

The Skeleton in the Cognitive Map: A Computational Hypothesis*

Benjamin Kuipers
Computer Science Department
University of Texas at Austin
Austin, Texas 78712 USA

February 2, 2001

Abstract

Experts seem to find routes in a complex environment by finding a connection from the starting place to a subset of major paths — the “skeleton” — then moving within the skeleton to the neighborhood of the destination, making a final connection to the destination. We present a computational hypothesis to account for the skeleton as an emergent phenomenon, arising from the interaction of three factors. (1) The topological map is represented as a bipartite graph of places and paths, where a path is an extended one-dimensional description of an ordered set of places. (2) Travel through the environment allows the traveler to incrementally accumulate topological relationships, including the relation of a place to a path serving as a dividing boundary separating two regions. (3) A bounding path is often a natural subgoal during way-finding search, meaning that paths rich in boundary relations are likely to appear in routes, which means they are likely to acquire more boundary relations. This positive-feedback loop leads to an oligarchy of paths rich in boundary relations. We propose computational and empirical tests for this hypothesis.

Keywords: cognitive map, way-finding, topological map.

*This work has taken place in the Intelligent Robotics Lab at the Artificial Intelligence Laboratory, The University of Texas at Austin. Research of the Intelligent Robotics lab is supported in part by NSF grants IRI-9504138 and CDA 9617327, and by funding from Tivoli Corporation.




Figure 1: A large-scale cognitive map has a skeleton of major paths. The graphical conventions for emphasizing major streets and highways on a printed map are related, but not identical to the skeleton of the cognitive map.

Expert way-finders in a complex large-scale environment use a “skeleton” of important paths and places to guide their problem-solving [7, 8, 1]. How is this skeleton represented? How is it acquired? And how does it help in way-finding?

In this note, I describe a preliminary computational hypothesis to explain the phenomenon of the skeleton, based on the concepts in the Spatial Semantic Hierarchy [6] (which extends the TOUR Model [4, 5]). I also describe a set of computational and cognitive tests that could be applied to this hypothesis. This note proposes a hypothesis and raises questions for investigation, rather than providing definitive answers.

1 The Skeleton

Researchers who have studied expert way-finders such as experienced taxi drivers [8, 1, 11, 3] have observed a common strategy. Such an expert knows a large number of places and paths, but much of their travel occurs within a small subset of “major” paths, which is sometimes called the “skeleton” (Figure 1). When given a way-finding problem, the expert first finds a route from the initial point to the nearest point on the skeleton, then finds a route within the skeleton to a point near the destination, and finally finds a route from that point to the destination itself.

This sketch raises several questions. How are the paths and places in the skeleton selected from the larger set the expert knows about? Is there a qualitative difference between the skeleton and the rest of the map, or is the role of the skeleton an emergent behavior of some uniform mechanism applied to the entire cognitive map?

The hypothesis presented here is that the skeleton is an emergent phenomenon, arising from the interaction between:

1. the topological representation for places and paths;

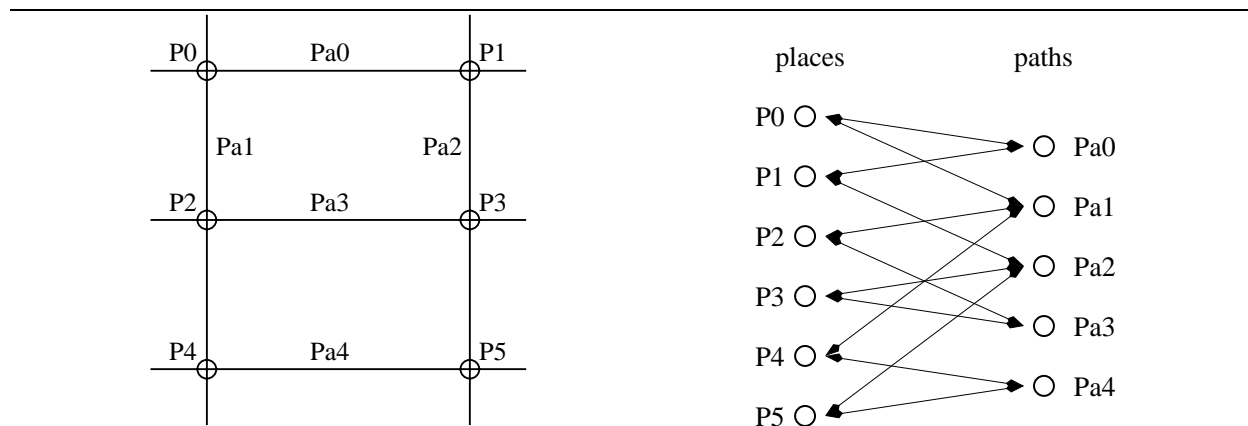


Figure 2: The rectangular block environment (left) is described in the topological map (right) as a bipartite graph of places and paths.

2. the incremental, opportunistic learning of “boundary relations” during travel; and
3. the use of boundary relations to provide subgoals during way-finding.

2 The Topological Map

It is widely accepted [7, 10] that the “cognitive map” includes a topological level of description, in which places (0-D), paths (1-D) and regions (2-D) are symbolically described and linked by relations such as connectivity, order and containment. Metrical relationships such as distance and direction may also be associated with the topological map, but there is typically no single global frame of reference, and metrical errors in a variety of tasks are much more common than topological errors.

The Spatial Semantic Hierarchy (SSH) [6] is a computational model of knowledge of large-scale space, consisting of a lattice of different but related representations for space. The control level consists of knowledge of control laws for taking the agent from one “distinctive state” within the environment to another. A state (position plus orientation) is distinctive if it is the stable point of a local “hill-climbing” control law that eliminates moderate amounts of accumulated error by bringing the agent to a particular state from anywhere in its local neighborhood. The causal level of the SSH abstracts the control laws to actions and represents behavior in the environment as a set of discrete causal schemas $\langle S, A, S' \rangle$, describing the relation between a state, an action, and the resulting state. The topological level posits places, paths and regions to account for the experienced regularities in the causal description. Local pieces of metrical information can be used throughout the other levels, but a global metrical model with a single frame of reference can only be created after all the other descriptions exist. The TOUR Model [4, 5] has been incorporated into the causal and topological levels of the SSH.

In the SSH topological map, a path describes an extended one-dimensional structure such as a street. The topological map is a bipartite graph (Figure 2), with nodes corresponding to places and paths, and arcs corresponding to the assertion that a particular place is on a particular path. A place is on one or more paths, and it includes the circular order of departing directed paths. A path includes an ordering on its set of places. The benefit of the bipartite graph of places and paths is that physically distant places on the same path may




Figure 3: Boundary relations. A path serves as a boundary separating places on the left from those on the right. Boundary relations can be inferred from local travel patterns.

be close in the topological map, making way-finding easier.

3 Boundary Relations

The SSH topological map can represent more than connectivity and order. A path is a one-dimensional subset of the environment, with a direction implied by the order on its places. A directed path is described by (Pa, dir) , where dir is either *pos* or *neg*, and $-dir$ is the other one. If a directed path extends to infinity, it divides the places in the environment into three subsets: those on the path, those on the right, and those on the left. Note that “right” and “left” are used here as topological terms. If the path curves, a place that is topologically to the right may occasionally be visible to the traveler’s egocentric left.¹

The assertion that a place P lies to the right of a directed path (Pa, dir) is called a *boundary relation*: $right_of(Pa, dir, P)$. (See Figure 3(left).) We define $left_of$ similarly, and provide the axiom

$$right_of(Pa, dir, P) \equiv left_of(Pa, -dir, P).$$

Boundary relations can be acquired incrementally during travel by simple local rules. For example (Figure 3(right)):

If the traveler moves along a path $(Pa1, dir1)$ from place $P0$ to place $P1$, and takes a right turn at $P1$ onto $(Pa2, dir2)$, and travels along $(Pa2, dir2)$ to reach place $P2$, then we can conclude $right_of(Pa1, dir1, P2)$.

We can also conclude that $right_of(Pa2, dir2, P0)$.

We are making some relatively weak and plausible assumptions. We assume that the finite length of the boundary $Pa1$ does not lead us astray. We assume that path $Pa2$ does not intersect or cross $Pa1$. These inferences are implemented as *default rules*, so that if there is contrary evidence, no conclusion is drawn [9].

¹Example: When traveling east along the Charles River separating Boston and Cambridge, Boston is topologically to the right. However, because of the curve of the river, the highly visible John Hancock Tower in Boston can sometimes be seen in the distance to the traveler’s left. The Boston area is a treasure-trove of spatial paradoxes for the cognitive map theorist.




Figure 4: Way-finding using boundary relations.

The above rule applies only when there is a direct connection from $Pa1$ to $P2$, but it is straight-forward to handle more complex connections. Meanwhile, we are not assuming that $Pa1$ or $Pa2$ are straight. We are not assuming that a right turn is a 90 degree turn. We are not assuming that $P2$ is close to $Pa1$, since the path $Pa2$ can be quite extended.

Using local rules like the one above, any experience traveling through the environment will lead the topological map to accumulate boundary relations among places and paths experienced during travel. This is the first link in a positive-feedback system to ensure that paths that are used frequently tend to be used more frequently.

4 Way-Finding Using the Boundary Heuristic

Way-finding is the process of finding a route from an initial place to a destination place. At the SSH topological level, a route is an alternating sequence of places and paths, each connected to its neighbors. Once a topological route is found, it can be refined for execution, first to an alternating sequence of states and actions at the causal level, and finally as a sequence of control laws at the control level.

There are a number of graph search algorithms that can find paths in a topological map [2]. Metrical information such as estimates of path segment lengths can be used to guide heuristic search in the A^* and Dijkstra algorithms. However, the boundary relation can be used as the basis for a purely qualitative heuristic to guide way-finding search (Figure 3).

If we are searching for a route from place A to place B , and
 if there is a path Pa such that $left_of(Pa, dir, A)$ and $right_of(Pa, dir, B)$, or vice versa,
 then consider Pa a subgoal, and search for routes from A to Pa , and from Pa to B .

When places A and B are on opposite sides of path Pa , the route connecting them must necessarily cross the boundary. The heuristic can also be useful in case both places are on the same side of the boundary, though of course the route could be inefficient. Unfortunately, we can not in general expect the same boundary to be related to both endpoints of the desired route. A more general form of the boundary heuristic is (Figure 4):

If we are searching for a route from place A to place B , and

if there are paths $Pa1$ and $Pa2$ such that A shares a boundary relation with $Pa1$, and B with $Pa2$,
 then propose the subgoal of finding a connection from $Pa1$ to $Pa2$.

In the simplest case, $Pa1$ and $Pa2$ can be connected by sharing a place (Figure 4(left)). Since paths are extended 1-D subsets of the environment, this will not be uncommon.² In more complex cases (Figure 4(right)), we can search for a connection from $Pa1$ to $Pa2$, using the same heuristics.

The places A and B may have multiple boundary relations with different paths. If there are several possible boundaries, order them according to the number of boundary relations they have with other places. This will increase the probability of finding a useful connection earlier in the search. It is the second link in the positive feedback cycle that leads to the emergence of the skeleton.

5 A Positive Feedback Cycle

There is a positive feedback cycle between the inference of boundary relations and the effect of boundary heuristics on way-finding search.

- Travel along a path Pa makes it likely that a boundary relation, say $left_of(Pa, dir, P1)$, will be observed and inferred.
- The existence of a boundary relation $left_of(Pa, dir, P1)$ increases the probability that path Pa will be used in the solution to a way-finding problem, even if place $P1$ is not involved in the route, but of course more so if it is.
- Following the newly-found route, travel along the path Pa increases the probability that a new boundary relation, say $right_of(Pa, dir, P2)$ will be observed and inferred.

This is a self-reinforcing, “rich get richer” process, leading to an oligarchy of paths (the skeleton) rich in boundary relations. The skeleton perpetuates itself because way-finding most easily finds routes using paths within that subset. Note that there is no qualitative distinction between paths within the skeleton and those outside. There is simply a distribution of boundary relations among the paths in the cognitive map.

The Boundary Relation Hypothesis is:

The empirical phenomenon of the skeleton — that expert way-finders in an environment preferentially use a small set of important paths — is explained computationally by the positive feedback cycle between inference of boundary relations during travel and the use of the boundary heuristic during way-finding.

6 Research Questions

The Boundary Relation Hypothesis suggests a number of computational experiments that can be carried out on a simulated model of a real or artificial urban street network.

²Interesting statistical question: for a typical urban street plan, compare the probability that two randomly selected places share a path with the probability that two randomly selected paths share a place. It seems plausible that the former is orders of magnitude smaller than the latter, which would help explain the power of the boundary heuristic.

Implement a simulated agent that travels from place to place in the simulated environment model. As it travels, it builds its own topological map of the environment, including both connectivity relations between places and paths, and boundary relations.

The agent travels to a randomly selected sequence of destinations. When it can, it solves the way-finding problem itself, and follows the route it found. Otherwise, an “oracle” provides a set of route directions (“Take the first right, then the second left . . .”), which the agent follows, extending its cognitive map.

This process continues until the agent’s cognitive map covers the entire environment, so it is consistently able to solve new way-finding problems.

We can construct a variety of different experiments within this framework.

- The Boundary Relation Hypothesis states that, all else being equal, paths with larger numbers of boundary relations will be more frequently used in routes. After the cognitive map is complete, measure its behavior to see whether it has created a skeleton, and compare the distribution of path use with the distribution of boundary relations in the cognitive map.
- Applying the simulator to models of different geographical environments, compare the skeletal structures of paths developed in different geographical areas with the skeletal structures observed in human experts (e.g., taxi drivers) in those areas. This addresses the role of length of street, density of places on a street, presence of bottleneck points that many routes must pass through, and other geographical factors in determining which paths become part of the skeleton.
- Using the same simulator and evaluation metrics, examine the effect of the distribution of destinations in the learning phase on the cognitive map constructed by the traveler. A traveler with a wide range of travel will clearly learn a different cognitive map, and presumably a different skeleton, from a traveler with a more constrained region of travel.
- Compare the cognitive maps and skeletal structures developed by the simulator with those developed by human experts in the same environment, in order to evaluate and refine the rules and heuristics for inferring and using boundary relations.

These questions can be investigated in artificial environments where geographical variation is controlled, but in natural environments, we do not yet know when all else will be equal.

7 Extension to Topological Grids

So far, the only boundary relation we have considered is the relation between a boundary path B and a place P which lies on one side of B . However, it is also possible for a path C to lie entirely on one side of the boundary B . If paths B and C do not intersect, and travel from one to the other matches certain behavioral patterns, we can describe B and C as “topologically parallel” (Figure 5(left)). Inference of this description from local evidence must be defeasible, since extended paths can curve around so that future observations reveal very different properties.

The relation $path_right_of(B, dir, C, dir2)$ means that the region to the right of path B facing in direction dir includes path C , and that direction $dir2$ along C is “the same direction” as direction dir along B .




Figure 5: Building topological grids from boundary relations.

$$\begin{aligned}
 \text{path_right_of}(B, \text{dir}, C, \text{dir}) &\equiv \text{path_left_of}(C, \text{dir}, B, \text{dir}) \\
 \text{path_right_of}(B, \text{dir}, C, -\text{dir}) &\equiv \text{path_right_of}(C, \text{dir}, B, -\text{dir}) \\
 \text{path_right_of}(B, \text{dir}, C, \text{dir}) \wedge \text{right_of}(C, \text{dir}, P) &\rightarrow \text{right_of}(B, \text{dir}, P) \\
 \text{path_right_of}(B, \text{dir}, C, \text{dir}) \wedge \text{left_of}(B, \text{dir}, P) &\rightarrow \text{left_of}(C, \text{dir}, P) \\
 \text{path_right_of}(B, \text{dir}, C, -\text{dir}) \wedge \text{right_of}(C, -\text{dir}, P) &\rightarrow \text{right_of}(B, \text{dir}, P) \\
 \text{path_right_of}(B, \text{dir}, C, -\text{dir}) \wedge \text{left_of}(B, \text{dir}, P) &\rightarrow \text{left_of}(C, -\text{dir}, P)
 \end{aligned}$$

The value for way-finding of the “topologically parallel” relation between paths is in the systematic nesting of their boundary regions, and hence the relationship among their boundary relations. A set of paths that are pairwise topologically parallel provide an structured set of subgoals for finding routes. In the same way that an ordered set of places provides a qualitative measure of position along a path, a set of topologically parallel paths provides a qualitative measure of position within a larger area.

If there are two intersecting bundles of parallel paths, then together they provide a qualitative description of two-dimensional position, to the granularity of the spacing between paths (Figure 5(right)). The clearest examples of this sort of topological grid are of course the geometrical grids of streets that are laid out in many cities. However, the topological grid description applies, and is useful for way-finding, even when the topological grid is not a geometrical grid.

References

- [1] W. G. Chase. Spatial representations of taxi drivers. In D. R. Rogers and J. A. Sloboda, editors, *Acquisition of Symbolic Skills*. Plenum Press, New York, 1982.
- [2] R. J. Elliot and M. E. Lesk. Route finding in street maps by computers and people. In *Proc. 2nd National Conf. on Artificial Intelligence (AAAI-82)*, Cambridge, MA, 1982. AAAI Press/The MIT Press.
- [3] R. G. Golledge, editor. *Wayfinding Behavior: Cognitive Mapping and Other Spatial Processes*. The Johns Hopkins University Press, Baltimore MD, 1999.

- [4] B. J. Kuipers. Modeling spatial knowledge. *Cognitive Science*, 2:129–153, 1978. Reprinted in *Advances in Spatial Reasoning, Volume 2*, Su-Shing Chen (Ed.), Norwood NJ: Ablex Publishing, 1990, pages 171–198.
- [5] B. J. Kuipers. The ‘Map in the Head’ metaphor. *Environment and Behavior* 14, pages 202–220, 1982.
- [6] B. J. Kuipers. The spatial semantic hierarchy. *Artificial Intelligence*, 119:191–233, 2000.
- [7] Kevin Lynch. *The Image of the City*. MIT Press, Cambridge, MA, 1960.
- [8] J. Pailhous. Representation de l’espace urbain et cheminements. *Le Travail Humain*, 32:87–140, 239–270, 1969. (As cited by Chase (1982).).
- [9] E. Remolina and B. Kuipers. Boundary region relations. In *Cognitive Robotics, Papers from the 1998 AAAI Fall Symposium*, pages 117–124, Menlo Park, CA, 1998. AAAI Press. AAAI Technical Report FS-98-02.
- [10] A. W. Siegel and S. H. White. The development of spatial representations of large-scale environments. In H. W. Reese, editor, *Advances in Child Development and Behavior*, volume 10. Academic Press, New York, 1975.
- [11] Sabine Timpf, Gary S. Volta, David W. Pollock, and Max J. Egenhofer. A conceptual model of wayfinding using multiple levels of abstraction. In A. U. Frank, I. Campari, and U. Formentini, editors, *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*, volume 639 of *Lecture Notes in Computer Science*, pages 348–367. Springer-Verlag, Berlin, 1992.