

# The topology of boundaries

Margaret M. Fleck \*

*Department of Computer Science, University of Iowa, Iowa City, IA 52242, USA*

Received November 1990; revised May 1993

---

## Abstract

High-level representations used in reasoning distinguish a special set of boundary locations, at which function values can change abruptly and across which adjacent regions may not be connected. Standard models of space and time, based on segmenting  $\mathbb{R}^n$ , do not allow these possibilities because they have the wrong topological structure at boundaries. This mismatch has made it difficult to develop formal mathematical models for high-level reasoning algorithms. This paper shows how to modify an  $\mathbb{R}^n$  model so as to have an appropriate topological structure. It then illustrates how the new models support standard reasoning algorithms, provide simple models for previously difficult situations, and suggest interesting new analyses based on change or non-change in scene topology.

---

## 1. Introduction

Most work in artificial intelligence (including not only “core” AI, but also peripheral areas such as vision, robotics, and natural language) involves reasoning about arrangements of objects in space or arrangements of events in time. The most popular model for time is the real line ( $\mathbb{R}$ ) and the most popular model for space is real Euclidean space  $\mathbb{R}^n$ . However, these models have the wrong topological structure at object, event, or region boundaries. In the symbolic models used in reasoning, adjacent regions need not be connected to one another and otherwise continuous functions can display abrupt changes in value across region boundaries. Standard models of regions based on  $\mathbb{R}^n$  do not allow either of these possibilities.

Practical implementations in artificial intelligence can typically work around this problem. An algorithm may not make extensive or sophisticated use of topological

---

\* E-mail: [mfleck@cs.uiowa.edu](mailto:mfleck@cs.uiowa.edu).

concepts. Or the author may adopt an model which works in his particular application but may not work for other applications, e.g. a robot motion planner may assume that there are only two types of regions in the scene, obstacles and free space. However, the topological mismatch between  $\mathbb{R}^n$  and what is desired for artificial intelligence creates a pattern of difficulties when researchers come to build theoretical, mathematical analyses of their algorithms. Thus, researchers waste time battling with inappropriate theoretical models, time that could be better spent designing practical algorithms.

Alternatively, some authors [2, 3, 19, 20, 22, 23, 36] base their analyses on a set of symbolic axioms without a specific concrete model. This is mathematically risky. It is difficult to prove an axiom set consistent except by exhibiting (at least) one concrete model. It is difficult to derive possible models from axiom sets (cf. [52]). And it is almost impossible to extend an axiom set without a clear picture of the object that these axioms are intended to describe.

In this paper, I will show how to modify an  $\mathbb{R}^n$  model so it has the correct topology. As the issues involved are relatively subtle, I will proceed straight into a review of basic definitions from topology (Section 2). I will then discuss what topological structure should be assigned within regions and across boundaries, and why the structure of  $\mathbb{R}^n$  is not correct (Section 3). I will describe the new model in Section 4 and discuss alternative models in Section 5. Section 6 works through some extended examples from the literature, to show how they would be handled in the new model.

## 2. What is topology?

A satisfactory model of space and time must define a suitable topological structure, because the topological structure determines two concepts of great importance in practical reasoning:

- which functions are continuous, and
- which regions and paths are connected.

In this section, I will review the definitions of basic topological concepts.<sup>1</sup> See any topology text (e.g. [38, 39]) for additional details.

### 2.1. Open sets and connectivity

Readers will be familiar with terms such as “connected region”, “open set” and “continuous function” from courses in calculus or analysis. Topology generalizes these definitions so that they can be applied to a wider range of models for space and time. In particular, readers unfamiliar with topology may believe that topological notions are closely tied to a metric (e.g. the usual distance function for points in space) and/or an ordering (e.g. the usual linear order on points in time). Although this is true for some familiar examples, it is not true in general.

---

<sup>1</sup> I will resist the temptation to describe the fascinating range of derived topological concepts such as homology groups, Euler characteristics, and knot polynomials, because they would distract from the main point of this paper.

Two pieces of information are required to specify a particular topological space:

- what points does the space contain, and
- what sets of points are “open”.

So, for example, in the usual topology on the real line  $\mathbb{R}$ , the open sets are the familiar open intervals, e.g.  $(1, 2)$ ,  $(-57, 3)$ ,  $(202, \infty)$ , and unions of such intervals. The real line can also be given other topologies, see [38]. Open sets can also be specified indirectly, e.g. via order relations, by cell complex constructions [39], or by taking a subset of some established space such as  $\mathbb{R}^n$ . In particular,

**Definition 1.** Let  $A$  be a subset of a topological space  $B$ . In the *subspace topology* on  $A$ , a subset  $S$  of  $A$  is open if and only if  $S = X \cap A$ , where  $X$  is some open subset of  $B$ .

So, the subset  $A = [0, 1) \cup (1, 5]$  of  $\mathbb{R}$  is typically given the subset topology. The open sets of  $A$  are the intersections of  $A$  with open sets of  $\mathbb{R}$ .

To be a well-formed topological structure, the open sets must satisfy the following conditions:

- the null set and the entire space are both open,
- the union of any (finite or infinite) collection of open sets is open, and
- the intersection of any finite collection of open sets is open.

A set is said to be *closed*, in some specified topology, if its set complement is open. It is possible for a set to be both open and closed (e.g. the whole of  $\mathbb{R}$ ) or neither (e.g. a half-open interval of  $\mathbb{R}$ , in the usual topology). The *closure*  $\bar{A}$  of a set  $A$  is the smallest closed set containing  $A$ .

A region in a topological space is said to be *connected* if it is not the union of two disjoint open sets. A region  $A$  is said to be *path-connected* if any two points  $x$  and  $y$  in  $A$  can be connected by a continuous path lying entirely in  $A$ . Formally, there must be a continuous function  $p : [0, 1] \rightarrow A$ , such that  $p(0) = x$  and  $p(1) = y$ . Informal models of space do not distinguish these two notions. Fortunately, this does not matter: connectedness and path-connectedness will be equivalent for the spaces discussed in this paper.

## 2.2. Continuous functions

The topological definition of continuity is a generalization of the familiar  $\varepsilon$ - $\delta$  definition from calculus:

**Definition 2.** A function  $f : X \rightarrow Y$  is *continuous* if and only if  $f^{-1}(U)$  is open for any open set  $U \subseteq Y$ .

In order to apply this definition, we must specify a topology for both the domain and range of the function. Within a connected region, continuous functions behave as you expect them to. If, however, the domain consists of two non-connected components (e.g. the domain is  $[0, 1) \cup (1, 5]$ ), the value of a continuous function can jump abruptly as one passes between the two components.

Continuous functions are also used to match representations topologically, e.g. determine whether one object can be obtained from another via stretching, bending, and deformation. Formally, we define:

**Definition 3.** A function  $f : X \rightarrow Y$  is a *homeomorphism* if  $f$  is continuous and has a continuous inverse.  $X$  and  $Y$  are said to be *homeomorphic* if there exists a homeomorphism from  $X$  onto  $Y$ .

To match two complex situations, as in stereo matching (cf. [14]), a stronger constraint can be useful [28, 45, 46]:

**Definition 4.** Two subsets  $X$  and  $Y$  of  $\mathbb{R}^n$  are *isotopic* if there is a continuous family of homeomorphisms  $f_i : X \rightarrow \mathbb{R}^n$ ,  $i \in [0, 1]$ , such that  $f_0(X) = X$  and  $f_1(X) = Y$ .

Two homeomorphic spaces have identical topological properties. Two isotopic spaces are, in addition, embedded “the same way” in  $\mathbb{R}^n$ , so that one can be continuously deformed into the other.<sup>2</sup>

### 3. The desired topological structure

Reasoning algorithms, both high-level and low-level, distinguish a particular set of locations as “boundaries”. Boundaries demarcate objects in 3D space, distinctive regions in 2D camera images, and events, processes, or states in 1D representations of time. Although algorithms for locating boundaries are still flawed and there is no generally accepted formal definition of a boundary, there is considerable consensus among researchers as to where boundaries should be placed in most practical situations. Many artificial intelligence applications depend on having a suitable topological structure for space or time. However, as we shall see in this section, the structure required in the interior of regions seems to differ greatly from that required across boundaries. By contrast,  $\mathbb{R}^n$  has the same structure at all locations.

#### 3.1. Boundaries

Before discussing the topological facts, we must first draw a clear distinction between boundaries and the objects, regions, or intervals they delimit. First, two parts of the same object can fail to be locally connected. The object might consist of more than one connected component. Or, alternatively, a connected region may double back so as to touch itself, as illustrated in Fig. 1. When you touch your thumb to your finger, they do not typically fuse on contact. Similarly, the boundaries of 3D objects and 2D regions need not form closed curves, as illustrated in Fig. 2.

Second, two distinct regions can be connected to one another. For example, people distinguish a finger from the hand connected to it, despite the lack of any clear boundary

<sup>2</sup> I am abusing notation slightly to make the presentation accessible to readers unfamiliar with topology. Strictly speaking, an isotopy relates two *embeddings* of a space  $X$  into another space such as  $\mathbb{R}^n$ .

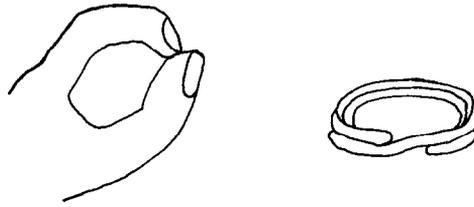


Fig. 1. Examples of a region touching itself, with an intervening boundary: a finger and thumb, a split ring.



Fig. 2. The boundary between two adjacent fingers ends abruptly, both in 3D and in its 2D projection.

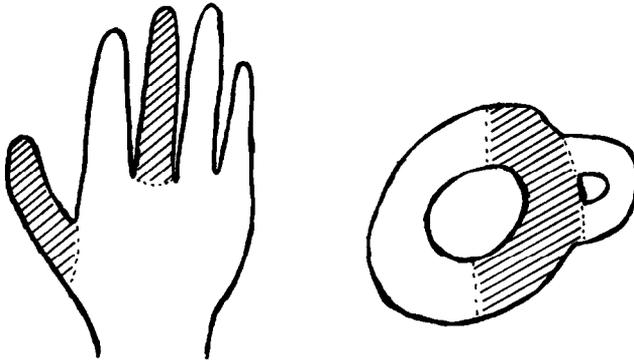


Fig. 3. A region need not be completely delimited by boundaries: (left) a finger and a thumb are parts of a connected hand, (right) a region of complex shape may be subdivided for the convenience of internal algorithms.

separating them (Fig. 3). Similarly, a description such as Sentence 1 need not refer to the maximal connected interval over which its predicate (“John is in town”) holds, i.e. the verb may refer to an interval which does not extend all the way to the boundaries at which the state starts and ends. Finally, reasoning algorithm may subdivide a connected region for convenience of processing. For example, a 2D region may be represented by a set of cells in a digitized map. Or, a surface of complex topology may be decomposed into subsets of simple topology, e.g. a body and “handles” [45]. For all these reasons, it is necessary to regard boundaries and regions as distinct concepts, only loosely related.

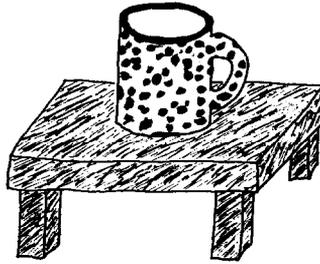


Fig. 4. Properties such as texture, color, and material composition vary continuously within objects, but can change abruptly at object boundaries. So, here we see a transition between a cup (spotted, white, ceramic) and a table (stripy, brown, wood).

**Sentence 1.** *John was in town from 2:00 to 3:00.*

### 3.2. Behavior of functions

Many reasoning algorithms assume that, during the course of any specified process or event, properties change continuously over time. Sharp changes are allowed only across a limited set of event boundaries. For example, Sentence 2 describes two constant states, separated by an abrupt change. Sentences 3-4 describe abrupt changes between two continuous processes. Similarly, in 2D images or 3D scenes, properties such as texture, color, and material composition vary smoothly within regions but change abruptly across region boundaries (Fig. 4).

**Sentence 2.** *Boris and Martina were married yesterday.*

**Sentence 3.** *At the boundary of the prism, the light ray's direction changed from  $(x, y)$  to  $(x', y')$  due to refraction.*

**Sentence 4.** *The temperature of the water rose to 100 degrees, and then it started to boil.*

This contrast between continuous and abrupt change is used throughout artificial intelligence. Computer vision algorithms assume that image values [7, 35], texture distributions [16, 17], motion vectors, or stereo disparities [14, 43] vary continuously within regions.<sup>3</sup> These algorithms hypothesize boundaries where values change abruptly. Sharp spikes in property values, e.g. due to dark cracks or phase discontinuities in texture (Fig. 5), can also indicate boundaries [16]. Descriptions of boundaries and regions, typically more compact than the original image (cf. [42]), are then passed to later vision algorithms such as shape analysis and object identification.

Qualitative reasoning algorithms [19, 31, 33, 34, 54, 56] use continuity within regions to summarize behavior over extended intervals of time into succinct representations

<sup>3</sup> In fact, these algorithms make the stronger assumption that values have bounded derivatives within regions.

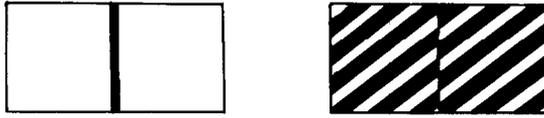


Fig. 5. A boundary can separate two regions with identical properties: (left) a dark crack, (right) a phase change.

(processes). They then identify a sparse set of locations at which the qualitative state changes abruptly. At such a location, the reasoner must do a more detailed examination to determine which process will occur over the next interval of time. In general, some quantities may change continuously across the boundary and some change abruptly. For example, when a ball hits a wall, the position of the ball changes continuously but its velocity changes abruptly.

It can be argued that many “abrupt” changes are gradual if examined at high enough resolution [56]. This may be appropriate for some applications. However, it is not clear that it holds for all reasonable models of low-level physics. Further, and more importantly, the models used in reasoning are abstractions from actual physical reality and allowing abrupt changes often results in simpler models. There is a strong community in planning and linguistic semantics (e.g. [2, 10, 21, 30, 58]) who believe that certain events should be modelled as abrupt changes in state. Gradual change explanations are particularly unconvincing for examples involving human-defined abstract properties: in Sentence 2 there is no such state as “half-way married”.

### 3.3. Connectivity

Region connectivity is also important in reasoning about physical systems. First, analysis of mechanical systems requires a representation for the material connectivity of objects. Material connectivity is an abstraction that allows a reasoner to reason collectively about the behavior of many small patches of stuff by grouping them together into one “object”. In standard freshman physics, each object moves as a rigid unit, with one collective position, orientation, and velocity. Even for non-rigid objects, material connectivity implies transmission of forces along the object: if we pull on one end of an ethernet cable, eventually either the cable will become taut or the other end will move.

Material connectivity, however, fails across boundaries. For example, a telephone may rest on a desk without being connected to it: if we pick up the phone, the table will not be pulled along with it. It may be tempting to see lack of connectivity as a side-effect of differences in material properties. This may be true in some low-level representation of the molecular structure. However, in high-level more abstract models, it seems necessary to allow two regions of identical properties to come into contact without fusing together (Fig. 6).

Many objects extracted by people for high-level reasoning seem either to be connected or to be composed of a small number of connected components. Because of this, computer vision systems often compute shape properties only for connected image regions or connected curves (e.g. [6]), because this seems to increase the chance

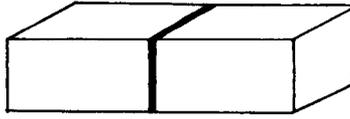


Fig. 6. Two metal blocks, with identical properties, can come into contact without fusing together.

that the whole region or boundary belongs to one real-world object. Similarly, most reasoning systems use only connected intervals to model events, processes, or states. Certain linguistic constraints on overlaps between actions can be modelled using a distinction between connected intervals and boundaries (Section 6.3). Short disruptions in an action (cf. [57]), or clutter obscuring a region in an image, do not necessarily force one to drop connectedness assumptions, because they might be eliminated by low-level processing.

It is also necessary to represent the connectivity of empty space. Objects move [26, 59] and fluids flow [19, 23, 31, 33, 34] only through connected regions of empty space and are halted by (insufficiently porous) objects. Some path planning algorithms [59] use concepts closely tied to connectivity, e.g. *retractions* and *homotopies*. Abstract flows such as heat, light, and electricity also flow through connected paths, where connectivity is defined by abstract barriers such as changes in conductivity, refractive index, density, or heat capacity. At these barriers, they may stop, lose amplitude, or change direction. Well-chosen barriers can allow reasoning to consider only a limited subsection of the world without excessive interference from outside forces [23].

Connectivity information can also be used to compensate for errors in shape information. For example, suppose we want to represent a floppy sack. If we represent the possible locations of the interior and exterior surfaces of the sack, the variability in shape leaves open the possibility that the exterior surface lies inside the interior surface. To prevent this, the representation must contain explicit topological information, such as “the sack plus its mouth is bounded by two sphere-shaped boundaries, one inside the other, separating the sack’s inside from its outside”. This clarifies the relationship between the two sides of the sack and explains how the sack can contain objects. For similar reasons, topological information is sometimes used to augment algebraic representations of surfaces, because numerical problems in the algebraic models may cause apparent gaps between surfaces that should be connected.

### 3.4. What happens in $\mathbb{R}^n$ ?

Now, suppose that we try to build a formal model for a set of objects in space, or a set of events over time. The traditional approach is to use  $\mathbb{R}^n$  as the model of space or time. Each object or event is modelled as a subset of  $\mathbb{R}^n$ . An expectation implicit in most traditional modelling is that

the set of regions in an  $n$ -dimensional model must completely cover  $\mathbb{R}^n$ .

In computer vision jargon, we are building a “segmentation” of the scene. In philosophy jargon, we are not allowing “truth gaps”.

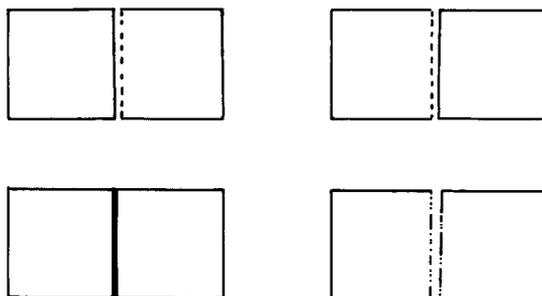


Fig. 7. Several possibilities for assigning boundary points to two regions in  $\mathbb{R}^n$ . Points can go with lefthand region (top left), go with righthand region (top right), belong jointly to both (bottom left), or be divided between them (bottom right).

It should be clear from the above discussion that  $\mathbb{R}^n$ , with its usual topology, is a good model for the interior of regions. At boundaries, however, there is a problem. Because we have assumed that the regions cover all of space or time, there can be no gap between adjacent regions. Fig. 7 shows some possible shapes for the abutting edges of two adjacent regions. No matter which of these models is used, the two regions (jointly) occupy a connected region of  $\mathbb{R}^n$ , so they are connected topologically. Furthermore, because they are connected, continuous functions cannot have abrupt jumps in value between the two regions. This problem has not been noted before: previous authors (e.g. [2, 8, 9, 21, 22, 30, 41]) have concentrated instead on the relative merits of the different possibilities illustrated in Fig. 7.

#### 4. The new model

So, an exhaustive segmentation of  $\mathbb{R}^n$  into regions, together with the standard definitions of connectedness and continuity, does not correctly model topological properties at boundaries. To model the effect of boundaries on the underlying continuous space, we must “cut” the topological structure of  $\mathbb{R}^n$  along the boundaries. The simplest way to do this is to delete all points in the boundaries from  $\mathbb{R}^n$  (Fig. 8). That is, these points no longer form part of any region and they are no longer in the domain or range of any function.<sup>4</sup> The rest of this section will discuss the details of this new model and explain how it gives scenes the correct topological properties.

##### 4.1. The boundaries

In order to create well-formed regions, it is necessary to impose some constraints on the set of boundaries. However, these constraints do not need to be as strong as those in some previous models. In particular, I will require that

<sup>4</sup> It may be convenient to reference these points in data structures used in computer implementations. This is legitimate so long as the overall implementation emulates the behaviors described below.

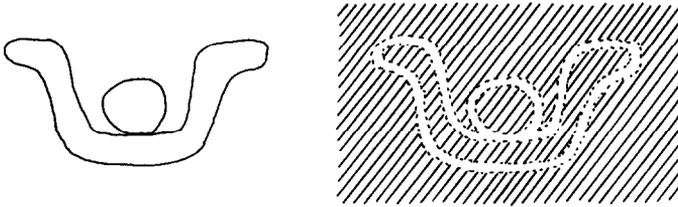


Fig. 8. A set of boundaries (left) and the region topology they induce (right).

**Constraint 1.** *The boundaries are a (topologically) closed set of points.*

Thus, if a line segment is in the boundaries, its vertices must also be in the boundaries; if a 2D region is in the boundaries, its edges must also be in the boundaries. This forces the regions to be a manifold of the same dimension as the original space.<sup>5</sup> In particular, a region cannot be only one point wide (infinitely thin), either for its entire extent or at an isolated cutpoint.

In most applications, it is probably best to require that

**Constraint 2.** *The boundaries in any bounded subset of  $\mathbb{R}^n$  are isotopic to the union of a finite number of semi-algebraic sets.*

Most readers are probably familiar with semi-algebraic sets: they are sets defined by the intersection of polynomial constraints (for precise definitions, see [59]) and are used extensively in robotic and visual modelling. The extension to boundaries which are only isotopic to semi-algebraic sets allows one to make boundaries smooth (derivatives of all orders exist) by rounding sharp corners. This is a technical condition to prevent one from constructing a variety of exceedingly nasty topological spaces (e.g. letting the boundaries be a Cantor set or the Alexander Horned Sphere, cf. [46]). Such situations cannot occur in real life and their existence may complicate theoretical proofs.

Boundaries are not, however, required to be thin, i.e. of dimension  $\leq (n-1)$  in a scene of dimension  $n$ . For example, the boundaries in a 2D scene can contain some 2D regions as well as 1D curves and isolated points. Such thick boundaries are useful in low-level vision for representing blurred transitions and complex intersections, for compatibility with traditional on-cell representations of boundaries [15], and as a technical device for implementing isotopies [14] (see Section 6.5). Furthermore, such dimensionality restrictions would not significantly simplify the theoretical model.

This model extends easily to digitized representations, e.g. digitized camera images. In a 2D rectangular digitization, each digitized location is identified with a rectangular subset of the plane (a cell). The cells for adjacent locations share a common edge. Boundaries can then contain entire 2D cells, edges between cells, and/or vertices where

<sup>5</sup> Any open subset of a manifold is a manifold.

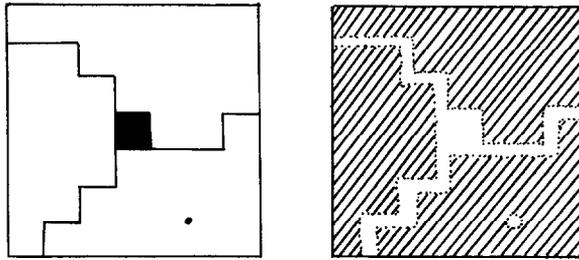


Fig. 9. In a digitized model, the boundaries must be made up of cells, edges, and vertices from the digitization. However, the region topology (right) is derived exactly as in non-digitized models.

four cells meet (Fig. 9).<sup>6</sup> This has been used in implementations of several computer vision algorithms [13-15, 17] and Section 6.5.

#### 4.2. The topology

The new models are given the subspace topology, as defined in Section 2.1. Thus, in the interior of regions, the new models look just like  $\mathbb{R}^n$  and formal definitions of continuity and connectedness work in the familiar way. However, two regions separated by a boundary are no longer connected to one another, because of the line of missing points between them. Because adjacent regions are not connected, the values of a continuous function can jump abruptly between the two regions. This is exactly the behavior specified in Section 3.

In this model, the edges of regions look "open".<sup>7</sup> That is, neighborhoods on the edges of a region are homeomorphic to  $\mathbb{R}^n$ . Some authors [23] have the intuition that the edges of a region should, instead, look "closed", i.e. be homeomorphic to a closed half-space of  $\mathbb{R}^n$ . Models similar to the new one, but with "closed" edges, can be built using a slightly more complicated mathematical construction [12, 13]. However, I have not been able to find any convincing practical or mathematical reason why a reasoning algorithm should need "closed" edges. Furthermore, the "closed" models do not allow two regions to touch (be infinitely close but not fused), whereas the "open" models support a three-way distinction between nearby regions, regions in contact, and fused regions (Sections 4.5 and 6.4).

Since the model of a scene is always a subspace of  $\mathbb{R}^n$ , we can transform one model continuously into another model via isotopy (see Section 2.2). An isotopy can move a boundary continuously, shrink a region so as to widen the adjacent boundary, or expand a region so as to narrow the adjacent boundary. It cannot change the handedness of a region that differs topologically from its mirror image, nor can it remove a region from inside another, ring-shaped, region. In reasoning about motion and collision of objects in space (as in freshman physics), requiring an isotopy between models at successive times would concisely capture the fact that each object moves continuously, without

<sup>6</sup> The extension to other dimensions and non-regular digitizations is beyond the scope of this paper. See [12-15].

<sup>7</sup> The regions are both open and closed topologically, so I am not using this word in its technical sense.

breaking up, developing cracks, or merging with other objects. This constraint has also been used in stereo matching: see Section 6.5.

#### 4.3. Functions

In the new model, abrupt changes in the value of a function are represented by adding a boundary to the domain of the function. Since underlying boundaries seem to account for all cases in which functions in reasoning applications behave discontinuously, I propose the following constraint on representations used by reasoning algorithms:

**Constraint 3.** *All functions are continuous.*

In practical terms, this means that if the reasoning algorithm observes apparently discontinuous behavior from one of its functions, it is required to restructure its representation.

Restructuring may be accomplished in several ways. The reasoner may hypothesize a boundary in the domain of the function to account for the function's behavior. Or, the change can be modelled as gradual, by changing thresholds on what constitutes "abrupt" or adding intermediate values to the function's range. A higher-resolution representation of the function may be obtained, in which the change may appear more gradual. The abrupt change may be attributed to sensor noise and removed by changing (smoothing) the values of the function. Finally, the reasoner may decide that this entire function is not relevant to the task at hand and remove it from the model.

The appropriate method of restructuring depends on the application. For example, the apparently simple event of turning on a light can reasonably be modelled as either an abrupt change or a continuous one. Unlike some previous researchers [56], I believe that both models are appropriate, in different contexts, and that a reasoner may need to maintain several qualitatively distinct models of a situation (e.g. a gas may be modelled as a set of molecules or as a fluid following the ideal gas law). However, each model must be internally consistent. We will see (Sections 5.1, 6.1, and 6.3) that internal consistency severely constrains what models can be constructed.

#### 4.4. Region borders

Many algorithms in computer vision track boundaries, so as to convert maps of boundary locations into connected curves for later shape analysis. Similarly, reasoning algorithms often refer to the locations at which sharp changes occur. Since the boundary points are deleted in the new models, attention naturally shifts to the *borders* of regions, i.e. thin strips of region adjacent to the region boundaries (Fig. 10). Specifically, we define

**Definition 5.** An open set  $E$  of points is a border of a set of regions  $R$  if  $R - \bar{E}$  is isotopic to  $R$ .

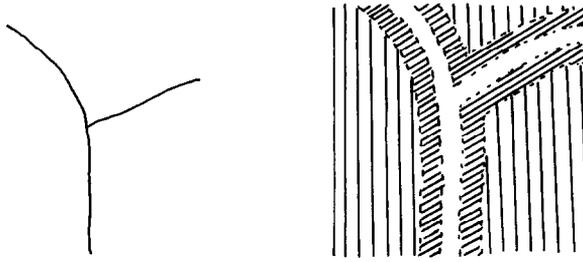


Fig. 10. Because boundaries (left) are deleted in the new models, attention shifts to region borders (right, diagonally shaded), thin strips of region adjacent to the boundaries.

In other words, deleting the border points from the regions shrinks the regions slightly but does not change their topological shape.

Within the topological constraints, borders are free to come in a variety of widths. In fact, purely topological conditions can't prevent a border from having strange protrusions that don't reflect the boundary shape. However, the intent is that borders be made "narrow" and approximately "uniform" in width, using definitions of "narrow" and "uniform" that are appropriate to the individual application. Low-level sensory data, as in computer vision, is spatially sampled (digitized). For such data, it is convenient to use borders whose width is one sample location or one observation. In high-level representations, the borders might constitute "the smallest interval over which the change of state has clearly taken place" [10, p. 141]. For many purposes, it is not necessary to specify the border width exactly.

Traditional algorithms which refer to boundaries can typically be converted to use borders instead. In implementation, a section of border can be represented as a section of boundary plus an inward/outward distinction, or as a piece of region plus the direction in which the adjacent boundary lies. Tracking borders in 2D images is actually easier than tracking boundaries [15] (cf. also [40]). A digitized border segment always has a unique successor, whereas boundary tracking is poorly defined at junctions and when boundaries are thick. High-level vision algorithms, such as shape analysis and object identification, can accept borders in place of boundaries with only minor revisions. In some cases [18], they become simpler. The code used in our introductory vision laboratories is entirely border-based. Section 6 shows how some typical types of qualitative reasoning might be rewritten using borders.

#### 4.5. *Touching*

Many authors in AI and the allied fields [2, 3, 10, 21-23] consider it important to allow two intervals in time or two patches of space to "touch" one another. In order for two sets of points to "touch", three key properties must be satisfied:

- the two sets do not overlap,
- there is nothing between them, and
- the distance between them is insignificant, preferably zero.

So, in Sentence 5, the second event occurs “right after” the first and in Sentence 3 (above), the state change occurs “in an instant”.

**Sentence 5.** *The robot moved forward two inches, then grabbed the block.*

Allowing regions to touch makes it easier to model the physics of objects in space. For example, if we put a cup on a table, gravity will maintain it at zero distance from the tabletop. If we push something towards a rigidly fixed table, e.g. a paint scraper, additional forces will be generated as the two come into contact. Such effects have a great impact on low-level robot control algorithms, which can go unstable when asked to plan motions involving contact. For high-level planning algorithms, objects in contact have fewer degrees of freedom than objects surrounded by empty space.

The key observation, which seems not to have been made explicit by previous authors, is that the motivating examples require touching only at region or event boundaries. Some existing axiom systems (such as [2, 3]) allow regions of space or intervals of time to touch at any location in the model. However, no one has produced a convincing example of touching at a location in the middle of a homogeneous region or action.

The new models inherit their metric (distance function) and differential structure from  $\mathbb{R}^n$ , just as they inherit their topological structure. Thus, if two adjacent regions are separated by a thin boundary, they are zero distance apart.<sup>8</sup> As the regions do not overlap and there is nothing between them, the adjacent regions meet the above specifications for touching. If thick boundaries are used, adjacent regions will be some nonzero distance apart. However, we can still identify adjacent regions formally:

**Definition 6.** Two regions  $A$  and  $B$  are adjacent if there are points  $a \in A$  and  $b \in B$  such that all points on the line segment connecting  $a$  and  $b$  (not including any deleted boundary points) belong to  $A$  or  $B$ .

Perhaps adding an application-dependent constraint on the distance between  $A$  and  $B$ .

#### 4.6. Summary

Thus, the new model of boundaries and regions faithfully models the desired topological properties described in Section 3. It allows regions to touch without overlap, required for certain applications in artificial intelligence. It also allows isotopy to be used as a constraint on relations between two models (e.g. changes over time). And it allows us to insist that all functions used in reasoning models be continuous.

### 5. Some approaches that don't work

The new model for boundaries departs significantly from current practice in artificial intelligence, because it suggests that, in our intuitive models, space itself changes as

<sup>8</sup>Using the standard definition of the distance between two sets  $A$  and  $B$ :  $\inf\{d(x,y) \mid x \in A, y \in B\}$ , where  $d(x,y)$  is the distance between two points.

the scene changes. One might ask whether a less drastic solution could be used. The recent literature suggests several alternative approaches. Although they do not work, it is instructive to see how they differ from the new model proposed in Section 4 and why they fail to have the desired properties.

### 5.1. Selective violations of continuity

The violations of function continuity have been the most conspicuous failure of the  $\mathbb{R}^n$  model. We could handle them (as in [56]) by allowing functions to be only piecewise continuous, i.e. continuous except at a limited set of locations. For any function individually, it is impossible to distinguish this analysis from one in which functions are continuous and space is disconnected. Differences appear when multiple functions have the same domain, e.g. color, depth, and texture data for a common 2D image, and when region connectedness is considered.

In reasoning and computer vision applications, it is common for functions with the same domain to exhibit sharp changes at roughly the same locations. For example, between a coffee cup and the table below it, there may be changes in color, intensity, texture, material composition, and (if the cup has just been put there) temperature. In the new model, a boundary in the domain licenses (but does not require) abrupt changes in the values of *all* functions at the boundary. Thus, it produces a compact representation for the cup and table, and suggests the (plausible) prediction that other properties might also change at their boundary. In the piecewise continuity model, abrupt changes in one function are independent of abrupt changes in other functions and discontinuities would most naturally be stored as properties of the individual functions. This yields a less compact representation of the cup and table, no reason to expect different functions to change at a common location, and no prediction about the behavior of other functions. One might, of course, incorporate some additional co-occurrence constraint into the piecewise continuity model but this seems *ad hoc*.

A second limitation of the piecewise continuity model is that it says nothing about region connectivity. In the cup and table example, the cup also fails to be materially connected to the table. The break in connectivity occurs at the same location as the sharp changes in function values. In raw image data, sharp changes in function values can occur where there are no breaks in material connectivity, e.g. shadows and surface markings. However, it is an accepted goal of computer vision algorithms to identify and remove such boundaries, giving only “real” object boundaries to later object identification and reasoning.<sup>9</sup> In the new model of boundaries, continuity and connectedness must fail simultaneously, so it predicts situations like the cup and table and represents them compactly.

The piecewise continuity model requires a distinct—and currently non-existent—model of region connectedness and some mechanism for relating breaks in connectivity to discontinuities in functions. Some high-level models may contain a primitive CONNECTED, used to specify whether two regions are connected. However, this is too crude a mechanism to describe the full range of different connectivities, e.g. two regions con-

---

<sup>9</sup> Current algorithms are, of course, not particularly successful at doing this, but that is beside the point.

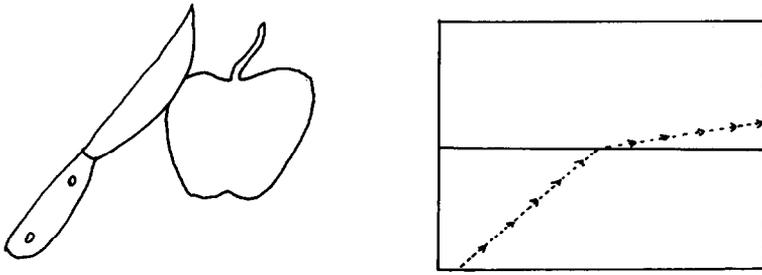


Fig. 11. There may be more than two types of regions in a scene. There may be two types of objects e.g. edible and non-edible (left), plus the background. Two regions of free space may differ in properties, e.g. a change in refractive index creates a boundary at which light rays change direction (right), but remain distinct from regions filled by objects.

nected in two distinct places or along a ring, a region touching itself along an internal boundary.

### 5.2. Same points as $\mathbb{R}^n$

Suppose that we keep the same points as  $\mathbb{R}^n$  and just alter the open sets. When two regions are adjacent, the points in their common boundary must be divided between them. The points must belong to one of the two regions, otherwise we create an extra region with strange properties: most processing ignores its existence and functions have undefined values on its points (a “truth gap”) [2, 10, 21, 30]. If there is a change in property values between the two regions (e.g. they are different colors), a point cannot belong to both regions because this creates two conflicting values. Furthermore, the regions then overlap.<sup>10</sup>

Thus, there are two basic ways to allocate the boundary points (see Fig. 7). We can give all boundary points to one region (the *asymmetrical model*). Or, we can divide them in some way between the two regions, e.g. 50% to each region, chosen randomly (the *random model*). We then remove, from our inventory of open sets, all sets that cross a boundary. In both cases, the resulting space is not a subspace of  $\mathbb{R}^n$ , a problem since this prevents us from using isotopy to model continuous change (Section 4.2).

The random model shares the further technical defect that the regions are not manifolds. The asymmetrical model requires an algorithm for deciding which region is to get the boundary points. Some authors [22, 59] suggest assigning the points to object regions and not to empty space regions. Digital topology can be interpreted as assigning the points to the region currently of interest. However, these algorithms do not work when there are more than two types of regions (Fig. 11), contact between regions of the same type (Fig. 6), or regions which touch themselves (Figs. 1-2).

Alternatively, one could [36, 41] assign boundary points based on directions, e.g. give the points to the region above or to the left. However, this means that a region changes its topology if rotated. Alternatively, if there are no self-touching regions, we could assign an integer to all regions and give the points to the region with the lower number.

<sup>10</sup> Redefining the term “overlap” as in [8, 9] is possible but *ad hoc*.

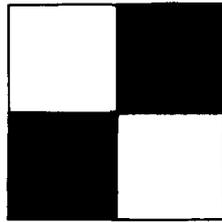


Fig. 12. At a four-way junction, the black regions are connected if the center point is assigned to them. Otherwise, the white regions are connected.

However, many regions will not be manifolds and the topology of a region may change as it breaks or comes into contact with other regions. At four-way junctions (Fig. 12), changes in the assignment of the center point alter which pair of regions is connected.<sup>11</sup>

### 5.3. Integers and hyperreals

Integers or integer grids are often used in computer vision [41,44] and occasionally proposed [48] for high-level reasoning. These models allow abrupt changes in function values across boundaries and adjacent regions are not connected to one another. However, the standard definitions of topological concepts do not accurately model the intuitive notions within regions. In particular, all functions from the integers to another space are continuous and no nontrivial subset of the integers is connected. It is possible to impose other topological structures on integer grids [32] but this simply causes them to act like  $\mathbb{R}^n$  for the purposes of the current discussion.<sup>12</sup>

An interesting recent proposal for qualitative reasoning [55] involves using the hyperreals as a model of time. In the order topology (as assumed in [55]), the hyperreals behave much like the integers. For each real value  $x$ , the set of points infinitely close to  $x$  is open. Thus, any set of real points, together with the hyperreals that are infinitely close to them, is open. Therefore, no region of non-infinitesimal length is connected, because it can be decomposed into two disjoint open intervals. Furthermore, continuous functions on the hyperreals can have abrupt jumps in value "at" any real number, e.g. a continuous function can have the value 1 on all reals in  $(-\infty, a)$  but 0 on all reals in  $[a, \infty)$ .

This conclusion may seem at odds with the standard hyperreal method for proving a real function continuous. The continuity proofs depend on a transfer principle which guarantees that certain statements about the reals are true if and only if the corresponding statements about the hyperreals are true. This transfer principle, however, is restricted to statements which quantify over the reals, whereas the standard topological definitions require quantification over *sets* of reals. Furthermore, the transfer principle only constrains the behavior of those hyperreal functions which happen to correspond to real functions, not arbitrary hyperreal functions. More general versions of the transfer

<sup>11</sup> In the new model, neither pair of regions is connected.

<sup>12</sup> A full discussion of such models is interesting but beyond the scope of this paper.

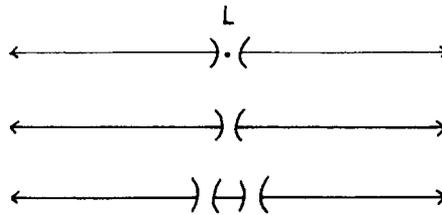


Fig. 13. Qualitative reasoning systems model a limit point  $L$  using three regions: below  $L$ , at  $L$ , and above  $L$  (top). The new models can represent these values using an immediate transition from above  $L$  to below  $L$ , i.e. two qualitative regions separated by a boundary (middle), or a third region, near  $L$ , can be added (bottom).

principle exist [27] but they do not alter the basic fact that the transfer principle is limited.

### 6. Worked examples

The new model handles boundaries in an unfamiliar way, allowing some new possibilities in modelling and making some familiar methods difficult. This section will explore how some standard analyses can be re-written using the new models and how the models can be used to develop new, simpler analyses of data that was previously puzzling.

#### 6.1. Crossing limits without limit points

A system for reasoning about the qualitative behavior of processes over time [19, 20, 22, 31, 33, 34, 36, 55, 56] is given the initial state of a system and must work out all qualitative states that could immediately follow it. This is done by finding the next limit point, i.e. important qualitative transition point, that each property value could reach. A typical limit point might be the boiling point of water or the location of an obstacle that a moving object might hit. The system establishes all possible sequences of events by working from limit point to limit point, building up a graph of possible states.

These systems typically represent a limit point  $B$  by breaking the space of values into three qualitative regions:  $(-\infty, B)$ ,  $[B]$ , and  $(B, \infty)$ . The new models cannot create an infinitely thin region. Instead, one can use two qualitative regions  $(-\infty, B)$  and  $(B, \infty)$  separated by a boundary. This would be appropriate when values change rapidly through the limit value, or when some discrete change will indicate the relationship of the value to  $B$  (e.g. the opacity of water indicating it is below the freezing point). When values are expected to remain within measurement error of  $B$  for nontrivial periods of time, one would instead use three regions:  $(-\infty, B - \epsilon)$ ,  $(B - \epsilon, B + \epsilon)$ , and  $(B + \epsilon, \infty)$  (Fig. 13). This is a sensible model for the boiling point of a water, because boiling water in a container will vary noticeably in temperature. Also, low-level data analysis algorithms (e.g. in computer vision) can easily produce 3-valued sign data of this type.

Some readers may wonder why I have separated distinct qualitative regions by boundaries. Most examples in this paper have concentrated on the topological structure of space

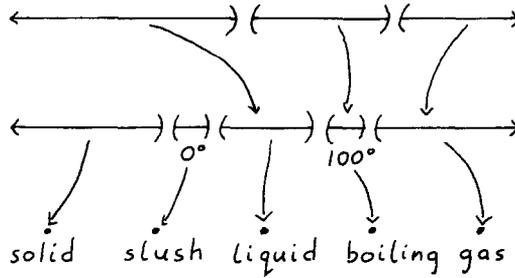


Fig. 14. The process of heating bulk water until it boils away can be modelled using a map from time to temperatures and a map from temperatures to states of matter.

and time, but similar arguments apply also to property spaces. Ranges of temperatures are distinguished as qualitatively different because some other property changes discontinuously as a function of temperature. For example, Fig. 14 shows how the physical state of bulk water depends on temperature. Since the physical states are a finite, discrete set, this can only be a continuous map if there are boundaries in the space of temperatures. The boundaries of qualitative regions in temperature space then induce event boundaries in time.

To use the new models, we could describe the qualitative reasoner as a process which gradually adds boundaries to time. Initially, time is represented as having no boundaries. The reasoning finds the limit point that each continuous process is moving towards. A continuous process moving towards a limit point cannot cross it, because the image of a connected set under a continuous function is connected. Nor can it halt or turn back, because that would require a change in the qualitative sign of its derivative. Therefore, the representation is not consistent.

Each step in limit analysis detects when the system could next become stuck in this way. For each process that could get in trouble, it then hypothesizes that the system crosses the limit point and adds a boundary to time to license the abrupt change in state. It then figures out what else might change at this temporal boundary and starts up the new processes for the next interval of time. In this analysis, all qualitative states last for intervals of nontrivial length and all changes between states occur instantaneously.

### 6.2. Equilibria without exact equality

Reasoning about equilibria can also be rewritten so as to use region edges rather than boundaries. For example, Forbus [19] represents an equilibrium as a cycle of qualitative states in which values are perturbed from exact equality for “an instant” and then returned equally fast to exact equality. In the new models, we cannot create a single-point region representing exact equality. Instead, we must place a boundary at the limit point, dividing the property space into two qualitative regions (Fig. 15).

To explain why the equilibrium is stable, we first extract the borders of the two regions,  $L_-$  below the limit point and  $L_+$  above it. The length of the borders is set so that  $L = L_- \cup L_+$  contains all points indistinguishable from the limit point given

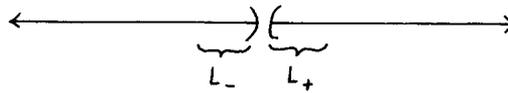


Fig. 15. The new representation of equilibrium splits the property space with a boundary at the point of exact equality. We then identify two regions,  $L_-$  and  $L_+$ , which are borders adjacent to the boundary.

the prevailing errors. Errors will occur both due to limitations in measurement devices and approximations inherent in the theoretical model (e.g. water in a container is not perfectly level). In  $L_-$ , a restoring process moves the value towards the boundary with  $L_+$ . In  $L_+$ , a restoring process moves the value towards the boundary with  $L_-$ . Therefore, no matter which of the two sub-states it is in, the system cannot leave  $L$  but can only shuttle between  $L_-$  and  $L_+$ . In other words, the system maintains equilibrium up to measurement error.

### 6.3. Verb classes

Recent work in linguistic semantics [2, 10, 37, 47, 50, 51, 53] has classified English verbs and verbal constituents into four types. A verbal constituent may be either:

- a state,
- an activity (a pattern of change in state over time),
- a state change (traditionally called an “achievement”), or
- an accomplishment (an activity ending in, and causing, a state change).

Some examples of the different types are:

- *states*: be green, like mathematics, own a car,
- *activities*: swim, drive a car, shred documents,
- *state changes*: reach the finish line, find a free terminal, pass an exam,
- *accomplishments*: build a house, do a problem set, make a pot of coffee.

The type of a verb or verb phrase affects the syntactic and semantic contexts it can occur in (see in particular [10, pp. 55–60 and 184] and [51]). Based on a fuller analysis of tense and aspect phenomena [12, 13], I treat progressive, perfect, and habitual (simple present) verb forms as states and homogeneous actions such as “fall” or “stand” as activities.

When two actions occur in the same sequence of events, there are constraints on how they can overlap in time [2, 11, 24, 25, 49, 51, 57]. These constraints depend on the classes of the verb phrases involved. State changes and accomplishments can only overlap in limited ways, whereas overlap among states and activities is largely unconstrained. Previous descriptions of these phenomena, though insightful, require separate cases for the two types of predicates. Using the new models of boundaries, we can generate unified explanations.

Specifically, I will model a state or activity as occupying a connected region of time (no internal boundaries), whereas a state change or an accomplishment occupies an interval containing one or more internal boundaries (Fig. 16). A state change occupies the two borders adjacent to a boundary, just enough time to observe the states before and after the change. An accomplishment occupies a connected interval of nontrivial length,

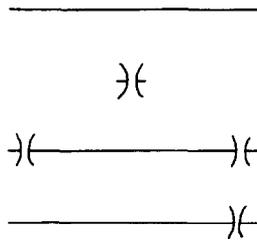


Fig. 16. Models of the four verb classes: state or activity (top), state change (upper middle), accomplishment (lower middle, bottom).

ending in a state change and possibly starting with one.<sup>13</sup> These models are similar to those used by previous authors except for the topological details around the boundaries. It is important that the representations of accomplishments include the state resulting from the action (here represented by the border region following the state change) in order to describe the semantics of perfect verb forms [12, 13, 29, 57].

In the new models, the presence or absence of boundaries affects how actions can fit together in any one model of a situation. Let us hypothesize that

**Constraint 4.** *A description of a sequence of events in English can typically be interpreted using only one model, with a consistent assignment of boundary locations.*

If this constraint holds, a boundary specified by one action cannot fall during an interval that another action specifies as connected, since this would require constructing two situation models.

In English discourse, a sequence of events can be indicated using explicit temporal connectives [24, 57] or implicitly by the order of sentences [11, 25]. Let us consider the case of the temporal connective “when”. We can represent its meaning as:

“when  $X$ ,  $Y$ ” or “ $Y$ , when  $X$ ” means that  $X$  and  $Y$  occur over overlapping intervals of time  $I_X$  and  $I_Y$  such that no part of  $I_Y$  is earlier than all of  $I_X$ .

The temporal relations allowed by “when” follow from the interaction of this meaning with the constraints imposed by the boundary models.

Because states and activities contain no internal boundaries, they can overlap freely. As a verb can refer to any interval which matches its meaning, not just the maximal such interval, this means that two states or activities connected by “when” are merely required to overlap in time (Sentence 6). Similarly, a state or activity must simply overlap a state change: the specific verbs in Sentences 7–8 favor two different temporal patterns. A state or activity can overlap the end state of an accomplishment and/or fill its main interval (Sentences 9–10). It cannot start or end during the main interval of the accomplishment, because that would imply interrupting the main interval with a state change.

<sup>13</sup> It is not clear whether there is a state change at the start of an accomplishment. Or, whether there are two types of accomplishments, one with the state change and one without it.

**Sentence 6.** *The heating was on when I was in the library.*

**Sentence 7.** *Mirabelle was married when she passed her driving test.*

**Sentence 8.** *Mirabelle was happy when she passed her driving test.*

**Sentence 9.** *When Mike ran up the stairs, he was exhausted.*

**Sentence 10.** *When Mike ran up the stairs, they were covered in snow.*

By contrast, if two state changes or accomplishments overlap, they must agree about the presence or absence of boundaries. So Sentence 11 has two possible readings. The two state changes might have happened at exactly the same time. Or, the start of the second action overlaps the ending state of the first, i.e. the actions happened in quick succession. In either case, they would naturally be interpreted as having tight causal links to one another.

**Sentence 11.** *The bomb exploded when Curtis pushed the button.*

Consistency of boundary assignment also explains the “subinterval” property also used to distinguish states and activities from state changes and accomplishments [5, 37, 49, 50]. If a homogeneous state or activity holds over an interval, it holds over every subinterval. By contrast, if a state change or accomplishment holds over an interval, it holds only over those subintervals which still have the correct pattern of boundaries. Similar facts distinguish mass nouns, which refer to types of stuff, from count nouns, which refer to objects: a subset of water is still water, but a subset of a table is not necessarily a table (cf. [1, 50]). A slight complication is that many activity verbs and mass nouns impose minimum length requirements on their intervals. So, for example, a region of space only contains “fruitcake” if it contains a representative selection of cake, raisins, and nuts [50].

#### 6.4. *Splitting and merging*

One feature of the new boundary model is that it makes a sharp distinction between two regions which are fused together (connected) and two regions which are merely adjacent (not connected). The process of grafting trees, illustrated in Fig. 17, shows the distinction clearly. The pieces of plant involved in the graft only fuse together sometime after they are brought into contact with one another. Other materials (e.g. pudding) make this transition much more quickly, in some cases (e.g. water) effectively fusing instantaneously upon contact. Similar processes occur in reverse, e.g. a plate cracks in the microwave but the two pieces separate only when someone tries to pick it up.

This distinction is quite deeply embedded in our intuitive models of the world. Many English verbs clearly indicate which type of action they refer to:

- *changes in contact only*: put into, stack, hang on, remove from, take off of,
- *changes in connectivity*: fuse, glue, merge, cut, saw off, split, tear.



Fig. 17. In grafting, a tree branch is brought into contact with a stump. After a while, they grow together.

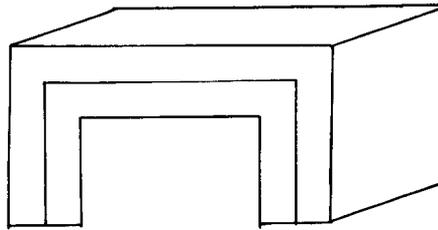


Fig. 18. Two nested tables do not form a single table. Even if their exterior boundary has the correct shape, there is an extraneous internal boundary.

We can model this distinction by providing two types of high-level operations for altering models. The first type moves regions and/or boundaries, shrinks or enlarges regions, and reshapes or rearranges regions using isotopies. The second type adds or deletes sections of boundary.

This distinction helps explain the “concatenation” property sometimes used to distinguish count nouns from mass nouns, and to distinguish state changes and accomplishments from states and activities [4, 37, 49, 51]. In general, the union of two objects is not an object of the same type, e.g. two adjacent tables do not form a table. The two can be distinguished by whether an internal boundary is present (Fig. 18). By contrast, two adjacent piles of sand form a single pile of sand, because piles of sand fuse on contact. Similarly, if you build a house over interval  $I$  and build another house over the adjacent interval  $J$ , you cannot be said to have built a (single) house over interval  $I \cup J$ . However, if you sing Bach over interval  $I$  and then again over interval  $J$ , you can also be said to have sung Bach over interval  $I \cup J$ .

### 6.5. Digitized isotopy

As discussed above, because the new models are subspaces of  $\mathbb{R}^n$ , we can use the notion of isotopy to model constraints on deformation of one representation into another.

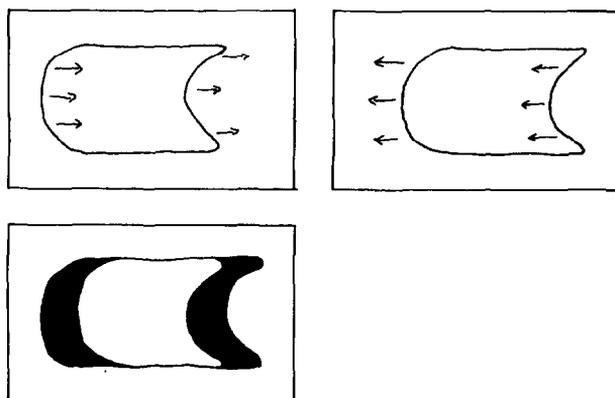


Fig. 19. Two nearly aligned images (top) are brought into correspondence by enlarging the boundaries in each image, using isotopies of the regions. This process converts both images to a common form (bottom), so the images must be isotopic to one another.

This idea has been implemented in a stereo matching system [14]. When a pattern of surface markings is projected onto the left and right images of a stereo pair, the regions in the two images must be isotopic. Together with other constraints (e.g. dark regions must match dark regions), this can be used to identify when a patch of the left image is a possible match to a patch of the right image.

In general, it is quite difficult to tell whether two scenes are isotopic. However, in the bottom level of stereo matching, we can assume that the two images are aligned almost correctly. Thus, two corresponding regions occupy overlapping regions of the plane. The matcher tries to eliminate mis-matches between the images by enlarging the boundaries in each image (Fig. 19). Where mis-matches are completely eliminated by this process, there must exist (locally) an isotopy between the two scenes. The digitized boundary model given in Section 4.1 provides an exact correspondence between the digitized operations used in implementation and continuous isotopies of the underlying images.

This technique has also been used in building a fast algorithm for finding boundaries in textured images [17]. Because this algorithm requires regions to be at least 12 pixels wide, boundaries are initially detected using a sub-sampled version of the image. These poorly localized boundaries are expanded to thick boundaries at full image resolution. The boundary locations are refined on the basis of high resolution intensity information, using the digitized isotopy operations to prevent changes in the image topology.

## 7. Conclusions

We have seen how topological properties have been, and could be, used in a wide range of reasoning algorithms. Standard models of space and time, based on  $\mathbb{R}^n$ , were shown to have the wrong topological structure at object or event boundaries. New models with an appropriate topology were constructed by deleting points in the boundaries. A

number of worked examples were presented, illustrating how the new models support standard reasoning algorithms (e.g. qualitative reasoning), simplify analysis of some phenomena (e.g. verb classes), provide formal models for situations that could not be properly modelled by previous authors (e.g. regions in contact), and suggest interesting new analyses based on change or lack of change in scene topology (e.g. isotopy-based matching, split and merger of regions).

### Acknowledgements

This research was carried out at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology, the Department of Engineering Science of the University of Oxford, and the Department of Computer Science of the University of Iowa. Support for the MIT Artificial Intelligence Laboratory's artificial intelligence research was provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-85-K-0124. The author was also supported by the Fannie and John Hertz Foundation, by the AT&T Bell Laboratories Graduate Research Program for Women, and by a junior research fellowship funded by British Petroleum.

This paper covers a wide range of material and many people have given me pointers and suggestions. Particularly memorable help came from Mike Brady, David Chapman, Bruce Donald, George Fleck, Ken Forbus, Ken Hale, Jim Henle, Rich Zippel, the topology group at Iowa, and an anonymous reviewer. None of them, of course, are responsible for any errors in this paper.

### References

- [1] K. Allan, Nouns and countability, *Language* **56** (3) (1980) 541-567.
- [2] J.F. Allen, Towards a general model of action and time, *Artif. Intell.* **23** (2) (1984) 123-154.
- [3] J.F. Allen and P.J. Hayes, A common-sense theory of time, in: *Proceedings IJCAI-85*, Los Angeles, CA (1985) 528-531.
- [4] E. Bach, The algebra of events, *Linguist. Philos.* **9** (1) (1986) 5-16.
- [5] M. Bennett and B. Partee, *Toward the Logic of Tense and Aspect in English* (Indiana University Linguistics Club, 1978).
- [6] J.M. Erady and H. Asada, Smoothed local symmetries and their implementation, *Int. J. Robotics Res.* **3** (3) (1984) 36-61.
- [7] J.F. Canny, A computational approach to edge detection, *IEEE Trans. Comput. Vision Pattern Recognition* **8** (6) (1986) 679-698.
- [8] E. Davis, An ontology of physical action, Technical Report #123, Computer Science, Courant Institute, New York University, New York (1984).
- [9] E. Davis, Shape and function of solid objects: some examples, Technical Report #137, Computer Science, Courant Institute, New York University, New York (1984).
- [10] D.R. Dowty, *Word Meaning and Montague Grammar* (Reidel, Dordrecht, Netherlands, 1979).
- [11] D.R. Dowty, The effects of aspectual class on the temporal structure of discourse: semantics or pragmatics?, *Linguist. Philos.* **9** (1) (1986) 37-61.
- [12] M.M. Fleck, Boundaries and topological algorithms, Ph.D. Thesis, Electrical Engineering and Computer Science Department, MIT, AI Laboratory TR 1065, Cambridge, MA (1988).

- [13] M.M. Fleck, Topological models for space and time, OUEL Report No. 1860/90, Oxford University, Department of Engineering Science (1990).
- [14] M.M. Fleck, A topological stereo matcher, *Int. J. Comput. Vision* **6** (3) (1991) 197–226.
- [15] M.M. Fleck, Some defects in finite-difference edge finders, *IEEE Trans. Pattern Anal. Mach. Intell.* **14** (3) (1992) 337–345.
- [16] M.M. Fleck, Texture: plus ça change . . ., in: *Proceedings European Conference on Computer Vision*, Lecture Notes in Computer Science **588** (Springer-Verlag, Berlin, 1992) 151–159.
- [17] M.M. Fleck, Practical edge finding with a robust estimator, in: *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, Seattle, WA (1994) 649–653.
- [18] M.M. Fleck, Shape and the wide-angle image, Technical Report 94-04, Department of Computer Science, University of Iowa, Iowa City, IA (1994).
- [19] K.D. Forbus, Qualitative process theory, *Artif. Intell.* **24** (1984) 85–168.
- [20] K.D. Forbus, The logic of occurrence, in: *Proceedings IJCAI-87*, Milan, Italy (1987) 409–415.
- [21] C.L. Hamblin, Instants and intervals, in: J.T. Fraser, F.C. Haber, and G.H. Müller, eds., *The Study of Time, First Conference of the International Society for the Study of Time* (Springer-Verlag, Berlin, 1972) 324–331.
- [22] P.J. Hayes, Naive physics 1: ontology for liquids, in: J.R. Hobbs and R.C. Moore, eds., *Formal Theories of the Commonsense World* (Ablex, Norwood, NJ, 1985) 71–107.
- [23] P.J. Hayes, The second naive physics manifesto, in: J.R. Hobbs and R.C. Moore, eds., *Formal Theories of the Commonsense World* (Ablex, Norwood, NJ, 1985) 1–36.
- [24] O. Heinämäki, Semantics of English temporal connectives, Ph.D. Thesis, University of Texas, Austin, TX (1974); distributed by the Indiana University Linguistics Club (1978).
- [25] E. Hinrichs, Temporal anaphora in discourses of English, *Linguist. Philos.* **9** (1) (1986) 63–82.
- [26] J.E. Hopcroft and D.B. Kraftt, The challenge of robotics for computer science, in: J.T. Schwartz and C.K. Yap, eds., *Algorithmic and Geometric Aspects of Robotics*, Advances in Robotics **1** (Lawrence Erlbaum, Hillsdale, NJ, 1987) 7–43.
- [27] R.F. Hoskins, *Standard and Nonstandard Analysis* (Ellis-Horwood, Chichester, England, 1990).
- [28] J.F.P. Hudson, *Piecewise Linear Topology* (Benjamin, New York, 1969).
- [29] M.R. Johnson, A unified temporal theory of tense and aspect, in: P.J. Tedeschi and A. Zaenen, eds., *Tense and Aspect, Syntax and Semantics* **14** (Academic Press, New York, 1981) 145–175.
- [30] H. Kamp, Events, instants and temporal reference, in: R. Bäuerle, U. Egli and A. von Stechow, eds., *Semantics from Different Points of View* (Springer-Verlag, Berlin, 1979) 376–417.
- [31] J. de Kleer and J.S. Brown, A qualitative physics based on confluences, *Artif. Intell.* **24** (1984) 7–83.
- [32] T.Y. Kong and R. Kopperman, A topological approach to digital topology, *Amer. Math. Monthly* **98** (10) (1991) 901–917.
- [33] B.J. Kuipers, Commonsense reasoning about causality: deriving behavior from structure, *Artif. Intell.* **24** (1984) 169–203.
- [34] B.J. Kuipers, Qualitative simulation, *Artif. Intell.* **29** (1986) 289–338.
- [35] Y. Leclerc and S.W. Zucker, The local structure of image discontinuities in one dimension, *IEEE Trans. Pattern Anal. Mach. Intell.* **9** (3) (1987) 341–355.
- [36] D. McDermott, A temporal logic for reasoning about processes and plans, *Cognitive Sci.* **6** (1982) 101–155.
- [37] A.P.D. Mourelatos, Events, processes and states, in: P.J. Tedeschi and A. Zaenen, eds., *Tense and Aspect, Syntax and Semantics* **14** (Academic Press, New York, 1981) 191–212.
- [38] J.R. Munkres, *Topology: A First Course* (Prentice-Hall, Englewood Cliffs, NJ, 1975).
- [39] J.R. Munkres, *Elements of Algebraic Topology* (Addison-Wesley, Menlo Park, CA, 1984).
- [40] J.A. Noble, Descriptions of image surfaces, D.Phil. Thesis, Department of Engineering Science, University of Oxford (1989).
- [41] T. Pavlidis, *Structural Pattern Recognition* (Springer-Verlag, Berlin, 1977).
- [42] D.E. Pearson and J.A. Robinson, Visual communication at very low data rates, *Proc. IEEE* **73** (1985) 795–812.
- [43] S.B. Pollard, J.E.W. Mayhew and J.P. Frisby, PMF: a stereo correspondence algorithm using a disparity gradient limit, *Perception* **14** (1985) 449–470.
- [44] A. Rosenfeld, Digital topology, *Amer. Math. Monthly* **86** (6) (1979) 621–630.

- [45] C.P. Rourke and B.J. Sanderson, *Introduction to Piecewise-Linear Topology* (Springer-Verlag, Berlin, 1982).
- [46] T.B. Rushing, *Topological Embeddings* (Academic Press, New York, 1973).
- [47] G. Ryle, *The Concept of Mind* (Barnes and Noble, London, 1949).
- [48] Y. Shoham, Reasoning about change: time and causation from the standpoint of artificial intelligence, Ph.D. Thesis, Report YALEU/CSD/RR #507, Department of Computer Science, Yale University, New Haven, CT (1987).
- [49] Y. Shoham, Temporal logics in AI: semantical and ontological considerations, *Artif. Intell.* **33** (1) (1987) 89–104.
- [50] B. Taylor, Tense and continuity, *Linguist. Philos.* **1** (1977) 199–220.
- [51] C.L. Tenny, Grammaticalizing aspect and affectedness, Ph.D. Thesis, Department of Linguistics and Philosophy, MIT, Cambridge, MA (1987).
- [52] J.F.A.K. van Benthem, *The Logic of Time* (Reidel, Dordrecht, Netherlands, 1983).
- [53] Z. Vendler, *Linguistics in Philosophy* (Cornell University Press, Ithaca, NY, 1967).
- [54] D.D. Weld, The use of aggregation in causal simulation, *Artif. Intell.* **30** (1) (1986) 1–34.
- [55] D.D. Weld, Exaggeration, *Artif. Intell.* **43** (1990) 311–368.
- [56] B.C. Williams, Qualitative analysis of MOS circuits, *Artif. Intell.* **24** (1984) 281–346.
- [57] E. Woisetschlaeger, A Semantic Theory of the English Auxiliary System, Ph.D. Thesis, Department of Linguistics and Philosophy, MIT, Cambridge, MA; distributed by the Indiana University Linguistics Club (1976).
- [58] G.H. von Wright, *An Essay in Deontic Logic and the General Theory of Action* (North-Holland, Amsterdam, 1968).
- [59] C.K. Yap, Algorithmic motion planning, in: J.T. Schwartz and C.K. Yap, eds., *Algorithmic and Geometric Aspects of Robotics*, Advances in Robotics **1** (Lawrence Erlbaum, Hillsdale, NJ, 1987) 95–145.