MODELLING SPATIAL KNOWLEDGE

Benjamin Kuipers
The Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Abstract: A person's cognitive map, or knowledge of large-scale space, is built up from observations gathered as he travels through the environment. It acts as a problem-solver to find routes and relative positions, as well as describing the current location. The TOUR model describes the multiple representations that make up the cognitive map, the problem-solving strategies it uses, and the mechanisms for assimilating new information. The representations have rich collections of states of partial knowledge, which support many of the attractive qualities of common-sense knowledge.

Descriptive terms: spatial knowledge, cognitive maps, common-sense knowledge, partial knowledge, representations.

1. Introduction

The Puluwat Islanders of the South Pacific are famed for their navigational abilities. They frequently set out in small canoes, without the aid of modern instruments, and arrive reliably at a small island destination after travelling over a hundred miles of open ocean. Much less dramatically, each of us finds his way consistently between home, work, and other familiar destinations without the use of map or compass, and usually without conscious thought. A "sense of direction" is the way certain kinds of spatial abilities are often described. It is not often noted, however, that it is possible to have a sense of direction within two restricted regions, say Harvard Square and downtown Boston, while lacking any notion of how they *are* related. This suggests that a "sense of direction" is not a property just of the individual, but also of the location.

What these anecdotes have in common is that they deal with a person's cognitive map: the mental description which a person maintains of his environment. This research is restricted to common-sense knowledge of large-scale space obtained without the use of maps or instruments. From a functional point of view, the cognitive map is viewed as a black box which takes as input simple observations made while following a route, and produces as output a description of the current position, and the answers to route-finding and position-finding questions. In this case, however, the contents of the black box change as a consequence of the observations.

The scientific problem addressed by this research is to describe the contents of the cognitive map, how it is created from available observations, and how it answers questions about routes and positions. The answer to this problem is a descriptive formalism for the knowledge in the cognitive map, and a collection of inference rules: 1F-A0DED rules to accomplish the assimilation of new observations into the cognitive map, and IF-NEEDEO rules to answer questions.

A small set of metaphors for the cognitive map can illuminate the descriptive methods needed to characterize it. The first is the "Map in the Head" theory, which states that the cognitive map is just like a printed map, only examined by the mind's eye. This cannot be true of human cognitive maps because they often lack the global consistency of different relative position vectors that characterizes a printed map.

The Tatchwork Map" theory modifies this to suggest that the cognitive map is like an atlas of printed maps, each of which contains consistent relative position vectors, but which cannot necessarily be compared. This is quite a good theory, as we shall see below in the "Orientation" section, but ignores the capacity of human cognitive maps to represent route descriptions independently of relative position information.

The "Street Network" theory focusses on the ability of the cognitive map to represent the topological properties of route descriptions and street patterns, omitting geometrical information. Like the "Patchwork Map" theory, this is quite a good metaphor in that it recognizes an important kind of partial knowledge that people are prone to have in their cognitive maps. However, the use of uniform algorithms from graph theory to explain route-finding is unwarranted because it ignores geometical strategies that people *are* quite conscious of using.

Thus, neither of these metaphors alone can explain the cognitive map. In this paper, I describe a hybrid theory, taking some aspects of each of these mini-theories. A great deal of additional mechanism is required to make such a theory work, and it has been extended to explain acquisition of the cognitive map as well as its use in problem-solving.

This paper focusses on the representation for spatial knowledge in the cognitive map, and on the inference rules used for assimilation. These representations have been implemented as a large LISP program and tested on extensive examples. This research is reported in greater detail in [Kuipers 1977].

A great deal of data on human spatial cognition has been collected by psychologists and urban planners [Lynch, 1960; Downs and Stea, 1973; Siegel and White, 1975; Moore and Golledge, 1976]. Inevitably, these observations often do not directly address the issues involved in the creation of a computational model. Thus, in many cases, design decisions were made from general computational characteristics of common-sense knowledge. A representation for common-sense knowledge should support many states of partial knowledge, so that initial acquisition of information is very easy, assimilation into more elaborate representations can take place through a series of small steps, and performance can degrade gracefully when resources are restricted.

2. The TOUR model

The model of spatial cognition (called the "TOUR model") contains three different classes of representations for knowledge:

1. Representations for knowledge about a particular environment;
2. A description of the current position;
3. Representations for inference rules which manipulate knowledge of the other two kinds.

The purpose of the TOUR model is to show how knowledge in these different representations is combined, and how the inference rules translate it from one representation to another. Many of the inferences are organized around the process of following a route description through the cognitive map.

Knowledge about particular environments is encoded by the TOUR mods! in five different representations:

1. An imperative route representation directs a traveller along • particular route through the map.

2. A topological representation for local properties of street networks, including the ordering of places on a street and the local geometry of streets at an intersection.

3. Frames of reference, called "orientation frames", define relative positions of objects. Positions defined with respect to different orientation frames are not comparable unless a relationship between the orientation frames is represented explicitly.

4. Dividing boundaries provide a qualitative representation for position that can easily be transformed into a route.

5. A structure of containing regions provides levels of abstraction for stating relations among their elements.

These representations consist of descriptions of route-instructions, places, paths, regions, and orientation-frames. A description is made up of a number of properties and their values, implemented in LISP as the property list of a generated atom and a collection of associated access functions.

The current position of the traveller is represented by a small working memory called the "You Are Here" pointer, which describes the current position in terms of place, path, one-dimensional orientation on that path, current orientation frame, and two-dimensional orientation with respect to that orientation frame.

```
YOU ARE HERE:
        PLACE:  < place description >
        PATH:  < path description >
        DIRECTION:  < 1-D orientation:  +1 or -1 >
        ORIENT:  <orientation-frame description>
        HEADING:  < 2-D orientation:  0 to 360 >
```

Most manipulations of knowledge in the TOUR model take place through an interaction between the environmental descriptions and the "You Are Here" pointer. Furthermore, the only environmental descriptions which are accessed are typically the ones referred to by the "You Are Here" pointer and the current route instruction. These amount to a focus of attention for the inference rules that manipulate the descriptions.

Both the "You Are Here" pointer and the environmental descriptions may be incompletely specified. In most cases the TOUR model will function with incompletely specified descriptions, although with degraded performance. For example, it may be possible to follow a poorly specified route description from source to destination, but not to add information to place or path descriptions or to maintain orientation knowledge in the "You Are Here" pointer.

The inference rules that manipulate knowledge embedded in these various representations are represented as productions: simple modules that wait for a certain set of conditions to be true and then perform some action. They are organized around a process that follows a route description as a sequence of instructions to move the "You Are Here" pointer through the environmental description. Since this process resembles a computer executing a computer program, the collection of inference rules is known as the "TOUR machine." The inference rules also fall into several categories, corresponding roughly to the kinds of environmental representations.

1. Rules which compare the current route instruction, the "You Are Here" pointer, and the topological descriptions of the environment. They can act to fill gaps in each representation with information from the others. In particular, this is how the topological description is originally created from Information in the route description.

2. Rules for maintaining the current heading, or two-dimensional orientation, with respect to the current orientation frame. They operate with the relation between the one- and two-dimensional orientations represented in the "You Are Here" pointer and in the current place and path descriptions.

3. Rules which detect special structural features of a part of the environment, such as paths which act as dividing boundaries separating places. These rules act within the focus of attention provided by the current route instruction and the "You Are Here" pointer.

4. Rules which solve route-finding and position-finding problems using knowledge in the hierarchy of regions and in the descriptions of orientation frames and boundaries.

3. Route and Topological Representations

Information is first presented to the TOUR model as incomplete route descriptions. These are assimilated into a topological description of the environment in terms of paths and places. The more global descriptions of the environment are built on the foundation of this topological description. [An exception to this is a strategy for exploring unknown territory which is discussed in section 8 below.]

The topological representation consists of PLACE and PATH descriptions. The PATH description includes a partial order of PLACEs which are on that path. This partial order represents partial states of knowledge about the total order which places actually have on a path. A PATH has a one-dimensional Orientation with respect to this order: *1 represents facing in the direction of the order, and -1 represents facing against the order. The partial order datastructure is a list of sequences such that each place is in at least one sequence.

```
PATH:
        NAME:  <name>
        ROW:  <partial order datastructure>
```

A PLACE description includes a description of the local geometry of paths which intersect at that place. This local geometry describes the relations among paths, their one-dimensional orientations, and their radial headings in the local frame of reference of this intersection. The local geometry makes it possible to predict the results of turns.

```
PLACE:
        NAME:  <name>
        ON:  <list of PATHs>
        STAR:  <local geometry datastructure>
```

For the purposes of the topological representation, the "You Are Here" pointer describes three aspects of the current position: the current place, the current path, and the current one-dimensional orientation on that path. Some of these may be left unspecified, of course.

```
YOU ARE HERE:
        PLACE:  <place description>
        PATH:  <path description
        DIRECTION:  < +1 or -1 >
```

The route description is a sequence of GO-TO and TURN instructions leading from one place to another. GO-TO instructs the TOUR machine to move the "You Are Here" pointer from one given place to another, in a given direction along a given path. TURN instructs the TOUR machine to move the "You Are Here" pointer from one path and direction, through a certain number of degrees, ending on another path and direction, though at the same place.

The TOUR model treats an instruction as a simple observation by assuming that the instruction is realizable: both places linked by a GO-TO are on the given path with the given direction between them, and two paths have the relation required by the TURN instruction involving them.

When following a route description, the TOUR machine initializes the "You Are Here" pointer to the source of the route. Each instruction has the effect of changing the "You Are Here" pointer to its destination until the end of the route has been reached. Meanwhile the inference rules that make up the TOUR machine take fragments of information from one description and put it into parts of the others that have been left unspecified. These inferences are of several kinds:

1. Inference rules that take information about the current position in the "You Are Mere" pointer and use it to fill unspecified parts of the current instruction. For example, the "You Are Here" pointer may have a DIRECTION component provided by previous inferences, and it can be used to supply the missing 0IRECT10N component of the current GO-TO instruction.

2. Inference rules that take information from the current instruction and add it to the description of the current place or path (specified by the "You Are Here" pointer). For example, if a GO-TO instruction states that two places are related by a given direction on a given path, this information can be added to the partial order of places which is part of that PATH description. At a more basic level, these rules add a new PLACE or PATH description to the cognitive map when required by the current instruction.

3. Inference rules that take information from the current place or path description (specified by the "You Are Here" pointer) to provide missing information for the current instruction. For example, the current PATH description may be able to supply missing information about the direction component of a GO-TO instruction.

4. A fourth kind of inference fills gaps in a route description by posing them as problems for the problem-solving component of the TOUR model. The solution to the problem can itself be an incomplete route-description, requiring further calls on the problem-solver.

The states of partial knowledge in the route and topological representations result from the partial order and local geometry descriptions, and from the ability of the TOUR machine to tolerate underspecified descriptions of route and environment. When these descriptions are fully specified, the TOUR machine has a greater ability to fill in missing parts of new descriptions.

These states of partial knowledge make it possible for the individual inference rules to be *very* simple, so that a single pass through a route-description can assimilate useful amounts of information into the topological representations at low cost. Usually, several passes through a particular route-description are necessary before all the useful information is extracted. Even before this is done, the partially assimilated knowledge can be used.

4. Topological Assimilation Example

This example shows how the route description and the environmental descriptions interact to fill unspecified parts in each other. The scenario takes place near Central Square in Cambridge, whose simplified map is:

```
                                    Broadway
             ------------|---------------------------X
                         |                      Kendall
                         |                        Sq.
                         |Prospect St.
                         |
         Putnam          |
         Circle          |          Mass Ave.
X------------X-----------|---------------------------X
Harvard               Central                      MIT
  Sq.                   Sq.
```

We begin this scenario at Central Square, on Prospect Street, having come from Broadway. The current position is in the "You Are Here" pointer:

```
YOU ARE HERE:
        PLACE:  [PLACE2: Central Square]
        PATH:  [PATH3: Prospect Street]
        DIR:  +1
```

The relevant parts of the cognitive map are the descriptions of Central Square, Mass Ave, and Prospect Street.

```
PLACE2:
        NAME:  Central Square
        ON:  [PATH1: Mass Ave]
             [PATH3: Prospect Street]
        STAR:  (0. PATH1 -1)
               (90. PATH3 -1)
               (180. PATH1 +1)

PATH1:
        NAME:  Mass Ave
        ROW:  ([PLACE1: Harvard Square]
               [PLACE2: Central Square]
               [PLACE3: MIT])

PATH3:
        NAME:  Prospect Street
        ROW:  ([PLACE4: Broadway & Prospect Street]
               [PLACE2: Central Square])
```

The first instruction of the tiny route we will follow is:

"Turn right."

whose internal description has six elements: the place, path, and direction preceding the turn, the amount of the turn, and the path and direction resulting from the turn. In this case, most of the elements are left unspecified.

```
TURN:
        AT:
        ST1:
        DIR1:
        AMT:  90.
        ST2:
        DIR2:
```

The "You Are Here" pointer provides the current context to fill in the first three missing elements of the TURN instruction:

```
TURN:
        AT:  [PLACE2: Central Square]
        ST1: [PATH3: Prospect Street]
        DIR1: +1
        AMT:  90.
        ST2:
        DIR2:
```

At this point, the local geometry description in the STAR *properly* of PLACE2 can be used to predict the result of the turn. The local geometry associates a heading with certain (PATH DIRECTION) pairs as they radiate from a place, and the amount of the turn specifies the new heading which can, perhaps, specify a new (PATH DIRECTION) pair. The absolute values of these headings are meaningless, and can only be used to compute such differences. In this case, the result is to specify the TURN instruction completely:

```
TURN:
        AT:   [PLACE2: Central Square]
        ST1:  [PATH3: Prospect Street]
        0IR1: +1
        AMT:  90.
        ST2:  [PATH1: Mass Ave]
        DIR2: -1
```

The final operation is to update the "You Are Here" pointer to reflect the result of the TURN instruction.

```
YOU ARC HERE:
        PLACE:  [PLACE2: Centrel Square]
        PATH:   [PATH1: Mm Ave]
        OIR:    2±
```

The second sentence of our brief tour takes us to Putnam Circle, which is shown on the map above, but which is completely new to the program, so it must create a new PLACE description for it.

"Take Mass Ave to Putnam Circle."

```
GO-TO:
        FROM:
        TO:   [PLACE5: Putnam Circle]
        PATH: [PATH1: Mess Ave]
        OIR:
```

The "You Are Here" pointer again provides the current context, including the direction along Mass Ave that we are travelling, so we can fill in the missing parts of the GO-TO instruction:

```
60-TO:
        FROM: [PLACE?: Centrel Squere]
        TO:   [PLACE5: Putnam Circle]
        PATH: [PATH1: Mass Ave]
        OIR:  zi.
```

When processing the previous TURN instruction, the PLACE description was used to add information to the instruction. Here the completely specified GO-TO instruction will be able to add information to the description of the environment. First, we add to PLACE5 the fact that Putnam Circle is on Mass Ave.

```
PLACE5:
        NAME: Putnam Circle
        0N:   CPATH1: Mais Ave]
        STAR:
```

Second, since the GO-TO instruction gives an order relation (-1) between Central Square and Putnam Circle, we can add this information to the partial order in PATH1. Notice that we don't know where Putnam Circle is with respect to Harvard Square, but we do know that both are on the same side of Central Square.

```
PATH!:
        NAME: Mass Ave
        ROW:  (PLACE1 PLACE2 PLACE3)
              (PLACES PLACE2)
```

It is illuminating to consider the effect of a partially specified "You Are Here" pointer. If there had been no local geometry-information in PLACE2 about Central Square, for example, the direction of travel would have been unspecified in the GO-TO instruction, and the partial order in PATH1 showing the position of Putnam Circle would have been:

```
PATHI:
        NAME: Mess Ave
        ROW:  (PLACE 1 PLACE2 PLACE3)
              (PLACE3)
```

The route description still contains the information that a right turn from Prospect Street at Central Square points you toward Putnam Circle, but it is not represented in the more easily accessible PLACE and PATH descriptions.

This example has shown how the topological properties of places and paths are represented in their descriptions, and how information is assimilated from the relatively inaccessible route instructions into the more globally useful environmental descriptions. The assimilation process takes place through v*ery* simple, and computationally inexpensive, interactions between the current instruction, the environmental descriptions, and the "You Are Here" pointer. Notice that, since the only environmental descriptions accessed are those referred to by the current instruction and the "You Are Here" pointer, that processing time is independent of the amount of information represented.

5. Orientation

A sense of direction is the ability to define one's current heading (or two-dimensional orientation) with respect to the positions of remote (and often Invisible) places. The topological representation supports only a one-dimensional orientation with respect to the order on a particular path. Thus, we must augment the TOUR model to include descriptions of frames of reference, two-dimensional headings in the "You Are Here" pointer, and knowledge in the PATH and PLACE descriptions about the relation between one- and two-dimensional orientations.

An orientation-frame is a common frame of reference for the heading component of positions relating two places, and for the heading component of the "You Are Here" pointer. By having many different orientation-frames, a person may represent the positions of many places without the requirement that they fit into a single consistent framework. Thus the constraints on newly added position information are relatively weak, so the position representation has more states of partial knowledge. Knowledge about the relation between two orientation-frames is represented as part of the orientation-frame descriptions.

This independence of orientation-frames is required because it is quite common for a person to be well-oriented within each of two different regions, but have *yery* little rfotion of the relation between them. Knowledge about the relation between two orientation frames can be learned or forgotten separately from the position information within each one.

The "You Are Here" pointer must be augmented to include the current orientation-frame and the current heading with respect to that orientation-frame.

```
YOU ARE HERE:
          PLACE:  <place description>
          PATH:  <path description
          DIRECTION:  < +1 or -1 >
          ORIENT:  <orientation-frame description>
          HEADING:  < integer in [0,360] >
```

PLACE and PATH descriptions may contain information about the relationships between one-dimensional and two-dimensional orientation. The local geometry of a PLACE can be defined so that its headings are consistent with those of a particular orientation-frame. The PATH description may contain a number of different headings for travel along its +1 direction, defined with respect to a number of different orientation-frames.

As the TOUR machine drives the "You Are Here" pointer along the route, its problem is to maintain the current HEADING and to transfer orientation information between the "You Are Here" pointer and the PLACE and PATH descriptions. There are three Kinds of inference rules to accomplish this:

1. Inference rules that update the current heading for a TURN whose amount is known, *Bnd* that check to see that the PATH of a GO-TO is straight before allowing the heading to remain fixed.

2. Inference rules that set HEADING or DIRECTION in the "You Are Here" pointer by examining the current PLACE and PATH descriptions.

3. Inference rules that add information to the current PLACE •nd PATH descriptions about the interaction between the current HEADING and DIRECTION as they appear in the "You Are Here" pointer.

In addition to representing knowledge from visual observation or verbal report, knowledge of heading can be used to implement a "dead reckoning" technique for computing the relative positions of the source and destination of a given route. Dead reckoning requires the '"You Are Here" pointer to hold X and Y values for the current position in rectangular coordinates with respect to the current orientation frame. The distance travelled by a GO-TO on a known heading can be converted to those rectangular coordinates, and the result at the end of the route converted back to polar coordinates.

6. Two-Dimensional Orientation Example

Information about the current heading is maintained and updated in much the same way as the topological information in the previous example. For example, in the "Turn right" instruction at Central Square, the enlarged "You Are Here" pointer would acquire a heading and orientation-frame from the local geometry. description in PLACE2.

```
PLACE*:
          NAHE:  Central  Square
          ON:   [PATHI: Mass Ave]
                CPATM3: Prosptct Strttt]
          STAR:   (0. PATHI -I)
                  (90. PATH3 -1)
                  (ISO. PATHI +1)
          ORIENT:   ORIENTS

YOU ARE HERE:
          PLACE:   [PLACE?: Ctntral Squart]
          PATH: [PATHI: Jit si Avt]
          DIR:   -1
          ORIENT:   0RIENT3
          HEADING:   0
```

The domain of a given orientation-frame can also propagate along lines of frequent travel. Assume that we arrive at Putnam Circle and make a turn, so that information must be added to its local geometry description. Rather than creating a new, local orientation-frame for just that place, its local geometry would be defined with respect to ORIENT3, and would thus be closely related to the local geometry of Central Square.

Two orientation-frames collide when one defines the local geometry of the current place, while the other defines the heading in the "You Are Here" pointer. A collision can have two outcomes. If one of them is local to its particular place, it is usually replaced by the orientation-frame with the larger domain, which therefore continues to propagate. If both have substantial domains, the relationship between them is stored in the orientation-frame descriptions. In either case, the heading in the "You Are Here" pointer can be maintained and updated along a route that travels quite far from where its orientation-frame was defined.

At the same time, the HEADING property of a PATH description can hold the heading of that path (in the +1 direction) with respect to a given orientation frame. In this case:

```
PATHI:
          NAHE:   Mtss  Avt
          ROW:   (PLACE1  PLACE2  PLACES)
                 (PLACE5  PLACE2)
          HEADING:   (0RIENT3  160)
```

If more than one heading is put into the same PATH description, we again have an opportunity to compute and store the relationship between two orientation-frames.

7. Boundaries

Dividing boundaries are very useful in specifying the location of a place. For example, I can describe the location of the Cambridge Public Library by saying that it is:

on the other side of Prospect Street,
on this side of Mass Ave,
between Broadway and Cambridge Streets,
beyond Trowbridge Street,
before Quincy Street.

In each part of this description, I am using a street to draw a boundary, dividing the world (or at least a small part of it) into two regions, then specifying which of those regions contains the place I am describing. A boundary in this sense is not a barrier: it acts to define groups of positions, not to impede travel. Naturally, a barrier like the Charles River can also function as a boundary.

When a path acts as a boundary, the sides on the right and the left when facing the +1 direction are represented by REGIONS (sets of PLACEs) in the RIGHT and LEFT slots of the PATH description.

```
PATH:
          NAME:  <name>
          ROW:  <par»ial order of PLACEs>
          HEADING:  <list of pairs:
                          (orientation-frame heading) >
          RIGHT:  <region>
          LEFT:  <region>
```

Thus, specifying where a place lies with respect to a dividing boundary provides partial knowledge about its position. This knowledge is particularly easy to acquire, easy to combine with other similar pieces of knowledge, end easy to apply to route-finding problems.

A single turn in a route description can specify where a place lies with respect to a dividing boundary. If the route to the Cambridge Public library involves a left turn off Mass Ave, then the Cambridge Public Library is on the left side of Mass Ave, no matter how tortuous the rest of the route is. (Providing, of course, that it doesn't cross back over Mass Ave.)

Although in principle the two regions defined by a boundary extend to infinity, in practice they include only places whose relationship to the boundary has been brought to the attention of the TOUR machine. Thus, if a street is involved in a large variety of different routes, it will have many boundary relationships with different and distant places. Otherwise the division it represents may apply only to the immediate neighborhood..

The description of the position of the Cambridge Public Library translates readily into useful constraints on a route from here to there. Prospect Street and Trowbridge Street are both potential intermediate subgoals, while Mass Ave and Quincy Street can act as barriers in case the route strays too far from the goal. In fact, any of the given streets can act as skeletons for a route, since they all have known relationships with both source and goal. If connections are found from source to skeleton and from skeleton to goal, then the skeleton street can join the pieces into a complete route.

If two boundaries are known to be parallel, then the regions on their sides have further useful relationships. For example, Prospect Street and Trowbridge Street are parallel and both lie between here and the Cambridge Public Library. By knowing the order of the two parallel boundaries, Prospect Street can be chosen as the first subgoal, followed by Trowbridge Street. A bundle of parallel streets allows subgoals to be ordered into a sequence of small steps that can be easily joined to construct a route.

Two streets can be found to be parallel by examining their headings. The two streets fall within the focus of attention of the TOUR machine when they are cross-streets encountered on a GO-TO instruction. This allows "local parallel" relations to be found linking two streets. When a particular street has several local parallel relations to other streets, a gathering operation is initiated to follow the local parallel links and gather up and order a bundle of parallel streets.

A rectangular grid structure on an area amounts to two such bundles, perpendicular to each other. This makes route-finding even easier because the sequence of subgoals can be found within one bundle, while the connections are in the other. Thus, the rectangular grid is a useful description of the geography because it allows certain route-finding strategies to be used very powerfully. Since this route-finding power is relatively ins'ensitive to irregularities in the geography, people are led to apply the grid description even when it is incorrect.

8. Exploration

Finally, we can discuss an interesting technique for exploring unknown territory. How does a person explore an unfamiliar area before he knows the topology of the street network? It is clear that accomplished explorers use their sense of direction to find the way back to familiar places while learning the new area. If an explorer can maintain his own heading with respect to a familiar street, and if he knows what side of that street he is on, he can always navigate back toward it when he wants to. What it the knowledge that permits him to do this?

A person in a new area can define an orientation frame by the position of a prominant landmark, for example the John Hancock Building in Boston. This orientation frame can allow him to define the heading of a familiar street. Then, while exploring, he must maintain his current heading with respect to that orientation frame, and remember what side of the familiar street he is on. He can maintain the heading by attending only to the amounts of turns and the curves of streets. If the unknown territory can be assumed to have a grid structure, the problem becomes much easier, because the heading can have only four values. Then he always knows the relative heading of the familiar street, and can always guide himself toward it, even with no knowledge of the street network.

Observations of people learning an area have revealed that while newcomers orient themselves with respect to conspicuous landmarks, long-time residents very seldom do. [Lynch 1960] Those with detailed cognitive maps of an area can orient themselves by local features of each place in the street network. Furthermore, they often have a sufficient stock of familiar routes that they need not orient themselves at all, but can just follow route descriptions.

This example shows how very small amounts of orientation and boundary knowledge can combine to make a powerful technique for exploring unknown territory. Only because the representations have rich collections of states of partial knowledge can such small amounts of information be used so effectively.

9. Regions

Regions allow places to be grouped and referred to collectively. As such, they provide levels of abstraction for stating facts and answering questions. For example, I can give a route for getting from the West Coast to the East Coast and hope that it can be used to solve problems concerning particular places in the two regions.

Regions are often defined in terms of legislative boundaries, visual texture, typical activities, ethnic composition, and other characteristics that are not strictly aspects of spatial cognition. Thus, unlike the other aspects of the cognitive map, I will not talk about how a region description is created, only how it is used. The problem is how the relationships among different region descriptions allow information to be stated at one level of abstraction and used at another.

Dave Rumelhart (1974) of UCSD has discussed the problem of sensitivity to geographical context in question-answering. When in San Diego, why is it far from UCSD to the airport, but Boston seems near to New York? The answer he provides in his "Room Theory" is that the geography is described with a hierarchy of nested regions. The first step in answering such a question is to find the ROOM: the smallest region containing all the places of interest (those in the problem plus the reference position of the conversation). Then the answer is stated in the context of that ROOM. If the questions are asked in San Diego, then the ROOM for considering the airport is quite small so the distance must be considered large. However, the ROOM containing San Diego, Boston, and New York is the whole country, within which Boston is quite near to New York.

A hierarchy of nested regions can be represented quite easily by having a region description point to the next larger containing region. [For purposes of this discussion, a PLACE and a REGION •*re* indistinguishable.] The set of nested regions about a given place forms a sequence, so it is easy to compare two to find the smallest common region. Unfortunately, it is difficult to add a new region to such a structure, since its relationship with all other regions must be known before it can be merged into the sequence. In the TOUR model, the nesting requirement is relaxed to provide many more states of partial Knowledge without great loss of performance.

In the TOUR model, a region description points to an unordered collection of containing regions. When a problem is posed concerning a particular place, the set of regions containing it is obtained by following these "upward pointers" and ordering them with the same partial order mechanism we have seen twice before. The longest totally ordered subset is taken from this partial order to be used for problem-solving. This makes it possible to miss relationships which are actually represented, but only in cases where the order is partial, and then only involving the least closely related regions.

To implement Room Theory problem solving, the sequences of containing regions from the places of interest *are* compared from the ^{H}top down" to find the smallest containing Room. The diverging parts of the two sequences are then examined to see if the desired relationship is found between regions contained in the Room. If several applicable relationships are found, the most specific is used. As well as being computationally efficient, this top-down access fits the experimental results of Stevens (1976), who tested the relative difficulty people encountered in answering questions about geography.

When an abstract relationship is found between regions containing two places of interest, it must still be matched to the concrete problem as originally stated. If my original problem was to get from Stanford to MIT, I might discover that a good route from Northern California to New England was to take 1-80 to Cleveland, then 1-90 to New England. To complete the details of the route, I must shift my focus of attention at the endpoints of the route to make them more specific.

This requires that a region description include information to permit a mapping from a general description of a place to a more specific one. The region Northern California does not correspond to a particular more specific place or region, but Northern California on 1-80 does correspond to a specific place in San Francisco (actually, several places). Then I can reformulate my problem to get from Stanford to that place. Similarly, I can get from 1-90 in New England to MIT.

Thus, the structure of containing regions permits generally useful information to be stated at an abstract level and used at a more specific one. The containment relations permit a particular place to have a partially ordered set of containing regions, rather than simply a sequence. This places fewer constraints on the creation of new region descriptions. Only when a problem is posed are the local containment relations merged into a unified structure. It is necessary to have a downward mapping, going from more abstract to more specific descriptions of places. This is impossible if it relates only PLACE descriptions, but can be implemented by mapping more abstract values of the "You Are Here" pointer to more specific ones.

10. Conclusions

To summarize all this, I have presented the TOUR model as a model of spatial knowledge, consisting of a number of representations with rich collections of states of partial knowledge. The TOUR model initially divides into environmental descriptions, the "You Are Here" pointer, and the inference rules that manipulate them. A second division of the representations distinguishes five different kinds of knowledge: route descriptions, topological street networks, orientation-frames for position knowledge, dividing boundaries and grids, and structures of containing regions.

Considered as a psychological theory, the TOUR model provides a plausible explanation for the observations that have been made about human spatial cognition. Most importantly, it provides a framework for stating detailed empirical questions and for representing their answers. This kind of data will allow the model to make detailed predictions about spatial understanding.

The empirical questions which can be asked include questions about the distortions and geographical paradoxes that people are prone to, the extent of individual variation in spatial abilities* the relative difficulties of different inferences, and the order in which spatial abilities, are acquired by children. The answers to these questions will be represented in the detailed structure of the environmental descriptions and the collection of inference rules that manipulate them.

Considered as a method for representing common-sense Knowledge, the TOUR model provides several useful representations with rich collections of states of partial knowledge. Furthermore, it shows how knowledge is assimilated from an initial, very local representation into more powerful and global descriptions. Since people often apply spatial metaphors to other kinds of knowledge, there is promise that these representations can be applied to other, non-spatial domains of common-sense knowledge.

11. References

R. M. Downs and D. Stea (Eds.). *Image and Environment.* Chicago: Aldine Publishing Company, 1973.

B. Kuipers. Representing Knowledge of Large-Scale Space. Cambridge, MA: MIT Artificial Intelligence Laboratory TR-418, July 1977.

K. Lynch. *The Image of the City.* Cambridge: MIT Press, 1960.

G. T. Moore and R. G. Golledge. (Eds.) *Environmental Knowing: Theories, Research, and Methods.* Stroudsburg, PA: Dowden, Hutchinson and Ross.

0. E. Rumelhart. The Room Theory. Unpublished computer listing. La Jolla, CA: University of California, San Diego, 1974.

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

A. L. Stevens. The Role of Inference and Internal Structure in the Representation of Spatial Information. Unpublished doctoral dissertation. Psychology Department, University of California at San Diego, La Jolla, CA, 1976.

This report describes research done 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.