University of Groningen

Computing Spatial Similarity by Games

Aiello, Marco

Published in:
EPRINTS-BOOK-TITLE

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2001

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Aiello, M. (2001). Computing Spatial Similarity by Games. In *EPRINTS-BOOK-TITLE* University of Groningen, Johann Bernoulli Institute for Mathematics and Computer Science.

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Computing Spatial Similarity by Games

Marco Aiello

Institute for Logic, Language and Computation, and
Intelligent Sensory and Information Systems
University of Amsterdam
Plantage Muidergracht 24 1018 TV Amsterdam, The Netherlands
aiellom@ieee.org

Abstract. The multi-modal logic $S4_u$, known in the field of qualitative spatial reasoning to be a decidable formalism for expressing topological and mereological properties, can also be exploited to define a distance measure among patterns. Here, we recall the notion of topological distance defined in terms of games over $S4_u$ models, and show how it is effectively computed for a specific class of models: the class of polygons of the real plane, a class of topological models widely used in computer science and AI applications. Finally, we briefly overview an implemented system based on the presented framework. This paper is the practical counterpart of, and continuation to [1].

Keywords: qualitative spatial reasoning, model comparison games, image similarity

1 Introduction

The core of the question we address is *How can I compute how similar two spatial patterns are?* The fundamental issues to solve in order to answer this question involve finding an agreement on spatial representation, finding an agreement on a language to describe spatial patterns, and finding an agreement on a measure of similarity. Our choice here falls onto modal logics, topologically interpreted, and equipped with adequate model comparison games. The language, called $S4_u$, is a multi-modal $S4^*S5$ logic interpreted on topological spaces equipped with valuation functions.

Spatial representation is not only interesting in itself, but also when considering its applications. It is essential in vision, in spatial reasoning for robotics, in geographical information systems, and many more related fields. Of paramount importance in applications is the comparison of spatial patterns, which must be represented in the same way. We consider similarity measures and look at their application to image retrieval. Image retrieval is concerned with the indexing and retrieval of images from a database, according to some desired set of image features. These features can be as diverse as textual annotations, color, texture, object shape, and spatial relationships among objects. The way the features from different images are compared, in order to have a measure of similarity among images, is what really distinguishes an image retrieval architecture from another

one. We refer to [11] for an overview of image retrieval and more specifically to [18] for image similarity measures. Here we concentrate on image retrieval based on spatial relationships at the qualitative level of mereotopology, that is, part-whole relations, topological relations and topological properties of individual regions (see for instance [3,19,8]). Other image retrieval systems are based on spatial relationships as the main retrieval feature [20,15,9].

The paper is organized as follows. First, we recall the basic facts of the spatial framework and in particular of the similarity measure. The work overviewed in Section 2 is based on [2] and [1], to which we refer the reader for details and examples. In Section 3, we present an algorithm to compute the similarity measure in the case of polygons of the real plane. The techniques described in the paper have been used to implement an image retrieval prototype named IRIS (Image RetrIeval based on Spatial relationships) which is overviewed in Section 4.

2 A general framework for Mereotopology

The framework we adopt to express spatial properties at the mereotopological level is the multi-modal logic $S4^*S5$, usually referred to as $S4_u$. The language is known in modal logics [13], and has been introduced into spatial reasoning by Bennett [5] to encode decidable fragments of the RCC calculus [17]. For the syntax, axiomatization, truth definition and topological expressive power we refer to [2], while for an analysis of the mereotopological expressive power and a comparison with RCC we refer to [1].

Let us only say that the modal logic is interpreted on topological spaces (à la Tarski [21]) instead of the usual Kripke semantics. Every formula φ of $S4_u$ represents a region. $\Box\varphi$ is interpreted as “interior of the region φ ” and $U\varphi$ as “it is the case everywhere that φ .”

For $S4_u$ it is possible to define a notion of equivalence resorting to an adequate notion of bisimulation ([6,14]), after all we are dealing with a modal logic... For a definition and proof of adequacy see [2]. This notion lets us answer questions like *When are two spatial patterns the same?* or *When is a pattern a sub-pattern of another one?* If topological bisimulation is satisfactory from the formal point of view, one needs more to address qualitative spatial reasoning problems and computer vision issues. If two models are not bisimilar, or one does not simulate the other, one must be able to quantify the difference between the two models. Furthermore, this difference should behave in a coherent manner across the class of all models. Informally, one needs to answer questions like: *How different are two spatial patterns?*

To this end, we defined an adequate notion of model comparison game in the Ehrenfeucht-Fraïssé style. The idea is that two players challenge each other on two models. One player (Spoiler) is attempting to prove the difference of the models, while the other one (Duplicator) wants to prove their equivalence. The moves available to the players are those of deciding on which model to play, which type of round to engage, and that of picking points and opens on the two

models. A game is played to a fixed number of rounds n . We denote a game by $TG(X, X', n)$, where X and X' are two topological models, i.e., a topological space $\langle X, O \rangle$ equipped with a valuation function ν , and n is the number of rounds. For the precise definition we refer, again, to [2]. The *multi-modal rank* of a $S4_u$ formula is the maximum number of nested modal operators appearing in it (i.e. $\Box$, $\Diamond$, U and E modalities). The following adequacy of the games with respect to the mereotopological language holds.

Theorem 1 (Adequacy). Duplicator has a winning strategy (w.s.) in n rounds in $TG(X, X', n)$ iff X and X' satisfy the same formulas of multi-modal rank at most n .

Various examples of plays and a discussion of winning strategies can be found in [2]. The interesting result is that of having a game theoretic tool to compare topological models. Given any two models, they can be played upon. If Spoiler has a winning strategy in a certain number of rounds, then the two models are different up to a certain degree. The degree is exactly the minimal number of rounds needed by Spoiler to win. On the other hand, one knows (see [2]) that if Spoiler has no w.s. in any number of rounds, and therefore Duplicator has in all games, including the infinite round game, then the two models are bisimilar.

A way of comparing any two given models is not of great use by itself. It is essential instead to have some kind of measure. It turns out that topo-games can be used to define a distance measure.

Definition 1 (isosceles topo-distance). Consider the space of all topological models T . *Spoiler's shortest possible win* is the function $spw : T \times T \rightarrow \mathbb{N} \cup \{\infty\}$, defined as:

$$spw(X_1, X_2) = \begin{cases} n & \text{if Spoiler has a winning strategy in } TG(X_1, X_2, n), \\ & \text{but not in } TG(X_1, X_2, n-1) \\ \infty & \text{if Spoiler does not have a winning strategy in} \\ & TG(X_1, X_2, \infty) \end{cases}$$

The *isosceles topo-model distance* (*topo-distance*, for short) between X_1 and X_2 is the function $tmd : T \times T \rightarrow [0, 1]$ defined as:

$$tmd(X_1, X_2) = \frac{1}{spw(X_1, X_2)}$$

In [1], it is shown that indeed the above definition is a distance measure on the class of all models for the language:

Theorem 2 (isosceles topo-model distance). tmd is a distance measure on the space of all topological models.

3 Computing similarities

The fundamental step to move from theory to practice has been taken when shifting from model comparison games to a distance. To complete the journey

towards practice one needs to identify ways of effectively computing the distance in cases actually occurring in real life domains. We do not have an answer to the general question of whether the topo-distance is computable for any two topological models or not. Though, by restricting to a specific class of topological models widely used in real life applications, we can show the topo-distance to be computable when one makes an ontological commitment. The commitment consists of considering topological spaces made of polygons. This is common practice in various application domains such as geographical information systems (GIS), in many branches of image retrieval and of computer vision, in robot planning, just to mention the most common.

Consider the real plane $\mathbb{R}^2$, any line in $\mathbb{R}^2$ cuts it into two open half-planes. We call a half-plane *closed* if it includes the cutting line, *open* otherwise.

Definition 2 (region). A *polygon* is the intersection of finitely many open or closed half-planes. An *atomic region* of $\mathbb{R}^2$ is the union of finitely many polygons.

An atomic region is denoted by one propositional letter. More in general, any set of atomic regions, simply called *region*, is denoted by a $S4_u$ formula. The polygons of the plane equipped with a valuation function, denoted by $M_{\mathbb{R}^2}$, are in full rights a topological model as defined in Section 2, a basic topological fact. A similar definition of region can be found in [16]. In that article Pratt and Lemon also provide a collection of fundamental results regarding the plane, polygonal ontology just defined (actually one in which the regions are open regular).

From the model theoretic point of view, the advantage of working with $M_{\mathbb{R}^2}$ is that we can prove a logical finiteness result and thus give a terminating algorithm to compute the topo-distance. The preliminary step is thus that of proving a finiteness lemma for $S4_u$ over $M_{\mathbb{R}^2}$ models.¹

Lemma 1 (finiteness). There are only finitely many modally definable subsets of a finite set of regions $\{r_i | r_i \text{ is an atomic region}\}$.

Here is a proof sketch. We work by enumerating cases, i.e., considering boolean combinations of planes, adding to an ‘empty’ space one half-plane at the time, first to build one region r , and then to build a finite set of regions. The goal is to show that only finitely many possibilities exist. We begin by placing a half plane denoted by r on an empty bidimensional space, Figure 1.a. Let us follow what happens to points in the space from left to right. On the left, points satisfy the formula $E(r \wedge \Box r)$ and its subformulas Er and $E\Box r$. This is true until we reach the frontier point of the half-plane. Either $E(\neg r \wedge \Diamond \Box r)$ or $E(r \wedge \Diamond \Box \neg r)$ are true depending on whether the half-plane is open or closed, respectively. Once the frontier has been passed to the right, the points satisfy $E(\neg r \wedge \Box \neg r)$ and its subformulas $E\neg r$ and $E\Box \neg r$, better seen in Figure 1.a. In fact, if we consider negation in the formulas the role of r and $\neg r$ switch. Consider now a second plane in the picture:

¹ Of course, in general this is not true. There are infinitely many non equivalent $S4_u$ formulas and one can identify appropriate Kripke models to show this, [7].

Fig. 1. Basic formulas defined by one region.

- **Intersection:** the intersection may be empty (no new formula), may be a polygon with two sides and vertices (no new formula, the same situation as with one polygon), or it may be a line, the case of two closed polygons that share the side (in this last case depicted in Figure 1.b—*slope*—we have a new formula, namely, $E(r \wedge \square \diamond \neg r)$).
- **Union:** the union may be a polygon with either one or two sides (no new formula), two separated polygons (no new formula), or two open polygons sharing the open side (this last case depicted in Figure 1.¬b—*crack*—is like the spike, one inverts the roles r and $\neg r$ in the formula: $E(\neg r \wedge \square \diamond r)$).

Finally, consider combining cases (a) and (b). By union, we get Figure 1.a, 1.c, 1.d. The only situation bringing new formulas is the latter. In particular, the point where the line intersects the plane satisfies the formula: $E(\diamond \square r \wedge \diamond(r \wedge \square \diamond \neg r))$. By intersection, we get a segment, or the empty space, thus, no new formula.

The four basic configurations just identified yield no new configuration from the $S4_u$ point of view. To see this, consider the boolean combinations of the above configurations. We begin by negation (complement):

Union straightforwardly follows (where a stands for both a and $\neg a$, as both configurations always appear together):

$\cup$	a	b	c	d
a	a, $\neg b$, $\neg d$	a, c, d	a, $\neg b$, c, d, $\neg d$	a, $\neg b$, d, $\neg d$
b	a, c, d	b	c, d	d
c	a, $\neg b$, c, d, $\neg d$	c, d	a, $\neg b$, c, d, $\neg d$	a, $\neg b$, c, d, $\neg d$
d	a, $\neg b$, d, $\neg d$	d	a, $\neg b$, c, d, $\neg d$	a, $\neg b$, d, $\neg d$

The table for intersection follows, with the proviso that the combination of the two regions can always be empty (not reported in the table) and again a and $\neg a$ are represented simply by a:

$\cap$	a	b	c	d
a	a, b, c, d	b	a, b, c	a, b, d
b	b	b	b	b
c	a, b, c	b	a, b, c, d	a, b, c, d
d	a, b, d	b	a, b, c, d	a, b, c, d

We call *topo-vector* associated with the region r , notation $\mathbf{r}$, an ordered sequence of ten boolean values. The values represent whether the region r satisfies or not the ten formulas $\{Er, E\neg r, E\Box r, E\Box\neg r, E(\neg r \wedge \Diamond\Box r), E(r \wedge \Diamond\Box\neg r), E(r \wedge \Box\Diamond\neg r), E(\neg r \wedge \Box\Diamond r), E(\Diamond\Box r \wedge \Diamond(r \wedge \Box\Diamond\neg r)), E(\Diamond\Box\neg r \wedge \Diamond(\neg r \wedge \Box\Diamond r))\}$. The ten formulas are those identified in Figure 1 which we have shown to be the only one definable by boolean combinations of planes denoting the same one region r . For example, the topo-vector associated with a plate—a closed square r in the plane—is {true, true, true, true, false, true, false, false, false, false}.

Adding half-planes with different denotations $r_2, r_3, \dots$ increases the number of defined formulas. The definition of topo-vector is extended to an entire $M_{\mathbb{R}^2}$ model: $\{ E\bigwedge_i [\neg]r_i, E\bigwedge_i \Box[\neg]r_i, E(\bigwedge_i [\neg]^+r_i \wedge \bigwedge_i \Diamond\Box[\neg]^*r_i), E(\bigwedge_i [\neg]^+r_i \wedge \bigwedge_i \Box\Diamond[\neg]^*r_i), E(\bigwedge_i \Diamond\Box[\neg]^+r_i \wedge \Diamond(\bigwedge_i [\neg]^+r_i \wedge \bigwedge_i \Box\Diamond[\neg]^*r_i)) \}$, where $[\]$ denotes an option and if the option $[\]^+$ is used then the option $[\]^*$ is not and viceversa. The topo-vector is built such that the modal rank of the formulas is not decreasing going from the positions with lower index to those with higher. The size of such a vector is $5 \cdot 2^i$ where i is the number of denoted regions of the model. The fact that the size of the topo-vector grows exponentially with the number of regions might seem a serious drawback. Though, as we shall show in a moment, the topo-vector stores all the information relevant for $S4_u$ about the model. Furthermore, the size of a topo-vector is most often considerably smaller than that of a topological model. In fact, a topo-vector is of exponential size in the number of regions, while a topological model is of exponential size in the number of points of the space because of the set of opens. As a final argument,

one should add that in practical situations the number of regions is always much smaller than the number of points of the space.

We are now in a position to devise an algorithm to compute the topo-distance between two topological models. The algorithm works by first computing the associated topo-vectors and then comparing them. By the comparison it is possible to establish which formulas differentiate the two models and therefore the distance between the two models. Here is the general algorithm (in pseudo-code) to compute the topo-distance between two topological models M_1 and M_2 :

```

topo-distance( $M_1$ ,  $M_2$ )
   $v_1$  = topo-vector ( $M_1$ )
   $v_2$  = topo-vector ( $M_2$ )
  align  $v_1$  and  $v_2$ 
  loop on  $v_1$   $v_2$  with index  $i$ 
    if  $v_1(i) \neq v_2(i)$ 
      return  $\frac{1}{\text{modal rank}(v_1(i))}$ 
  return 0

```

The idea is of retrieving the topo-vectors associated with the two input models and then looping over their elements. The inequality check can also be thought of as a **xor**, since the elements of the array are booleans. If the condition is never satisfied, the two topo-vectors are identical, the two-models are topo-bisimilar and thus the topo-distance is null. The **align** command makes the topo-vectors of the same length and aligns the formulas of the two, i.e., such that to the same index in the vector corresponds the same formula. If a topo-vector contains a formula that the other one does not, the entry is added to the vector missing it with a false value. To complete the description of the algorithm, we provide the function to compute the topo-vector associated with an $M_{\mathbb{R}^2}$ model:

```

topo-vector( $M$ )
   $v$  = initialized to all false values
  loop on regions  $r$  of  $M$  with index  $i$ 
    loop on atomic regions  $a$  of  $r(i)$  with index  $j$ 
      loop on vertices  $v$  of  $a(j)$  with index  $k$ 
        update  $v$  with the point  $v(k)$ 
        if  $v(k)$  is not free
          loop on intersections  $x$  of  $a(j)$  with all
            regions of  $M$  with index  $l$ 
              update  $v$  with the point  $x(l)$ 
  return  $v$ 

```

If a point $v(k)$ of an atomic region $a(j)$ is contained in any polygon different from $a(j)$ and it is not contained in any other

Fig. 2. Computing the topo-vector on a simple model.

region, then the condition $v(k)$ is **not free** is satisfied. Standard computational geometry algorithms exist for this task, [10]. When the **update v with the point p** function is called, one checks in which case p is (as shown after Lemma 1), then one considers the position of the corresponding topo-vector and puts in a true value. An obvious optimization to the algorithm is to avoid checking points for which all the associated formulas are already true. Consider the simple model of Figure 2 composed of two closed regions r and q . Since there are two regions, the topo-vector will be of size $5 \cdot 2^2 = 20$ elements: $\{E(r \wedge q), E(r \wedge \neg q), \dots E(\Diamond \Box \neg r \wedge \Diamond \Box \neg p \wedge \Diamond(\neg r \wedge \neg q \wedge \Box \Diamond r \wedge \Box \Diamond q))\}$. After initialization, the region r is considered and one starts looping on the vertices of its polygons, first the point 1. The point is **free**, it is the vertex of a full polygon (not a segment) and therefore the topo-vector is updated directly in the positions corresponding to $Er \wedge \neg q$, $E\Box r \wedge \Box \neg q$, $Er \wedge \neg q \wedge \Box r \wedge \Box \neg q$, $Er \wedge \neg q \wedge \Diamond \Box r \wedge \Diamond \Box q$. The points 2 and 3 would update the values for the same formula and are not considered. The point 4 falls inside the first polygon of r , the topo-vector does not need update. Intersections are then computed and the point 5 is found. The point needs to update the vector for the formula $E\Diamond \Box r \wedge \Diamond \Box \neg q \wedge \Diamond(r \wedge \neg q \wedge \Box \Diamond \neg r \wedge \Box \Diamond \neg q)$. Finally, the point 6 is considered and the point needs to update the formula $E(r \wedge \neg q \wedge \Diamond \Box \neg r \wedge \Diamond \Box \neg q)$. The algorithm proceeds by considering the second region, q and its vertices 7, 8, and 9. The three vertices all fall inside the region r and provide for the satisfaction of the formulas $Er \wedge q$, $E\Box r \wedge \Box q$, $\dots$

Lemma 2 (termination). The topo-distance algorithm terminates.

The property is easily shown by noticing that a segment (a side of a polygon) can have at most one intersection with any other segment, that the number of polygons forming a region of $M_{\mathbb{R}^2}$ is finite, and that the number of regions of $M_{\mathbb{R}^2}$ is finite. Putting this result together with Lemma 1 one gets the hoped decidability result for polygonal topological models.

Theorem 3 (decidability of the topo-distance). In the case of polygonal topological models $M_{\mathbb{R}^2}$ over the real plane, the problem of computing the topo-distance among any two models is decidable.

Given the definition of topo-distance, the fact that two models have a null topo-distance implies that in the topo-game Duplicator has a winning strategy in the infinite round game. In the case of $M_{\mathbb{R}^2}$, Theorem 3 implies that the two models are topo-bisimilar. Note that, in general, this is not the case: Duplicator may have a winning strategy in an infinite Ehrenfeucht-Fraïssé game adequate for some modal language and the models need not be bisimilar [4].

Corollary 1 (decidability of topo-bisimulations). In the case of polygonal topological models over the real plane, the problem of identifying whether two models are topo-bisimilar or not is decidable.

4 The IRIS prototype

The topo-distance is a building block of an image retrieval system, named **IRIS** Image RetrIeval based on Spatial relationships, coded in Java and enjoying a Swing interface (Figure 4). The actual similarity measure is built in **IRIS** to both index and retrieve images on the basis of:

1. the spatial intricacy of each region,
2. the binary spatial relationships between regions, and
3. the textual description accompanying the image.

Referring to Figure 3, one can get a glimpse of the conceptual organization of **IRIS**. A spatial model, as defined in Section 2, and a textual description (central portion of the figure) are associated with each image of the collection (on the left). Each topological model is represented by its topo-distance vector, as built by the algorithm in Section 3 and by a matrix of binary relationships holding between regions. Similarly, each textual description is indexed holding a representative textual vector of the text (right portion of the figure). In Figure 4, a screen-shot from **IRIS** after querying a database of about 50 images of men and cars is shown. On the top-right is the window for sketching queries. The top-center window serves to write textual queries and to attach information to the sketched regions. The bottom window shows the results of the query with the thumbnails of the retrieved images (left to right are the most similar). Finally, the window on the top-left controls the session.

We remark again the importance of moving from games to a distance measure and of identifying the topo-vectors for actually being able to implement the spatial framework. In particular, in **IRIS** once an image is place in the data-base the topo-vector for its related topological model is computed, thus off-line, and it is the only data structure actually used in the retrieval process. The representation is quite compact both if compared with the topological model and with the image itself. In addition, the availability of topo-vectors as indexing structures enables us to use a number of information retrieval optimizations, [12]. In **IRIS**, the similarity consists of three components:

$$sim(I_q, I_j) = \frac{1}{k_n} (k_u^{\text{topo}} \cdot d_{\text{topo}}(I_q, I_j) + k_u^{\text{b}} \cdot d_{\text{b}}(I_q, I_j) + k_u^{\text{text}} \cdot d_{\text{text}}(I_q, I_j))$$

Fig. 3. The conceptual organization of IRIS together with the indexing data structures.

where I_q is the query image (equipped with its topological model and textual description), I_j is the j -th image in the visual database, k_u^{topo} , k_u^{b} , and k_u^{text} are user defined factors to specify the relative importance of topological intricacy, binary relationships and text in the querying process, k_n is a normalizing factor, $d_{\text{topo}}(I_q, I_j)$ is the topo-distance between I_q and I_j , $d_{\text{b}}(I_q, I_j)$ and $d_{\text{text}}(I_q, I_j)$ are the distances for the binary spatial relationships and for the textual descriptions, respectively.

5 Concluding Remarks

We have recalled a general mereotopological framework. We addressed issues of model equivalence and especially of model comparison, thus, looking at mereotopology from a new angle. Defining a distance that encodes the mereotopological difference between spatial models has important theoretical and application implications, in this paper we have focused on the latter. We have shown the actual decidability of the devised similarity measure for a practically interesting class of models.

Having implemented a system based on the above framework is also an important step in the presented research. Experimentation is under way, but some preliminary considerations are possible. We have noticed that the prototype is very sensible to the labeling of segmented areas of images, i.e., to the assignment

Fig. 4. The result of querying a database of men and cars.

of propositional letters to regions. We have also noticed that the mereotopological expressive power appears to enhance the quality of retrieval and indexing over pure textual searches, but the expressive power of $S4_u$ is still too limited. Notions of qualitative orientation, shape or geometry appear to be important, especially when the user expresses his desires in the form of an image query or of a sketch. The generality of the framework described in the paper allows for optimism about future developments.

Acknowledgments

The author is thankful to the anonymous referees for their helpful comments. This work was supported in part by CNR grant 203.7.27.

References

1. M. Aiello. Topo-distance: Measuring the Difference between Spatial Patterns. In G. Brewka and L. Moniz Pereira, editors, *Logics in Artificial Intelligence (JELIA 2000)*, LNAI, 1919, pages 73–86. Springer Verlag, 2000.
2. M. Aiello and J. van Benthem. Logical Patterns in Space. In D. Barker-Plummer, D. Beaver, J. van Benthem, and P. Scotto di Luzio, editors, *Logic Unleashed: Language, Diagrams and Computation*, Stanford, 2001. CSLI.
3. N. Asher and L. Vieu. Toward a Geometry of Common Sense: a semantics and a complete axiomatization of mereotopology. In *IJCAI95*, pages 846–852. International Joint Conference on Artificial Intelligence, 1995.
4. J. Barwise and L. Moss. *Vicious Circles*. CSLI, 1996.
5. B. Bennett. Modal Logics for Qualitative Spatial Reasoning. *Bulletin of the IGPL*, 3:1 – 22, 1995.
6. J. van Benthem. *Modal Correspondence Theory*. PhD thesis, University of Amsterdam, 1976.
7. P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.
8. R. Casati and A. Varzi. *Parts and Places*. MIT Press, 1999.
9. S.K. Chang and S.H. Liu. Picture indexing and abstraction techniques for pictorial databases. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 6(7):475–484, 1984.
10. M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 2000. Second edition.
11. A. Del Bimbo. *Visual Information Retrieval*. Morgan Kaufmann Publishers, 1999.
12. W. Frakes and R. Baeza-Yates, editors. *Information Retrieval, Data Structures and Algorithms*. Prentice Hall, 1992.
13. V. Goranko and S. Pasy. Using the universal modality: gains and questions. *Journal of Logic and Computation*, 2:5–30, 1992.
14. D. Park. Concurrency and Automata on Infinite Sequences. In *Proceedings of the 5th GI Conference*, pages 167–183, Berlin, 1981. Springer Verlag.
15. E. Petrakis, C. Faloutsos, and L. King-Ip. ImageMap: An Image Indexing Method Based on Spatial Similarity. *IEEE Transactions on Knowledge and Data Engineering*, 2001. To appear.
16. I. Pratt and O. Lemon. Ontologies for Plane, Polygonal Mereotopology. *Notre Dame Journal of Formal Logic*, 38(2):225–245, 1997.
17. D. Randell, Z. Cui, and A. G. Cohn. A Spatial Logic Based on Regions and Connection. In *Proc. of Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'92)*, pages 165–176. San Mateo, 1992.
18. S. Santini and J. Ramesh. Similarity measures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(9):871–883, 1999.
19. O. Stock, editor. *Spatial and Temporal Reasoning*. Kluwer Academic Publishers, 1997.
20. H. Tagare, F. Vos, C. Jaffe, and J. Duncan. Arrangement: A spatial relation between part for evaluating similarity of tomographic section. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(9):880–893, 1995.
21. A. Tarski. Der Aussagenkalkül und die Topologie. *Fund. Math.*, 31:103–134, 1938.