REPRESENTING SPACE FOR PRACTICAL REASONING

Margaret M. Fleck

Dept. of Engineering Science, Oxford University
Parks Road, Oxford OX1 3PJ, England

ABSTRACT

This paper describes a new approach to representing space and time for practical reasoning. Unlike R^n, the new models can represent a bounded region of space using only finitely many cells, so they can be manipulated directly. Unlike Z^n, they have useful notions of function continuity and region connectedness. Finally, the topology of space is allowed to depend on the situation being represented, accounting for sharp changes in function values and lack of connectedness across object boundaries.

I INTRODUCTION

In their daily lives, people frequently reason about the shapes and arrangements of objects in space. This *practical reasoning* goes on at a variety of levels, from low-level visual processing, through identifying objects, up to reasoning about how an object could be manipulated. All these types of reasoning depend on representations of 2D and 3D space. Similarly, a representation for time is needed for reasoning about the relative ordering of events. There has been much discussion recently about the proper representations for time and space in fields including AI, computer vision and robotics, linguistics, and philosophy (van Benthem 1983, Dowty 1979, Allen and Hayes 1985, Hayes 1978a, Lee and Rosenfeld 1986).

A reasoner will need to construct for himself various models or descriptions of the world. It is useful to distinguish *symbolic descriptions*, such as "There is a desk against the wall," from *concrete models* of these descriptions, e.g. the sets of points of R^3 which comprise the desk and the wall. Concrete models can be inferred from sensory input, such as camera images, or produced by the reasoner "in his mind's eye" from symbolic descriptions. Symbolic descriptions can be derived from natural language input or from parsing concrete models.

Both symbolic descriptions and concrete models are useful in practical reasoning. Symbolic descriptions can consisely capture the relevant facts about a situation. This consiseness is important for remembering situations, identifying objects, describing situations which are too complicated to visualize all at once, and communicating in natural language. However, it is often simpler to use a concrete model for geometric reasoning tasks such as checking topological connectivity or measuring distances between features. Robot path planning can done using digitized representations of space (e.g. Brooks and Lozano-Pérez 1985). Finally, it is almost impossible to verify the consistency of a symbolic description except by exhibiting a concrete model which satisfies it.

In this paper, I will concentrate on the form of the concrete models. Standard models of time and space, such as R^n and Z^n, do not account for the way people do practical reasoning. Subsets of R^n must be manipulated symbolically because they typically contain infinite numbers of points. They cannot be directly stored by a reasoner. Secondly, it is difficult to represent two regions which are touching, because it is unclear which region contains the points along the common boundary of the two regions. The objects must overlap along the boundary, or the boundary points must belong to neither object, or else the boundary must be assigned arbitrarily to one of the objects (Allen 1984, Allen and Hayes 1985).

Models based on the integers avoid these problems with R^n, but at the cost of having no useful notion of function continuity or region connectedness. All functions from the integers are continuous and no subsets of the integers are connected. Most integer-based models handle only regular arrangements of points. A good concrete model for practical reasoning should use a finite density of samples, like integer-based models, but it should allow irregular arrangements of samples and it should provide notions of region connectedness and function smoothness like R^n. This paper develops such a model, derived from work by Poston (1971). My forthcoming thesis will supply technical details omitted in this paper, as well as algorithms for deriving concrete models from visual input, parsing concrete models into symbolic descriptions, and generating concrete models from symbolic descriptions.

II ADJACENCY SPACES

The first task in developing a finite-resolution concrete model is to set up the topological structure of empty space. Consider the set of cells in Figure 1. If the cells are packed against one another, they completely fill a section of space, the *physical realization* of the set of cells. If the cells do not fill the space, we can grow them so that they do fill it. We cannot, however, assume that these cells are arranged in a perfectly regular pattern. The cells in the human retina are not. However, even for irregularly spaced cells, the pattern of how the cells touch one another completely determines the topological structure of the region of space that they cover.

More specifically, the set of cells which touch at a point (such as A, B, C, and D), an edge (such as A and B), or a face will be called an *adjacency set*. In an N-dimensional situation, the *dimension* of an adjacency set is (N-M) if its cells touch along an M-dimensional face. For example,

The research described in this paper was done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology and at the Department of Engineering Science, Oxford University. The author is supported by the Fannie and John Hertz Foundation and by the Bell Laboratories Graduate Research Program for Women.

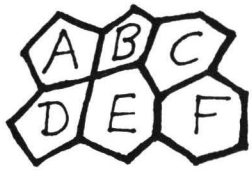

Figure 1. This set of cells has six adjacency sets of dimension 0: $\{A\}$, $\{B\}$, $\{C\}$, $\{D\}$, $\{E\}$, and $\{F\}$. The 1D adjacency sets are $\{A,B\}$, $\{A,D\}$, $\{B,E\}$, $\{D,E\}$, $\{B,C\}$, $\{C,F\}$, $\{C,E\}$, and $\{E,F\}$. The 2D adjacency sets are $\{A,B,D,E\}$, $\{B,C,E\}$, and $\{C,E,F\}$.

$\{A,B\}$ has dimension 1. For formal convenience, there is an adjacency set $\{X\}$ of dimension 0 for each cell X. A set of cells with adjacency sets and associated dimensions, will be called an *adjacency structure* (a modified form of Poston's local matroid structure). A pair of cells which belong to some common adjacency set (not necessarily one with two elements) are called *adjacent*. A *path from X to Y* or *connecting X and Y* is a finite ordered set of cells $X = W_0, ..., W_n = Y$, such that W_i and W_{i+1} are adjacent for every i $(0 \leq i < n)$. A set of cells $\mathcal{A}$ is *connected* if any two cells in $\mathcal{A}$ can be connected with a path.

In practical reasoning, it is necessary to represent objects and situations at more than one scale of resolution, i.e. using different densities of sample points. When the reasoner constructs a coarse-scale sampling from a finer-scale sampling, this can be done so that the coarse-scale cells are a subset of the fine-scale cells. In this case, the two adjacency structures can straightforwardly be related to one another. When two arbitrary samplings of the same object are created, e.g. when an object moves across the visual field, they are more easily related via a symbolic representation of the object.

Each situation used in practical reasoning has an intended dimension, e.g. a visual image is 2D. Subsets of the situation should have the same dimension as the original situation and the dimension of an object should not be altered by changes in the number of cells used to represent it. Thus, a ribbon one cell wide in a 2D situation is 2D. In such cases, the dimension of an adjacency set can be larger than or equal to the number of points in the set and the dimension of the adjacency set may be a range $[a,b]$ of non-negative integers, rather than just a single integer. For example, a one-cell subset of a 2D adjacency structure has a single adjacency set with dimensions $[0,2]$.

In order to be well-behaved, adjacency structures must meet some additional formal conditions. If $\mathcal{X}$ is an adjacency set with dimensions $[a,b]$ and the adjacency set $\mathcal{Y}$ with dimensions $[c,d]$ is a proper subset of $\mathcal{X}$, then $d < a$. If there is no adjacency set $\mathcal{W}$ distinct from $\mathcal{X}$ and $\mathcal{Y}$ with $\mathcal{Y} \subseteq \mathcal{W} \subseteq \mathcal{X}$, then $d + 1 = a$. Dimension 0 is restricted to singleton adjacency sets and dimension 1 to sets of no more than two elements. Since sampling should be at a finite resolution, a cell can belong to only finitely many adjacency sets and each adjacency set can contain only a finite number of cells. I also require that it be possible to embed in R^N the set of cells adjacent to any cell in an N-dimensional adjacency structure. Under these conditions, the adjacency structures are topologically equivalent to a subclass of regular cell complexes (Munkres 1984) which are also manifolds. Each cell of the adjacency space corresponds to a vertex in the cell complex. The cell complex is a deformation retract of the physical realization of the adjacency space. Thus the adjacency structure, unlike the pairwise adjacencies (Lee and Rosenfeld 1986), completely determines the topological structure of the space in the usual mathematical sense.

III FUNCTION SMOOTHNESS

Many algorithms in practical reasoning, such as interpretation of motion sequences, surface reconstruction, and reasoning about the behavior of physical objects, depend on the assumption that functions are "smooth." The idea behind function smoothness is that the value of a property should not change "too fast" as one moves through space or time. For R^n, there are a number of mathematical definitions corresponding to this intuitive concept, including continuity, smoothness, and bounded derivatives. The definition of function smoothness for adjacency spaces depends not only on the adjacency structure, but also on which cells *overlap*, i.e. sample overlapping patches of space. For example, in a CCD camera, the area sampled by a element overlaps areas sampled by elements which are several elements away, because of blurring and/or diffraction in the camera optics. Similar facts hold for the foveal area of the human visual system. This blurring before sampling reduces aliasing effects. In robot motion planning, it is essential that adjacent cells overlap, so that the entire area of space is covered by the cells and small objects cannot disappear from the representation.

Following Poston (1971), I call the overlap relation on a set of cells the *fuzzy*. An adjacency space with a fuzzy is called a *fuzzy space*. This relation is symmetric and reflexive, but not transitive. For a cell X in a fuzzy space, the *fuzzy neighborhood of X* is the set of cells overlapping X, including X itself. Each fuzzy neighborhood must be connected and have a finite number of cells. These definitions can be applied both to physical spaces and to abstract spaces, such as light intensities, temperatures, and distances. Two cells which overlap represent ranges of values which cannot reliably be distinguished. For example, I may have trouble distinguishing 10C and 15C, or 15C and 20C, but 10C and 20C are clearly different.

If $\mathcal{X}$ and $\mathcal{Y}$ are two fuzzy spaces, such as the visual field and grey-scale intensities, a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ is *smooth* if f(A) overlaps f(B) in $\mathcal{Y}$ whenever A and B overlap in $\mathcal{X}$. That is, in a region of smooth change, two overlapping cells in the visual field must have indistinguishable (overlapping) intensity values. If, for some overlapping A and B in X, f(A) and f(B) do not overlap, f has an *abrupt change in value* between A and B. This notion of smoothness depends on the fuzzies for the two spaces: if we extend the fuzzy on X so that more cells overlap, fewer functions are smooth. In reasoning about processes of change (Forbus 1984), it is necessary to consider the slope of a 1D function. For example, if the set of cells $\{X_i\}$ represents time and T the temperature of a room, we can define $DT(X_i)$ to be $T(X_{i-1}) - T(X_{i+1})$. We can now consider whether this new function DT is smooth, relative to appropriate fuzzies.

IV TOPOLOGICAL BOUNDARIES

Function smoothness and region connectedness are very important in practical reasoning. However, at a limited set of *natural boundaries* in a situation, such as at the edges of objects, functions can change abruptly and adjacent objects may not be, intuitively, connected. These locations of abrupt change are exactly the most interesting

parts of the situation for practical reasoning. For example, computer vision programs extract locations of abrupt changes in intensity and the reason using only descriptions of these boundaries. In analyzing processes such as heating liquids, properties and rates of change of properties can change abruptly when processes stop or when a substance undergoes a phase transition. These locations of abrupt change, along with summaries of behavior within regions of smooth change, can be used to predict the behavior of these systems (Forbus 1984).

Not only can functions have abrupt changes in value at natural boundaries, but the objects to either side of the boundary are not perceived as connected to one another. For example, it is necessary to distinguish whether two adjacent metal bars are physically connected in order to determine whether one bar will move if one pulls on the other. Pieces of metal do not merge on contact, although other substances (e.g. water) do. Connectedness can also be used to "limit causality" (Hayes 1978b). For example, if one can surround the situation of interest with boundaries across with nothing of interest is likely to flow, then reasoning can be limited to the region thus surrounded. Similarly, an event can only cause another event if the two are connected by a sequence of events.

Boundaries can be characterized by which types of functions change at them. For example, changes in lighting are important for reading but not for motion planning. Two pieces of metal can be physically but not electrically connected. However, in a particular situation, *boundaries relevant to different practical tasks tend to cluster.* In other words, the world, at any fixed scale of resolution, exhibits natural boundaries which are separated by large regions with no sharp changes. Connectivity boundaries and locations of abrupt changes in function values tend to coincide. This suggests that these natural boundaries are *topological* boundaries in situations. The usual explanation, that each function is discontinuous at a small number of places, fails to account for the clustering of boundaries and for the connectivity facts.

A program to detect boundaries in 2D camera images has been implemented, based on the model described in this paper. If there is a sharp change in intensity or in the slope of intensity between two adjacent cells, a boundary is marked between them. Each of the two cells is an *edge* of the region on its side of the boundary and a *border* of the region on the other side. An extended boundary creates two connected sets of edge/border cells. Boundaries should also be marked where there is reason to suspect lack of physical connection.

These pairwise boundaries define a new adjacency structure *relative to the boundaries (rttb)*, in which an adjacency set is removed if it contains two cells separated by a boundary. In an N-dimensional situation, if some adjacency set of dimension less than N is no longer a subset of any other adjacency set, its set of dimensions is extended to include N. For example, in Figure 1, if boundaries are added between C and F, and between E and F, the adjacency sets $\{C,F\}$, $\{E,F\}$, and $\{C,E,F\}$ are removed and $\{F\}$ is given dimensions [0,2]. If a boundary is added between A and E, B and D cease to be adjacent as a side-effect of removing the adjacency set $\{A,B,D,E\}$. Thus, adjacency structures avoid problems faced by representations based on pairwise adjacencies in adding boundaries where too many cells are adjacent (Lee and Rosenfeld 1986). From these new adjacencies, we can define paths (rttb) and region connectedness (rttb). The original fuzzy induces a fuzzy (rttb) in which each fuzzy neighborhood is restricted so as to be connected (rttb) and these fuzzies define function smoothness (rttb). Distances between cells, however, do not change when boundaries are added.

The natural boundaries in a situation can be used to define regions such as those people would use in describing the situation. For example, a local symmetry shape analysis (Brady and Asada 1984, Fleck 1985, 1986, Connell 1985) picks out sets of edge cells as the borders of elongated or round regions. These regions are typically connected relative to the boundaries which define them (Hayes 1978a). A region can be defined by boundaries which do not fully enclose it, e.g. the sides of an elongated region. A boundary can separate two parts of the same object, e.g. two adjacent fingers on the same hand. Thus, regions depend on locations of natural boundaries and not vice versa. Pairs of regions can be related using analogues of James Allen's (1983,1984) interval primitives. For example, a region $\mathcal{A}$ *touches* a region $\mathcal{B}$ if $\mathcal{A}$ and $\mathcal{B}$ are disjoint and there are two adjacent cells a and b such that a is an edge of $\mathcal{A}$ and a border of $\mathcal{B}$, and b is an edge of $\mathcal{B}$ and a border of $\mathcal{A}$. This relation corresponds to Allen's *meets.* Analogues of his other primitives can be defined similarly, using also an order for 1D spaces.

In doing practical reasoning, a reasoner must choose which boundaries are active during a particular piece of topological reasoning. Certain boundaries may not be relevant to the task at hand. It may be necessary to consider several different models, e.g. in deciding which of several electrical connections is broken. There may be more than one way to parse a situation. For example, the region occupied by a marble inside a cup can be seen as overlapping the interior of the cup, or as disjoint from the free space inside the cup. In other words, *the topology of space can be manipulated dynamically during reasoning.*

REFERENCES

Allen, James F., "Maintaining knowledge about temporal intervals," CACM 26:11 (1983) 832-843.

_____, 'Towards a General Model of Action and Time," *Artificial Intelligence* 23:2 (1984).

Allen, James F. and Patrick J. Hayes, "A Common-Sense Theory of Time," IJCAI (1985) 528-531.

van Benthem, J.F.A.K., *The Logic of Time*, D. Reidel, Dordrecht (1983).

Brady, J. Michael and Haruo Asada, "Smoothed Local Symmetries and Their Implementation," IJRR 3:3 (1984) 36-61.

Brooks, Rodney A. and Tomás Lozano-Pérez, "A Subdivision Algorithm in Configuration Space for Findpath with Rotation," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol SMC-15 (1985).

Connell, Jonathan H., "Learning Shape Descriptions," MIT AI Lab. TR–853 (1985).

Dowty, David R., *Word Meaning and Montague Grammar*, D. Reidel, Dordrecht (1979).

Fleck, Margaret M., "Local Rotational Symmetries," MIT AI Lab. TR–852 (1985).

_____, "Local Rotational Symmetries," CVPR (1986).

Forbus, Kenneth D., "Qualitative Process Theory," MIT AI Lab., TR–789 (1984).

Hayes, Patrick J., "Naive Physics 1: Ontology for Liquids," unpubl. (1978a).

_____, "The Naive Physics Manifesto," unpubl. (1978b).

Lee, Chung-Nim and Azriel Rosenfeld, "Connectivity Issues in 2D and 3D Images," CVPR (1986) 278-285.

Munkres, James R., *Elements of Algebraic Topology*, Addison-Wesley (1984).

Poston, Tim, "Fuzzy Geometry", Ph.D. thesis, Univ. of Warwick (1971).