what” system primarily deals with object recognition and therefore is central to processing landmarks. The “where” system, on the other hand, codes information concerning, among other things, the locations of objects in space relative to the observer. The information processed by the two systems is distinct; where the “what” system is detailed, the “where” system uses coarser information such as texture and crude shape. In a sense the “what” system deals with pure objects regardless of their spatial locations, while the “where” system deals with pure space regardless of the particular objects involved. A complete representation of a scene would involve taking the objects identified by the “what” system and putting them into the locations identified by the “where” system. Given that this location information is constantly being processed, and that it can usefully augment “what” information, it is natural that it be incorporated into the knowledge structure.

The topological representation described so far could be thought of as invested almost exclusively in the “what” system; it only requires the ability to recognize landmarks. Even a small amount of location information, especially encoded in the same type of economical structure as the network of landmarks, adds considerable functionality to the knowledge representation (Kortenkamp & Chown, 1993).

### Gateways

For sighted humans the primary source of location information is through the “where” portion of the visual system. One of the critical questions in understanding cognitive maps is how “where” information, acquired from a particular viewpoint and egocentrically organized, is synthesized with other views of the same environment. One might imagine trying to construct a coher-

ent representation of an environment from a few snapshots. A central conjecture of the PLAN architecture is that there are structures within environments which can serve to naturally organize mental representations; structures of this type are called gateways in PLAN.

Gateways, like landmarks, function as a kind of environmental marker. However, where a landmark acts as a mere indicator of place, gateways serve in that capacity and have additional uses as well. The two defining characteristics of navigational gateways are that they occur at places that are often visited, and that new information becomes available at them (in cognitive mapping such information is typically visual). Fortunately, in most environments, real and artificial, these places are the same. The defining characteristics ensure that gateways mark important transitional locations, which gives rise to their names, taken from Christopher Alexander’s design construct of the same name (1977).

Gateways represent natural transitions between one region of space and another. Typical examples include doorways, mountain passes, and entrances to forests. In indoor environments gateways such as intersections, doorways and junctions are relatively easy to recognize, even for a robot relying primarily upon sonar, because of the strong change in perceptual input associated with an opening (Kortenkamp, et al, 1992). There are gateways in other domains as well. In computer software, for example, a menu is often a gateway for a new environment. These would be strong examples of gateways because they meet both of the defining criteria.

Gateways serve several essential roles in cognitive structure. First, they provide a natural way to divide an environment up into smaller, more manageable, pieces. Gateways can be thought of as occurring at the transition between regions. Eventually gateways can be used to organize these regions in a more abstract, hierarchical structure. Second, since they are easily identifiable and frequently visited, gateways serve as ideal locations to anchor visually based structures.

The anchoring problem has proven difficult for robotic models in the past. There are several robotics systems which use a form of scene-based storage which have addressed this issue with mixed results. For the most part the disadvantages are because they fail to account for the environments in which they will be used. Asada (1988), for example, created scenes at fixed intervals. While this method ensures uniformity and a certain amount of predictability it has numerous drawbacks. Among these disadvantage is the fact there is no guarantee that intersecting paths will have a common point where scenes are stored. Yeap improved on this by introducing the concept of regions, which divide environments up into smaller pieces (1988). Central to Yeap’s region construct was the element of an exit which pro-


vided considerable inspiration for the development of the gateway structure.

A non-spatial example might help to further clarify this concept. In music there are gateways which help to organize understanding. These include breaks between movements, reintroductions of themes, etc. These gateways help listeners organize their mental representations. Part of the function of this process is predictive; the reintroduction of a theme, for example, brings certain expectations about what might happen next. Gateways are strong markers which automatically trigger certain cognitive responses and organization of structure. Alternatively, imagine that representation of musical structure were organized by some other principle, such as breaking a piece up after a fixed number of notes. Such a structure in no way resonates to the natural construction of the music, providing little, if any, basis for expectations and predictions. By being sensitive to musical gateways and transitions, composers are able to exert a certain amount of control over the listening experience. Some modern music takes advantage of these expectations in a different way, using the *lack* of gateways to provide a different type of feeling.

### Local Maps

The gateway structure provides the foundation for the second type of network in PLAN. Instead of having object representations as nodes, the other type of network has spatial representations as nodes. The new type of node, called a “local map,” is a schematic representation of a scene. In essence a local map is a “snapshot” of what can be seen from a gateway. The snapshot metaphor is actually a poor one though, because the information stored in a local map is schematic. Essentially, a local map can be thought of as a schematic view of a location with the head held rigidly in place. Seeing all the parts of the scene in such an arrangement requires moving the eyes to focus upon different relative locations. This information can be encapsulated as a grid where each square of the grid represents an eye location with the contents of the square representing what can be seen with the eye in that position (Figure 1). By linking the grid to the representations of the corresponding landmarks, the relative location of an object can be quickly found. Thinking of the landmark will associatively activate the relevant grid location yielding the relative location. The grounding of the representation in the physical system goes even further in PLAN as multiple scenes can be stored at a single location corresponding to what might be seen with different head and body positions (Figure 2).


One theory of cognitive mapping postulates that the functioning of cognitive maps takes place almost ex-



**Figure 1. A scene-based representation. The shaded regions represent the presence of objects in the field of view. The details of the objects are not shown in order to reinforce the idea that this information is locational in nature. In PLAN or R-PLAN the shaded locations would be associatively linked to the representations of the corresponding objects.**

clusively within the hippocampus (O’Keefe and Nadel, 1978; O’Keefe, 1989). Most of the evidence concerns the relationship between firing rates within the hippocampus as modulated by different locations and orientations within a location. These data appear to correspond strongly with the directional structures in PLAN. Trullier et al. (1997) review a number of models inspired by research on the hippocampus. A number of these models postulate view-based representations which are accessed by a combination of sensory input and head and body positions. Like the work in robotics, they differ from PLAN in determining how and when such views are constructed. Like PLAN, many of the models postulate that an important issue is the discrepancy between the current view and the views stored in memory. PLAN goes a step further in specifying the types of differences which are important (eg. new information can be seen) and by postulating that organisms are sensitive to particular environmental configurations. In conjunction with more recent evidence presented by Squire (1992), this suggests that one role of the hippocampus is as an interface between stored spatial memory structures and the currently perceived environment. Such an interface is necessary for a representation which is built around distinct locations and views scattered throughout an environment.

Local maps provide the same types of tradeoffs between uncertainty and planning as are found in the landmark representations. Instead of linking landmarks to



**Figure 2. Local Maps.** At a given location the potential exists for several scene-based representations based upon different head locations. In this representation the orientation of the body is taken to be fixed. Different parts of the representation correspond to different head locations. The location of an object in the representation, therefore, yields the relative eye and head locations necessary to view it from the fixed body position.

other landmarks, in local maps “objects” in a scene are linked to particular locations. The rules for linkage are the same as they were for the network of landmarks, ensuring that any given local map contains only those “objects” which are often seen. Further, the actual arrangements of objects within a local map need not be exact. Once a general heading is determined, the sensory system can take over and provide the needed precision. In this way the local maps reduce the load on the sensory system but still rely on it to combat uncertainty. As a further hedge against uncertainty, local maps are linked together using the same rules as with landmarks.

### Gateways and environmental structure

While the network of landmarks relies upon relatively weak assumptions about the structure of environments, the network of local maps relies upon much stronger assumptions concerning the nature of gateways. It is certainly possible, given these assumptions, that there are environments without natural gateways. This is true, and it is also true that there are many environments in which humans do not function effectively. However, the gateway is not the only knowledge structure available for navigation; there is also the associative network of landmarks. The two structures are functionally complimentary in many ways. The landmark structure can be learned relatively quickly and provides a rough and ready functionality with only a minimum of experience. The gateway structure, on the other hand, while potentially slower to learn, provides significant additional in-

formation and abstracts much more readily. Because the local map is invested in stronger assumptions about environmental structure than the landmark structure, it is less likely to strongly apply to a given environment, but more likely to yield significant payoff when it does. The local map structure relies on some type of natural gateways which may or may not be present in a particular environment. While not all environments may have good natural gateways, the ones that do are more easily abstracted into higher level constructs such as regions; in turn these abstractions allow for simpler planning and reasoning over large spaces. Given the conjecture that the two representations are complimentary, it is reasonable to suppose that organisms might rely on one or the other structure in different niches, and that having the two types of structure available is adaptive, not only because systems which can function relatively independently are useful in case one or the other breaks down, but because it means that the organism can flourish in a larger range of environments. The examples provided by the interior hallway environment navigated by R-PLAN and the islands around Puluwat show why this might be the case. In the hallway environment there are many possible gateways but few landmarks. The representations of R-PLAN are much more invested in the gateway structure due to the limited object recognition capability of robots. Since a hallway provides few natural landmarks anyway this is not a major limitation. On the other hand, in the open sea there are very few, if any, natural gateways. As Gladwin pointed out the Puluwat navigators were forced to rely almost exclusively upon their pattern recognition capabilities in this environment.

Gateways are useful because they mark transitions from one environment to another (or to a subset of an environment). Cognitively this allows the world to be broken up into smaller pieces, each of which, presumably, is more manageable than the whole. Furthermore, gateways are likely to occur at particularly important spatial locations, affording the ability, for example, to escape to a safer environment. The importance of being able to leave an environment points out another important attribute of an effective gateway, that it is both an exit and an entrance; being thrust into a new environment with no clear way to return is discomfiting at best.

Alexander (1977), whose use of the term gateway was applied in the design domain, recognized its importance in cognitive structure.

“Many parts of a town have boundaries drawn around them. These boundaries are usually in people’s minds. They mark the end of one kind of activity, one kind of place, and the

beginning of another. In many cases, the activities themselves are made more sharp, more vivid, more alive, if the boundary which exists in people's minds is also present physically in the world." (p. 277)

The point that Alexander is making is that when cognitive structure matches environmental structure people will operate more effectively and comfortably. Environmental features which naturally tend to organize cognitive structure, as gateways do, are central to this process. Environments, natural or designed, which lack such features will be more difficult to organize cognitively and effectiveness will be reduced.

One example of an artificial environment which can be used to examine the gateway concept is a computer operating system. As computer operating systems have developed through the years they have taken on the structure of gateways to a larger and larger degree. Early command-line operating systems were almost featureless and the user had to rely almost exclusively upon expertise and memorization. One simple innovation to the command-line system was to change the prompt when a new program was run. This change in information was a simple visual signal that somehow the environment had changed from the operating system to that of the running program. As graphical interfaces have become more prominent the gateway notion has developed further; new programs even have their own area on the screen. Even the name of these areas, "windows," indicates that new information can be seen in them. By creating easily discernible breaks between operations with different functionality, software becomes simpler to manage cognitively. In Alexander's terms, creating a new window for a program creates a physical boundary for the boundary already in people's minds between the operating system and the running program. Further, Alexander has pointed out that gateways ought to be solid "things" which emphasize the idea of transition. In software this has been done largely through the use of menus. Rather than typing a command and getting an abrupt transition the user finds a menu and is prepared for the transition. Menus have the added advantage that they usually provide users with built-in knowledge of exits which imply the safety of being able to easily escape.

Computer operating systems have naturally coopted the landmark structure as well. The most obvious example of this is the representation of programs and files as icons. Users are not required to know arcane commands to run a program, they are only required to recognize an icon. Naturally these icons fulfill a kind of dual role since they also act as gateways; users know

that if they click on a word processing icon they will enter a different software environment.

It is not surprising that the terminology of software systems often reflects the spatial nature of its structure; internet software systems are referred to as "navigators," for example, with the links from one page to another providing a kind of gateway. Gateways show up in numerous other artificial structures; the chapter headings of books inform readers of a transition to a new topic, television shows have opening and closing credits, and even multiple scenes in movies tend to be linked by gateways through fadeouts, panning, etc. In every case these artificial gateways are designed to provide context for a conceptual shift in information. Conversely, the lack of a gateway can be jolting and disconcerting. In terms of navigating in the real world, gateways tend to be both entrances and exits, and an abrupt transition may mean that there is no obvious exit available.

### Gateways and hierarchy

Gateways also provide the basis for much of the hierarchy in PLAN and R-PLAN. Because they naturally break environments into smaller pieces and because they tend to be visited often, gateways lend themselves naturally to abstraction. Regions can be thought of as the space between gateways. This structure imposes a natural hierarchy upon environments similar to that developed by Yeap (1988), though Yeap's definition of a region is more narrowly defined and is tied very closely to enclosure by visual barriers. The representations in PLAN which can take advantage of this structure are called "regional maps." Regional maps are structured exactly the same as local maps but contain extended fields of view by using the predictive power of the network structure.

As an environment becomes familiar, prediction, in the form of activating the next local map, becomes more efficient. For example, "around this corner I would expect to see . . ." In time the predictive effects of association will begin to be taken into account in the current local map, particularly if one pauses, as one might at a gateway. If one mentally runs through the next part of a journey, a larger structure containing not only the current local map, but parts of the neighboring local maps can be created. Such a representation would contain information beyond that which can be seen by taking advantage of the predictive power of a network structure. In terms of the representation the operation involved is extending the boundaries of stored scenes. If something farther away is thought of as being further "up" or "out" in the visual field, then more distant objects will be placed further than previous objects on the periphery of the grid structure. Such placement represents a distortion of the true visual field at the

location in question. However, the distortion can be resolved by considering the regional map to occur at a new point directly above the original point, corresponding to an oblique view of the mapped region. Taken further this means that as the area covered by the Regional Map expands, the perceived “height” of the map will rise. Thus an oblique, or bird’s eye, viewpoint emerges as one becomes more and more familiar with a large scale space.

The importance of entry and exit gateways only increases in such a structure. Since the structure is scene-based it only contains information in the field of view. However, since entry and exit gateways occur at the edges of regions, a gateway corresponding to those locations can contain the entire region in the field of view.

### Cognitive maps and abstract environments

One striking aspect of implementing PLAN on a mobile robot was that it could be made to work at all given the fundamental differences in the capabilities, particularly sensory capabilities, of humans and robots. Because robots do not possess general object recognition abilities, the emphasis in R-PLAN was on the local and regional map representations (as already mentioned, simulations of the network of landmarks have been run and compared to human performance (O’Neill, 1990; 1991)). One reason that R-PLAN could succeed was that it navigated in indoor environments created by humans. By their nature such environments tend to be rich in the kind of structures detailed in this article, particularly gateways. In fact, not only was it relatively simple for the robot to identify gateways, but once the robot had marked a location as being a gateway it could repeatedly return to the same spot within 3.5 degrees of orientation and 70 millimeters of position on subsequent visits (Kortenkamp, et al, 1993). Even if a network of gateways were the robot’s only representation it would afford the robot the ability to go from virtually any location to any other in its environment by traveling from one gateway to the next. Such a mode of travel is quite similar to that reported by Pailhouse in his study of taxi drivers (1969). Pailhouse found that Parisian taxi drivers tended to divide the city up into regions with at least one critical point in each region. For long trips the taxi drivers would plan their trips as being from one critical point to the next until they got close to their destination.

R-PLAN represents one step along a longer path. In R-PLAN the cognitive map structure was tested by implementing it on an agent with vastly different perceptual capabilities than humans. In R-PLAN, for example, gateways can be determined by sonar, a perceptual modality that humans do not possess. One of the

issues this raises is how gateways might be recognized, or even designed, in environments that are not visual in character. For example in story analysis there are verbal gateways which announce the transition from one part of a story to another; a simple example is the word “meanwhile” which cues the listener or reader that there will be a context shift. Interestingly, “meanwhile” is often accompanied by a phrase such as “back at . . .” which serves to take the listener or reader somewhere they have already visited. Of course stories are full of other conventions that fit into the spatial framework; indeed stories are often described as being “journeys,” “trips,” etc.

The next step of this research is to more fully explore how the patterns found in the navigation domain extend to other environments and niches. For this to be the case the same rudimentary structures must appear in such environments. Among the critical factors are:

- 1) They must be populated with uniquely identifiable objects.
- 2) Information flow within the environment should be sequential.
- 3) There should be easily identifiable transitions within the environment which mark context shifts. Further, these transition points should be simple to reach, generally permitting the possibility of exit.

Returning to the example of music, it is clear that musical pieces often match these three points very well. Objects, in this case, are not only individual notes, but can also be entire melodies. Certainly music is a sequential medium, and there are gateways such as pauses, crescendos, diminuendos, etc. which mark transitions between different structures.

If all of these points hold true, then the potential environment should prove to be “friendly” in that it should be reasonably simple to parse and learn. There are other factors, of course, but these three things should strongly predict how comfortable humans will find a given environment.

One important direction to take this research is to explore how these concepts map onto other domains and to other organisms. In some cases this should be relatively simple. For example, artificial worlds can easily be constructed which vary along any of the points described above. Such worlds can be used to examine concepts such as whether or not gateway structures make understanding large-scale environments inherently easier. The same sorts of tests can also be run with humans, both on spatial environments and on abstract ones, to test if the structure of an environment can truly predict how easy it is to learn. One interesting point to research is whether the gateway structure is only useful in

situations where abstract reasoning is called for, or whether simple organisms in simple niches have no need for such constructs.

## CONCLUDING REMARKS

There is a continuum of uncertainty in real environments and even within a specific niche. Objects may appear and disappear from one experience to the next or they might be moved or changed in appearance. On the other end of the continuum some objects are for all practical purposes permanent. Furthermore, there are patterns which are potentially exploitable across environments, such as the gateway pattern. An adaptive organism must be sensitive to as much of this continuum as is possible if it is to survive in diverse niches. Just as it is not reasonable to assume that there is no uncertainty in an environment, it is also a mistake to assume that there are not environmental features which can be used for effective prediction. A significant aspect to adaptive behavior is the ability of an organism to quickly acclimate itself to new environments and niches. If such an organism does not have to start with a completely blank slate in a given niche it has an adaptive advantage. Cognitive maps are one example of how organisms are able to take advantage of the fact that the world is not random, that there are certain patterns and structures which occur over and over in many different environments. Such patterns are not restricted to purely spatial domains either, and so it is hardly surprising, given the parsimony of the brain, that such spatial patterns could be exploited in arbitrary environments.

## ACKNOWLEDGMENTS

The author would like to thank Dave Kortenkamp and Steve Kaplan for making this paper possible, both through helpful suggestions and insightful ideas.

## REFERENCES

- Alexander, C., Ishikawa, S., & Silverstein, M. (1977). *A pattern language*. New York: Oxford University Press.
- Asada, M., Yasuhito, F. & Tsuji, S. (1988). Representing a global map for a mobile robot with relational local maps from sensory data. *Proceedings of the 9th International Conference on Pattern Recognition*, 520-523.
- Bates E.A., & Elman, J.L. (1992). Connectionism and the study of change. (Tech. Rep. No. 9202). San Diego, California: The University of California at San Diego, The Center for Research in Language.
- Bonasso, R. P., Kortenkamp, D. & Miller, D.P. (1995). Experiences with an architecture for intelligent, reactive agents. *Proceedings of the 1995 IJCAI Workshop on Agent Theories, Architectures and Languages*.
- Brooks, R.A. (1985). Visual map making for a mobile robot. *Proceedings of the IEEE Conference on Robotics and Automation.*, 824-829.
- Brooks, R.A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2, 1:14-23.
- Brooks, R.A. (1991). Challenges for complete creature architectures, in *From Animals to Animats, Proceedings of the First International Conference on the Simulation of Adaptive Behavior*, Meyer, J.-A. and Wilson, S.W. (eds.), MIT-Press.
- Bruner, J. S. (1957). On perceptual readiness. *Psychological Review*, 64, 123-152.
- Brunswik, E. (1956). *Perception and the representative design of psychological experiments*. Berkeley: University of California Press.
- Byrne, R.W. (1979). Memory for urban geography. *Quarterly Journal of Experimental Psychology*, 31 (1), 147-154.
- Chapman, D. (1987). Planning for conjunctive goals. *Journal of Artificial Intelligence*, 32 (3), 333-377.
- Chown, E. (1994). Consolidation and Learning: A Connectionist Model of Human Credit Assignment. Doctoral dissertation. The University of Michigan.
- Chown, E., Kaplan, S. and Kortenkamp, D. (1995). Prototypes, Location and Associative Networks (PLAN): Towards a unified theory of cognitive mapping. *Cognitive Science.*, 19, 1-52.
- Cimprich, B. (1993). Development of an intervention to restore attention in cancer patients. *Cancer Nursing*, 16, 83-92.
- Clark, A. (1989). *Microcognition: Philosophy, cognitive science, and parallel distributed processing*. The MIT Press.
- Connell, J.H. (1992). SSS: A hybrid architecture applied to robot navigation. *Proceedings of the IEEE International Conference on Robotics and Automation*, 2719-2724.
- Craik, K.J.W. (1943). *The nature of exploration*. London: Cambridge University Press.
- Gibson, J. J. (1966). *The senses considered as perceptual systems*. Boston, MA: Houghton Mifflin.
- Gladwin, T. (1970). *East is a big bird: Navigation and logic on Puluwat Atoll*, Harvard University Press.
- Gopal, S., Klatzky, R. L. & Smith, T.R. (1989). NAVI-GATOR: A psychologically based model of environmental learning through navigation. *Journal of Environmental Psychology*, 9 (4), 309-31.
- Gopal, S. & Smith, T.R. (1990). Human wayfinding in an urban environment: A performance analysis of a computational process model. *Environment and Planning*, 22 (2), 169-91.

- Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley.
- Hebb, D. O. (1963). The semiautonomous process, its nature and nurture. *American Psychologist*, 18, 16-27.
- Hebb, D. O. (1980). *Essay on Mind*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Hebb, D.O. (1982). The mechanisms of perception. In (H. A. Buchtel, ed.) *The conceptual nervous system*. Pergamon Press.
- Hirtle, S. & Jonides, J. (1985). Evidence of hierarchies in cognitive maps. *Memory and Cognition*, 13 (2), 208-217.
- Holland, J.H. (1986). Escaping brittleness: The possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In: R.S. Michalski, J.G. Carbonell & T.M. Mitchell, (Eds.), *Machine Learning: An artificial intelligence approach* (Vol. 2). Los Altos, CA: Morgan Kaufmann.
- Horswill, I. (1993). Polly: A vision-based artificial agent. *Proceedings of the National Conference on Artificial Intelligence (AAAI-93)*, 824-829.
- Kaelbling, L.P., Littman, M.L. & Moore, A.W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4, 237-285.
- Kaplan, R. (1993). The role of nature in the context of the workplace. *Landscape and Urban Planning*, 26, 193-201.
- Kaplan, S. (1973). Cognitive maps in perception and thought. In R. M. Downs and D. Stea (eds.), *Image and Environment*. Chicago: Aldine.
- Kaplan, S. & Kaplan, R. (1988). *The Experience of Nature*, Ann Arbor: Ulrichs.
- Kaplan, S. & Kaplan, R. (1989). The visual environment: Public participation in design and planning. *Journal of Social Issues*, 45, 59-86.
- Kaplan, S. & Peterson, C. (1993). Health and environment: A psychological analysis. *Landscape and Urban Planning*, 26, 17-23.
- Kaplan, S., Sonntag, M. & Chown, E. (1991) Tracing recurrent activity in cognitive elements (TRACE): A model of temporal dynamics in a cell assembly. *Connection Science*, 3, 179-206.
- Kortenkamp, D. (1993). Cognitive maps for mobile robots: A representation for mapping and navigation. Doctoral dissertation, The University of Michigan.
- Kortenkamp, D., Baker, L. D., & Weymouth, T. (1992). Using gateways to build a route map. *IEEE/RSJ International Conference on Intelligence Robots and Systems*.
- Kortenkamp, D. & Chown, E. (1993). A directional spreading activation network for mobile robot navigation. *From Animals to Animats 2, Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, Meyer, J.-A., Roitblat H. L. and Wilson, S.W. (eds.), MIT-Press.
- Kortenkamp, D., Chown, E., & Kaplan, S. (1996). Implementing a psychological model of cognitive mapping on a mobile robot: From theory to practice. Submitted.
- Kuipers, B. J. (1978). Modeling spatial knowledge, *Cognitive Science*, 2, 129-153.
- Lakoff, G. (1987). *Women, Fire, and Dangerous Things*, Chicago, IL: The University of Chicago Press.
- Levenick, J. R. (1991). NAPS: A connectionist implementation of cognitive maps, *Connection Science* 3, 107-126.
- Macphail, E. (1987). The comparative psychology of intelligence. *Behavioral and Brain Sciences*, 10(4), 645-695.
- Mahadevan, S. & Connell, J. (1992). Automatic programming of behavior-based robots using reinforcement learning. *Artificial Intelligence*, 55, 311-365.
- O'Keefe, J. (1989). Computations that the hippocampus might perform. In Nadel, L., Cooper, L. A., Culicover, P. and Harnish, R. M. (eds.), *Neural connections, mental computation*. Cambridge, MA: MIT Press.
- O'Keefe, J. & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.
- O'Neill, M. J. (1990). Computer simulation of the cognitive map: A validation study. In R. I. Selby, K. H. Anthony, J. Choi, and B. Orland (eds.) *Coming of Age: Proceedings of the Twenty-first Annual Conference of the Environmental Design Research Association*, Oklahoma City: EDRA.
- O'Neill, M. J. (1991). A biologically based model of spatial cognition and wayfinding. *Journal of Environmental Psychology*, 11, 299-320.
- Otto, I., Grandguillaume, P., Boutkhil, L., Burnod, Y. & Guigon, E. (1992). Direct and indirect cooperation between temporal and parietal networks for invariant visual recognition. *Journal of Cognitive Neuroscience*, 4, 35-57.
- Piaget, J., & Inhelder, B. (1967). *The Child's Conception of Space*, W.W. Norton and Company, New York.
- Shemyakin, F.N. (1962). General problems of orientation in space and space representations. In B.G. Anayev, editor, *Psychological Science in the USSR Vol. 1*, U.S. Joint Publications Research Service, Washington D.C., 186-251.
- Siegel, A. W., & White, S. H. (1975). The development of spatial representation of large-scale environments. In Reese, H. W. (ed.), *Advances in Child Development and Behavior*, Volume 10, Academic Press, New York, 1975.
- Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11, 1-74.



- Sonntag, M.L. (1991). Learning sequence in an associative network: A step towards cognitive structure. Doctoral dissertation. The University of Michigan.
- Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys and humans. *Psychological Review*, 99(2), 195-231.
- Tennessen, C.M. & Cimprich, B. (1995). Views to nature: Effects on attention. *Journal of Environmental Psychology*, 15, 77-85.
- Trullier, O., Wiener, S.I., Berthoz, A., & Meyer, J.A. (1997). Biologically based artificial navigation systems: Review and prospects. *Progress in Neurobiology*, 51, 483-544.
- Ungerleider, I.G. & Mishkin, M. (1982). Two cortical visual systems. In D.J. Ingle, M.A. Goodale, & R.J.W. Mansfield (Eds.), *Analysis of visual behavior*. Cambridge, MA: MIT Press.
- Weisman, G. D. (1981). Way-finding and the built environment: an evaluation of architectural legibility. *Environment and Behavior*, 13, 189ff.
- Yeap., W.K. (1988). Towards a computational theory of cognitive maps, *Artificial Intelligence*, 34, 297-360.

## ABOUT THE AUTHOR



### Eric Chown

Eric Chown is an Assistant Professor at Bowdoin College in Brunswick Maine. He received his Ph.D. from the University of Michigan in 1994 where he worked on a cognitive architecture called SESAME. From 1994 to 1998 he was a postdoctoral fellow at Oregon State University where he developed machine learning techniques for use in the construction and update of ecosystem models – particularly models used in global warming research. His current research focuses on architectures for modeling learning, as well as human cognitive mapping.