The "Map in the Head" Metaphor[*]

Benjamin Kuipers

Environment and Behavior **14**(2): 202–220, 1982.

Abstract

The "Map in the Head" metaphor states that knowledge of large-scale space is isomorphic to the information stored in a graphical map: That is, corresponding operations are used to store and retrieve information. The purpose of this essay is to look carefully at the "Map in the Head" metaphor to see the limits of its applicability. There are two types of experimental results that are difficult to accommodate within this metaphor. First, instead of being integrated into a single map, spatial knowledge can fall into disconnected components, with little or no relation between the components. Second, knowledge of routes (and other spatial facts) may be represented asymmetrically, so that a route can be followed in one direction but not in the other. The first set of results leads us to replace the simple "Map in the Head" with a more complex and sophisticated metaphor including separate metrical and topological components. The second set of results suggests that even the more sophisticated "Map in the Head" is built from computational structures that occasionally reveal their nonmaplike properties. A computational model is presented for assimilating observations gathered during travel, first into a description of the particular route, then into representations for the topological and metrical features of the environment.

THE COGNITIVE MAP

The study of environmental cognition focuses on the structure and content of the "cognitive map." While it is generally agreed that the cognitive map

---

[*] The research described here was performed in part while the author was at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Project Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643.

is not entirely maplike, the metaphor of the "Map in the Head" is the most persuasive and useful image for knowledge of large-scale space. However, metaphors and images must be treated very carefully in scientific investigations, lest their accidental properties be confused with the real properties of the phenomenon being studied.

The purpose of this essay is to look carefully at the "Map in the Head" metaphor to see the limits of its applicability. There are two types of experimental results that are difficult to accommodate within this metaphor. First, instead of being integrated into a single map, spatial knowledge can fall into disconnected components, with little or no relation between the components (Appleyard, 1970; Kosslyn et al., 1974; Lynch, 1960). Second, knowledge of routes (and other spatial facts) may be represented asymmetrically, so that a route can be followed in one direction but not in the other (Hazen et al., 1978; Piaget et al., 1960). The first set of results leads us to replace the simple "Map in the Head" with a more complex and sophisticated metaphor including separate metrical and topological components. The second set of results suggests that even the more sophisticated "Map in the Head" is built from computational structures that occasionally reveal their nonmaplike properties.

Rather than use the metaphorically loaded term *cognitive map*, we will define the cognitive spatial description to be that body of knowledge of a large-scale environment that is acquired by integrating observations gathered over time and is used to find routes and determine the relative positions of places. A *large-scale* environment is one whose structure is revealed by integrating local observations over time, rather than being perceived from one vantage point.

Notice that a graphical map of a city uses the (quite different) cognitive mechanisms of visual space to convey information about the large-scale space. Within a given visual image, a large number of "landmarks" are simultaneously visible, so relative distances and directions are easy to judge. Exploration methods that involve repeated visual searching and scanning are perfectly feasible. In large-scale space, by contrast, exploration is constrained by the speed of physical travel, so repeated searching and scanning are very expensive in time and effort. Relative distances and directions are difficult to determine from the perspective of the traveler. Furthermore, over the longer time intervals involved in travel, the explorer's cognitive processes are likely to be disrupted by more urgent demands. Because of the great differences in sensory access and processing demands, large-scale space can be considered a "different kind of space" from the visual space used to read a graphical map. The use of graphical maps to convey useful information about a large-scale space is therefore a complex process whose detailed structure remains beyond our grasp until we understand the structure of knowledge in each type of space. This article addresses the similarities and differences between the two types of spatial knowledge and their internal representations.

The "Map in the Head" metaphor states that spatial knowledge, as stored in the cognitive map, is isomorphic to the information stored in a graphical map. That is, information is added to, and retrieved from, the cognitive map using the same operations by which information is added to or retrieved from a graphical map. This is not to say that there must be a region of the brain onto which the environment is mapped, preserving two-dimensional spatial relations. Rather, there should be an exact correspondence between the input-output behaviors of the storage and retrieval functions appropriate to the two representations. Shepard and Chipman (1970) term this a "second-order isomorphism."

The simplest and strongest version of this metaphor says that the "Map in the Head," inspected by the "mind's eye," is functionally identical to a graphical map inspected by the physical eye. While there are many properties of interest to both types of eye, we will concentrate first on operations on vectors (i.e., distances and directions).

1. Given two locations, determine the relative-position vector pointing from one to the other.
2. Given two relative-position vectors, determine their sum.
3. Given two relative-position vectors, determine their difference in direction.

These operations, among others, clearly take place both during mental inspection of the cognitive spatial description and during visual inspection of a graphical map. The "Map in the Head" metaphor states that the functional behavior is the same in the two contexts.

This proposal provides a set of assumptions to support the study of spatial distortions: Since there is geometry inside the head and geometry outside, one can study the correspondence. The geographers and cartographers interested in environmental cognition have laid out these assumptions most clearly:

> If we can identify corresponding entities in the two images, we can construct a table of correspondences. One set represents coordinates of the associated point in one image, and the

> other set represents the coordinates of the associated point in the other image. We now assume that there exists a mapping between the sets, and we wish to investigate the properties of this transformation.... To obtain the distorted grid, we must use an interpolation procedure to go from the isolated observations to a field of data, in effect invoking two assumptions. The first assumption is that the functions are effectively, at least piecewise, continuous and everywhere defined. It is also convenient to assume that they are one-to-one and single valued. The conventional wisdom seems to admit to these types of assumptions [Tobler, 1975: 74].

Beck and Wood (1976) have studied sketch maps using these assumptions. After imposing a rectangular grid on a collected sketch map, they draw curves through the corresponding points on a cartographic map to demonstrate the nature of the distortion. The intended interpretation is that the distortion between the two maps reflects the distortion between the subject's cognitive spatial description and the actual geography. This interpretation relies heavily on the "Map in the Head" metaphor: The subject's cognitive spatial description possesses a two-dimensional structure (at least functionally) that can be placed in correspondence with his sketch map. In case the cognitive spatial description is more complex than that, the distorted grid interpretations are highly suspect.

I argue below that the cognitive spatial description consists of several very distinct representations for spatial knowledge: procedures for following routes, topological network descriptions, and metrical descriptions consisting of relative-position vectors. All can result in sketch maps. By applying default assumptions such as *"Draw a street as straight unless you know otherwise,"* a subject can put information on a sketch map that is not in his cognitive spatial description. Any mark on a sketch map acquires a position relative to all the other marks, but that relation may not reflect anything at all in the cognitive spatial description. A grossly distorted sketch map may result from a single incorrect relationship between two regions, each internally correct, rather than arising from widespread errors in relative position information. Improved methods for interpreting sketch maps must rely on a sufficiently rich model of the underlying knowledge representation and must deal explicitly with the processes that translate internal descriptions into marks on paper.

Experimental psychologists have also focused their attention almost uniformly on the "maplike" properties of the cognitive spatial description. Early work explicitly proposed that spatial knowledge consisted of an "imaginary map" (Trowbridge, 1913) or "field map" (Tolman, 1948) of the environment. Since the ethnographic work of Lynch (1960) in urban planning, and the developmental research of Piaget et al. (1960), many authors have recognized a distinction within the cognitive spatial description between "route maps" and "survey maps" as different types of spatial knowledge. However, virtually all experimental work since then (e.g., Kosslyn et al., 1974; Hardwick et al., 1976; Kozlowski and Bryant, 1977; Gärling, 1980) has dealt exclusively with the acquisition and use of knowledge of relative position. The "Map in the Head" metaphor has made nonmaplike aspects of spatial knowledge more difficult to define and to study.

Kosslyn et al. (1974) note the effect of travel distance on children's judged distances between places and suggest that children's cognitive spatial descriptions consist of "minimaps" less well-coordinated than those of adults. They leave open the question of how this coordination takes place. Hazen et al. (1978) report one of the few pieces of recent psychological research concentrating on knowledge of routes rather than of relative position. The assumption underlying the general concentration on knowledge of relative position seems to be that the ability to navigate from place to place is built on a foundation of knowledge of where things are. In other words, the irreducible core of the cognitive spatial description is the "Map in the Head." This essay argues that the opposite is true: Knowledge of where things are is built on a foundation of knowledge of how to navigate from one place to another.

In exploring the "Map in the Head" metaphor, we begin by focusing on the strong version stated above: A person's cognitive spatial description possesses a two-dimensional structure (at least functionally) that corresponds to the structure of the external world. Although this view is initially plausible, there are many phenomena it cannot account for, suggesting that there is some quite different underlying representation that simply appears maplike under some circumstances. In particular, two properties of vector operations on graphical maps are quite suspect in the context of the cognitive spatial description:

1. The map forms a single, global frame of reference within which any two vectors can be compared or added.
2. For any stored relationship (e.g., the vector from A to B) there is a corresponding reverse relationship represented in the map (the vector from B to A).

I claim that the comparability of all vectors is a property too strong for the human cognitive spatial description, which therefore cannot be a "Map in the Head" of this simple type. Many researchers have observed that cognitive spatial descriptions and sketch maps are frequently disconnected, with regions whose internal structure is well defined and reasonably accurate, but whose relation with the rest of the geography is simply unknown (Appleyard, 1970; Beck and Wood, 1976; Lynch, 1960). The reader should be able to perform a convincing thought experiment by selecting a pair of places whose relative orientation is not represented at all in his cognitive spatial description. (This is particularly easy in the Boston area.)

The simple "Map in the Head" metaphor would depict the cognitive spatial description as a structure in which every point must have a spatial relation to every other, just as two points cannot be drawn on a graphical map without an implied spatial relationship. However, the evidence indicates that the structure of the cognitive spatial description is such that two locations (or even two regions) can be represented and recognizable, but be completely without a defined spatial relationship. Thus the actual representations in the cognitive spatial description must be able to express more states of partial knowledge than is possible for the "Map in the Head."

When the "Map in the Head" is used as a metaphor of limited applicability, it serves the useful function of capturing some of the behavior of the cognitive spatial description. However, when it is treated as a theory of the underlying representation of spatial knowledge, or worse when it is tacitly assumed in the analysis of experimental data, its implications must be examined carefully. In this context, it is worth explicitly drawing the conclusion of global comparability of relative-position vectors, a conclusion that clearly stretches the metaphor beyond its proper domain.

TOPOLOGICAL RELATIONS IN THE "MAP IN THE HEAD"

A more sophisticated view of spatial knowledge can avoid many of the limitations of the simplest "Map in the Head" metaphor. This view arises from the recognition that topological relations are treated quite differently from metrical relations in the cognitive spatial description.

> However distorted, there was a strong element of topological invariance with respect to reality. It was as if the map were drawn on an infinitely flexible rubber sheet; directions were twisted, distances stretched or compressed, large forms so changed from their accurate scale projection as to be at first unrecogniz-

> able. But the sequence was usually correct, the map was rarely torn and sewn back together in another order. This continuity is necessary if the image is to be of any value [Lynch, 1960: 87].

The strict "Map in the Head" theorist (e.g., Tobler, 1975) interprets this observation to mean that the mapping between the inner and the outer geometry leaves topological relations invariant. An alternate interpretation is that topological relations such as connectivity, order, and containment are represented (or at least retrieved and manipulated) separately from metrical relations of distance and direction. In particular, this view holds that it should be possible to store a topological relation between two places in the absence of any metrical relation between them. This interpretation gains support from developmental studies (Hardwick et al.,1976; Hart and Moore,1973; Hazel et al.,1978; Kosslyn et al.,1974; Piaget et al.,1960; Siegel and White, 1975) and from studies of individual differences (Lynch, 1960), which encounter individuals or stages in which most spatial information consists of topological relations.

The metaphor of the "Map in the Head" can be extended to accommodate topological relations by adding storage and retrieval functions to create or trace connections between the nodes and edges of a street network. The cognitive spatial description could then answer questions about order of places on a street or containment within a region by scanning with the "mind's eye" along the street or boundary. One theory might explain a complete separation between topological and metrical relations by hypothesizing that storage and retrieval functions for distance and direction are developed by the child years after storage and retrieval functions for connection, order, and containment. Any new point added to the map would have implicit metrical relations with every other, but the child would not be equipped to detect them. In this version of the theory, it would still be necessary to explain why the adult does not eventually gain access to all possible relative-position vectors.

A different extension of the "Map in the Head" metaphor could avoid completely the problem of global comparability of vectors. Suppose that the "Map in the Head" were not a map but an atlas, with each sheet defining a separate frame of reference within which any two vectors could be compared. Relative-position vectors taken from different sheets could be compared only with effort, if at all, even by an adult with the normal capacity to store and retrieve vectors on the same sheet. Points on the different sheets might still be topologically related, connected by pieces of string, as it were, so that the storage and retrieval functions for connection and order could ignore the separation of the map into sheets. The atlas might even contain sheets drawn at different scales, so that the locations of points on one sheet can represent metrical relations between other entire sheets (cf. McDermott, 1980). This view of the cognitive spatial description predicts two different kinds of distortions on sketch maps: those reflecting distortions within a particular sheet and those reflecting a distorted relationship between sheets.

This is a far more sophisticated metaphor than the simple "Map in the Head" discussed above, and one that corresponds reasonably well with current theory and observations of the cognitive spatial description. How well does it stand up if we attempt to promote it from a metaphor to a theory of the representation of spatial knowledge? Unfortunately, not well.

When we criticized the strong analogy between cognitive spatial description and graphical map, the second suspect property of graphical maps was the fact that each spatial relation has a corresponding reverse represented in the map. If A is north of B, then B is south of A; if A is connected to B, then B is connected to A; if A precedes B on street S then B follows A on street S; and so on. While this is usually true for human cognitive spatial descriptions as well, exceptions are not at all rare. Developmental theorists, beginning with Piaget et al. (1960), have found that knowledge of routes is frequently not reversible in young children. That is, a child may know a route well enough to follow it in one direction, but not in the other. Lynch observed the same phenomenon as an individual variant in adults. Lynch's observations show us how this intermediate stage of knowledge can be possible:

> A sequential series of landmarks, in which one detail calls up anticipation of the next and key details trigger specific moves of the observer, appeared to be a standard way in which these people traveled through the city. In such sequences, there were trigger cues whenever turning decisions must be made and reassuring cues that confirmed the observer in decisions gone by [1960: 83].

Such a series of triggering landmarks will often not support travel on the route in the reverse direction. These connections are one-way links: A landmark may evoke the instruction for going from A to B, but learning the instruction leading from B to A is entirely independent and often relies on a different triggering landmark. The reader may be able to recall similar effects, particularly on infrequently traveled or newly learned routes.

Since this effect can be observed in route-learning experiments where the landmarks are equally visible from all directions (Hazen et al., 1978), it cannot be explained by the "accidental" circumstance that key landmarks might be obscured when traveling in the reverse direction. Rather, the irreversibility of routes reflects fundamental properties of the way humans access stored information. Even such an apparently symmetrical relationship as the distance between A and B is not always stored symmetrically. Lee (1970) reports that the distance from a metropolitan center to outlying locations is estimated as greater than the reverse distance.

The more sophisticated "Map in the Head" metaphor captures a large portion of the actual behavior of the cognitive spatial description and is therefore useful in organizing our observations and in making many successful predictions. However, it is inadequate as a theory of the representation of spatial knowledge because it makes a clearly false prediction about an important feature of the cognitive map: If a spatial relationship is stored in the cognitive spatial description, its reverse is also stored. Although the cognitive spatial description may frequently behave like a complex image scanned by the "mind's eye," we have shown that this behavior must be accomplished by some deeper level of representation that reconstructs the spatial relations from some other description. That deeper level occasionally reveals itself through behavior such as asymmetrically stored spatial relations.

COMPUTATIONAL MODELS AND THE ASSOCIATIVE LINK

We have seen that several quite different kinds of behavior come out of the cognitive spatial description. We would like to find a common basis for expressing the mechanisms that produce these behaviors and the relations between them. Computational models provide the most productive view of cognitive processing available in psychology today. These models are based on the assumption that there is a strong analogy between cognitive processes and computational processes: The manipulation of knowledge during learning and problem solving can be modeled by corresponding manipulations of data structured by computational procedures. The part of computer science known as *artificial intelligence* is dedicated to the study of intelligent processes, whether found in humans or created in machines. Research in artificial intelligence combines, in various proportions, constraints derived from human performance at a particular task, and the design of computational methods that accomplish tasks requiring intelligence.

An information-processing model describes the cognitive map as a complex data structure and a set of associated procedures[1] that exhibit a certain input-output behavior. The cognitive spatial description receives a sequence of observations, represented as structures produced by the perceptual system. The procedures in the cognitive spatial description assimilate those observations by changing or extending its data structures to include the information provided by the observations. Since the cognitive spatial description includes multiple representations for different types of spatial knowledge, there must be different types of data structures implementing the different representations. This focuses attention on a more specific question: Given what we know about the behavior of the cognitive spatial description in humans, what are the data structures by which it represents spatial knowledge, and what are the procedures by which it assimilates observations and solves problems?

A fundamental part of most models is the associative link between two pieces of data: If, given A, it is possible to retrieve B, then we say that A is linked to B (A → B). There are many functionally equivalent ways of implementing this link on a computer, but the most straightforward is to index the pair (A B) under the key A in a database. Notice that this link is inherently asymmetrical. If we wish to represent the reverse link (B → A), the pair (B A) must be indexed under the key B.

Using these associative links,[2] we can define an idealized model of the process by which a traveler learns a new route from observations. For the purposes of this argument, let us begin by assuming that the traveler's cognitive spatial description is empty, not only of particular spatial facts but even of such concepts as *place* and *path.* We will assume that the traveler's observations consist of a sequence of sensory and motor descriptions: *views* and *actions.* Views might be either stored images or complex prepositional descriptions. They need not be visual images: A blind person's "views" could be auditory, tactile, or even olfactory. We are concerned not with their internal structure, but with their role in the cognitive spatial description: to be matched for identity or used as keys for retrieval from a database. Actions are of two types: a *rotate* records the angle of rotation, and a *travel* records the distance traveled. Lynch (1960) observes that the cognitive spatial description contains relatively little detail between decision points. Since people seem to make little use of the fact that views change continuously during actions, our assumption of a discontinuous sensorimotor world appears justified.

As the attentive traveler is conducted along a route, he creates a link between each perceptual view and the action taken there to continue along the route (view → action). Holding both of these in mind until the next decision point is reached, he can also create the link to the result of the action ([view, action] → new-view). It is this second link that supports anticipation of upcoming landmarks, and thus mental review of the route in the absence of the environment. A complete set of these associations constitutes a complete description of the route, so that it can be followed independently or reviewed mentally from beginning to end. It is also asymmetrical, since the stored links support retrieval from beginning to end, but not from the end to the beginning. The *route description*, then, is effectively a procedure specifying a sequence of operations in terms of a set of sensorimotor primitives and associations between them.

The attentive traveler, concentrating on the construction of a complete route description, is the exception rather than the rule. In general, commonsense knowledge such as the cognitive spatial description functions in the background of many other processes that freely interrupt and disrupt its successful operation (Kuipers, 1979a). Under these circumstances, the route description will frequently consist only of associations of the first kind (view → action). Such a route description supports self-guided travel through the physical environment, but not mental review of the route or anticipation of upcoming landmarks, since the associative link can only be retrieved when the sensory system provides the current view to act as a retrieval key. In more impoverished situations, only some of the associative links are stored successfully, resulting in *states of partial knowledge* that support some, but not all, of what we call "knowing a route."

This type of performance is called *graceful degradation of performance under resource limitations* and is an important characteristic of commonsense knowledge. A knowledge representation can support graceful degradation by providing many states of partial knowledge, so that under resource limitations, a partially constructed or partially retrieved description can still be (1) meaningful and (2) useful. The same structural feature can support improvement of performance during learning. That is, under resource limitations, the assimilation of observations can take place in small increments, each resulting in a state of knowledge that is (1) meaningful and (2) incrementally more useful than its predecessor. A given knowledge representation can be analyzed to determine which states of partial knowledge it supports, what resources are required to change from one to another, and hence how well it supports graceful degradation or graceful improvement (Kuipers, 1979a, 1979b). Since there is great survival value in adequate performance under resource limitations, we can derive useful insights into the structure of human knowledge representations by considering these properties.

This representation for knowledge of routes demonstrates how a computational model can exhibit a much wider range of states of partial knowledge than would be possible for the "Map in the Head." Our attention was first drawn to this issue by the question of storing a spatial relationship without its reverse, but we have accomplished much more. The route representation is clearly irreversible, but can also be "patchy," with some segments of the route stored while others are missing, and can allow travel in the environment without the ability to anticipate upcoming landmarks. These phenomena (e.g., "I could take you there, but I can't tell you how.") cannot be explained by any other model of the cognitive map. It remains to show how representations for topological and metrical relations can be added to this computational model to account for the more "maplike" properties of the cognitive spatial description.

SIMULATING THE "MAP IN THE HEAD"

How can we extend the computational model of spatial knowledge to accommodate the more "maplike" aspects of the cognitive spatial description? The first step is to show how representations for fixed features of the environment such as places and paths can be derived from sensory experiences. The route representation given above has very few different *types* of knowledge: views, actions, and their sequence in time. Of course, in a rich environment, there will be many distinct instances of each type. As particular routes are learned, they constitute a new *type* of description created by aggregating views, actions, and the associations between them encountered while attempting to reach a particular goal.

This process of aggregating descriptions and associations and giving them a collective name is particularly important. The aggregate is a description of a new type, with properties that are distinct from any of its parts. For example, the attributes of a route description can include a total length, a number of actions, or a confusing segment, none of which could be meaningfully applied to the views, actions, or individual associations that are its parts. By giving the aggregate a name, it can be referred to in other descriptions, and it becomes a new unit in the memory structure of the cognitive spatial description.

A route description is formed by aggregating views and actions observed while pursuing a particular goal. A *place* description can be formed by aggregating views that are linked by actions that involve only rotation. If the perceptual universe consists of a large number of distinct views, this amounts to breaking that set of views into *equivalence classes*, taking two views to be equivalent if a rotation links them (i.e., V1 $\sim$ V2 if [V1, rotate] → V2). Each equivalence class is defined to be a place. An analogous process defines *paths* by taking two views to be equivalent when linked by an action involving only travel. The importance of these definitions is that the place and path aggregates can be given attributes and can have relationships that were meaningless for views, actions, or routes. In particular, topological relations such as connection can be defined between places and paths, and order relations can be defined among the places on each path. As new *types* of descriptions, the concepts of place and path extend the vocabulary for representing spatial information in the cognitive spatial description, allowing it to describe fixed features of the environment, rather than only the experiences of the traveler over time. It seems likely, of course, that the *concepts* of place and path are not newly invented by each child, but have been discovered during the evolution of the species and are passed on to each individual as part of its innate equipment.

The traveler through an environment receives a sequence of observations from which he must incrementally construct his cognitive spatial description. Therefore, the places and paths we have abstractly defined as equivalence classes must be gradually constructed by linking together views obtained under appropriate circumstances.[3] Once a view is linked to its place and path descriptions, we can recognize that a place and a path are topologically connected when they share a particular view. That connection is represented by an associative link (or *pointer*) from the place to the path and a corresponding *back-pointer* from the path to the place, together making a symmetrical connection. This symmetrical connection is the first step toward making the computational model behave like the "Map in the Head."

The order of places on a path is recognized and assimilated in a similar way: Individual views are linked to their place descriptions, so if two views are observed in sequence during travel, one piece of data about the order of their corresponding places can be added to the path description. Thus, a network of places and paths, linked by topological relations of connection and order, can be created from the sequence of views and actions. Problem-solving procedures, using well-known techniques for searching networks (Nilsson, 1980), can use this network to solve route-finding problems. This representation for topological relations accomplishes in the computational model the behavior specified by the topological parts of the "Map in the Head."

This line of reasoning can be pursued one step further in the same direction to account for relative position vectors by simulating the functional behavior of the two-dimensional "Map in the Head." The information needed to derive distance and direction information is available from the action descriptions: amount of rotation and distance of travel. Interpreting these as relative position information requires us to introduce the concept of a frame of reference. There are at least two distinct computational methods for representing relative position vectors and manipulating them to solve problems. From their study of visual recognition in three-dimensional space, Marr and Nishihara (1978) describe an explicit representation for vectors in terms of length and direction within an explicitly specified coordinate frame. As part of the cognitive spatial description, this proposal would require a description for each relative-position vector, whose parts would include a distance, a direction, and the frame of reference (cf. McDermott, 1980).

From a separately motivated study of mental imagery, Kosslyn and Shwartz (1977) propose representing relative positions by placing points on an internal two-dimensional array with a restricted scanning process. The frame of reference is implicit in the array. This amounts to implementing the simplest "Map in the Head" as a computational process serving as one component of a more complex cognitive spatial description, with a well-defined interface to the other components. Either proposal supports the definition, addition, and comparison of relative position vectors. To distinguish between the two proposals requires either more sensitive empirical data or a more careful consideration of existing data. It may well be that the cognitive spatial description uses both mechanisms for relative positive vectors under different circumstances, perhaps one for long-term storage and the other for spatial inferences in working memory.

Rather than describing the cognitive spatial description as a single representation for spatial knowledge, we have argued that it consists of a number of distinct representations (metrical, topological, procedural, and sensorimotor), all of which are implemented as components of a computational model. The original "Map in the Head" metaphor, which does capture some intuitions about spatial knowledge, reappears in the functional specifications for the metrical component of the cognitive spatial description. The power of the computational model is that it provides a common basis within which the different components can interact, in particular to allow assimilation of observational information from the route representation to the topological and relative position descriptions. The "Map in the Head" metaphor fails as a theory primarily because it cannot provide a common basis for the different types of information that the cognitive spatial description clearly includes and because it is not as rich in states of partial knowledge as the cognitive spatial description must be. It remains useful as a metaphor because it is an intuitively accessible model for an important part of the behavior of human spatial knowledge.

NOTES

1. The definition of a data structure must specify the operations that can be performed on the data. For example, the integers allowing addition and multiplication are quite different from the integers restricted to ordinal comparison, even though their storage formats in the computer might be identical.

2. Associations between concepts are nothing new to psychology, having been known at least since Aristotle. However, computational models provide a particularly rich way of explaining associations. and building complex structures from them to represent other types of knowledge.

3. These representations and the algorithms for assimilation are described in considerably more detail by Kuipers (1978, 1979b) as part of the TOUR model of the cognitive spatial description.

REFERENCES

APPLEYARD, D.(1970) "Styles and methods of structuring a city." *Environment and Behavior* 2: 100-117.

BECK, R. and D. WOOD (1976) "Comparative developmental analysis of individual and aggregated cognitive maps of London," in G. T. Moore and R. G. Golledge (eds.) *Environmental Knowing: Theories, Research, and Methods.* Stroudsburg, PA: Dowden, Hutchinson and Ross.

GÄRLING, T. (1980) Environmental Orientation During Locomotion: Experimental Studies of Human Processing of Information About the Spatial Layout of the Environment (Document D24:1980). Stockholm: Swedish Council for Building Research.

HARDWICK, D. A., C. W. MclNTYRE, and H. L. PICK (1976) "The content and manipulation of cognitive maps in children and adults." *Monographs of the Society for Research in Child Development* 41(3, Serial No. 166).

HART, R. A. and G. T. MOORE (1973) "The development of spatial cognition: a review," in R. M. Downs and D. Stea (eds.) *Image and Environment.* Chicago: Alpine.

HAZEN, N. L., J. J. LOCKMAN, and H. L. PICK, Jr. (1978) "The development of children's representations of large-scale environments." *Child Development* 49: 623-636.

KOSSLYN, S. M. and S. P. SHWARTZ (1977) "A simulation of visual imagery" *Cognitive Science* 1: 265-295.

KOSSLYN, S. M., H. L. PICK, Jr. and G. R. FARIELLO (1974) "Cognitive maps in children and men." *Child Development* 45: 707-716.

KOZLOWSKI, L. T. and K. J. BRYANT (1977) "Sense of direction, spatial orientation, and cognitive maps." *J. of Experimental Psychology: Human Perception and Performance* 3: 590-598.

KUIPERS, B. J. (1979a) "On representing commonsense knowledge," in N. V. Findler led.) *Associative Networks: The Representation and Use of Knowledge by Computers.* New York: Academic Press.

KUIPERS, B. J. (1979b) "Commonsense knowledge of space: learning from experience." Proceedings of the Sixth International Joint Conference on Artificial Intelligence. Stanford, CA: Stanford Computer Science Department.

KUIPERS, B. J. (1978) "Modelling spatial knowledge." *Cognitive Science* 2: 129-153.

LEE, T. R. (1970) "Perceived distance as a function of direction in the city." *Environment and Behavior* 2: 40-51.

LYNCH, K. (1960) *The Image of the City.* Cambridge: MIT Press.

MARR, D. and K. H. NISHIHARA (1978) "Representation and recognition of the spatial organization of three-dimensional shapes." *Proceedings of the Royal Society B* 200: 269-294.

McDERMOTT, D. V. (1980) "Spatial inferences with ground, metric formulas on simple objects." New Haven, CT: Yale University, Department of Computer Science Research Report 173.

NILSSON, N. (1980) *Principles of Artificial Intelligence.* Palo Alto, CA: Tioga.

PIAGET, J., B. INHELDER, and A. SZEMINSKA (1960) *The Child's Conception of Geometry.* New York: Basic Books.

SHEPARD, R. N. and S. CHIPMAN (1970) "Second-order isomorphism of internal representations: shapes of states." *Cognitive Psychology* 1: 1-17.

SIEGEL, A. W. and S. H. WHITE (1975) "The development of spatial representations of large-scale environments," in H. W. Reese (ed.) *Advances in Child Development and Behavior (Vol. 10).* New York: Academic Press.

TOBLER, W. (1975) "The geometry of mental maps," in R. G. Golledge and G. Rushton (eds.) *Spatial Choice and Spatial Behavior: Geographical Essays on Perception and Choice.* Columbus: Ohio State Univ. Press.

TOLMAN, E. C. (1948) "Cognitive maps in rats and men." *Psychological Review* 55:189- 208. (Reprinted in R. M. Downs and D. Stea (eds.) *Image and Environment.* Chicago: Aldine, 1973.)

TROWBRIDGE, C. C. (1913) "On fundamental methods of orientation and 'imaginary maps'." *Science* 38: 990.