

Similarity of Spatial Scenes*

H. Tom Bruns and Max J. Egenhofer

National Center for Geographic Information and Analysis
and
Department of Spatial Information Science and Engineering
University of Maine
Orono, ME 04469-5711, U.S.A.
{tbruns, max}@spatial.maine.edu

Abstract

Similarity is the assessment of deviation from equivalence. Spatial similarity is complex due to the numerous constraining properties of geographic objects and their embedding in space. Among these properties, the spatial relations between geographic objects—topological, directional, and metrical—are critical, because they capture the essence of a scene's structure. These relations can be categorized as a basis for similarity assessment. This paper describes a computational method to formally assess the similarity of spatial scenes based on the ordering of spatial relations. One scene is transformed into another through a sequence of gradual changes of spatial relations. The number of changes required yields a measure that is compared against others, or against a pre-existing scale. Two scenes that require a large number of changes are less similar than scenes that require fewer changes.

1 Introduction

Similarity of spatial scenes is a casual judgment people make frequently in everyday life. It is intuitive, subjective, and displays no strict mathematical models. Geographers have for a long time investigated methods to describe similarity of point sets for spatial analyses, addressing such properties as pattern, density, and dispersion (Unwin 1981). Recently, spatial similarity has found new champions with the advent of content-based image retrieval (del Bimbo *et al.* 1994; Faloutsos *et al.* 1994; Flickner *et al.* 1995) and sketch-based spatial query languages for geographic databases. While the principle remains the same—the quantification of deviations from equivalence for spatial configurations—the methods employed are different in order to accommodate cognitive considerations as well as database processing constraints. This paper focuses on similarity measures that are appropriate for the retrieval of similar spatial configurations, expressed in a spatial query language, from a geographic database. Such formalizations provide an integral of Naive Geography (Egenhofer and Mark 1995b).

An example of a spatial similarity query is a sketch, which approximates reality and often is sufficient for communicating meaning about salient properties of the objects and their geometries. The sketch is similar to reality and represents to the human observer a close approximation of reality. For a spatial database and spatial query processor, however, the sketch is not necessarily equivalent to any stored representation. There are too many spatial variations and differences between the geometries sketched and the geometries present in a geographic

* This work was partially supported by a Massive Digital Data Systems contract sponsored by the Advanced Research and Development Committee of the Community Management Staff and administered by the Office of Research and Development, by the National Science Foundation under grant IRI-9309230, and by Rome Laboratories under grant number F30602-95-1-0042. Max Egenhofer's work is further supported by NSF under grant number SBR-8810917 for the National Center for Geographic Information and Analysis, the Scientific and Environmental Affairs Divisions of the North Atlantic Treaty Organization, Intergraph Corporation, Environmental Systems Research Institute, and Space Imaging Inc.

database. The mere retrieval of the exact match between the sketch and the database content would therefore be insufficient in most cases. The same holds true for the comparison of scenes stored in a database. Rarely will two scenes match exactly: the spatial relations between objects are slightly different; the directions and distances are not exactly the same; and shape, if considered, may vary considerably. As a result, constraining results exact matches would almost always result in empty query results.

Spatial knowledge representation schemes absorb much of the differences among spatial scenes, allowing a focus on the important aspects and thereby facilitating comparison. These abstractions afford spatial reasoning to derive new knowledge from existing knowledge. In computer vision and image processing, the similarity of scenes is based on specific measures of objects. Their locations, shapes, and orientations have been addressed extensively, favoring metrical approaches to determine similarity (Chang and Lee 1991; Uhlmann 1991; Gudivada and Raghavan 1995). The problem addressed here is similarity of spatial scenes contained in geographic databases. A spatial scene is a set of geographic objects together with their spatial relations—topological relations, distance relations, and direction relations—and optionally other types of spatial characteristics, including such unary object descriptors as shape, ratio-type relations such as relative sizes (areas and lengths), or attributes specifying the semantics of the spatial objects. Assessing the similarity between different scenes involves comparing individual tuples or sets of tuples in the geographic database.

To retrieve the most similar configurations to a target scene, this paper employs the concept of *gradual change*, which imposes an order on sets of spatial relations. Gradual change originates from the gradual deformation of spatial objects until the spatial relation between them is changed. It applies to topological relations (Egenhofer and Al-Taha 1992) as well as to cardinal directions (Freksa 1992) and approximate distances (Sharma 1996). Two spatial relations that require little change to deform one into the other are more similar than relations that require more change. By extension, spatial scenes that require little change to deform the topological, distance, and direction relations of one into the other, are more similar than scenes that require more change.

The remainder of this paper discusses the design of the conceptual neighborhoods for the three types of spatial relations being considered: topology (Section 2), distance (Section 3), and direction (Section 4). The core is the process by which these concepts are integrated into an assessment of spatial similarity (Section 5). The paper concludes with comments on future work.

2 Similarity of Topological Relations

Topological relations often capture the essence of a spatial configuration—topology matters, metric refines (Egenhofer and Mark 1995b). Topological constraints are attractive for assessing similarity as they are largely immaterial to subtle geometric variations and when they get changed, usually significant alterations occur. If several of such significant changes occur, a chain reaction gets triggered. Initially two relations are slightly changed, or still just one, only a bit more dramatically. The new scene is still similar, only less so. As the number and extent of the changes increases, the new scene becomes less and less similar. The change is gradual, from equivalent, to highly similar, to less and less similar.

This concept of *gradual change* has been used to model conceptual neighborhoods of topological relations (Egenhofer and Al-Taha 1992; Egenhofer and Mark 1995a). Conceptual neighborhoods facilitate an ordering of topological relations, and support the determination of similar relations. Figure 1 shows the eight topological relations for simple areal objects (Egenhofer and Franzosa 1991). Relations connected by a line in the figure represent conceptual neighbors.

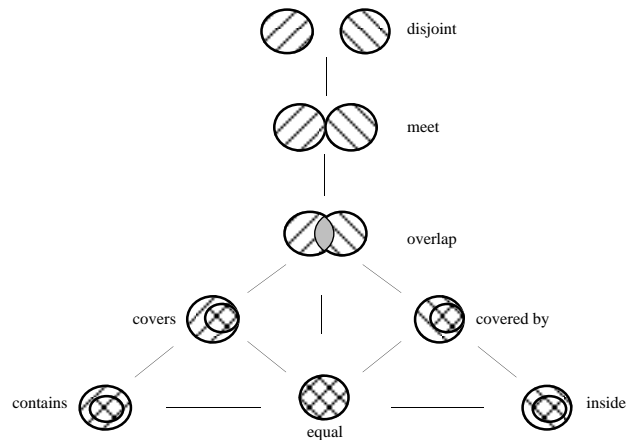


Figure 1: Conceptual neighborhood of topological relations.

The application of some gradual, spatially consistent process can change a relation into any of its conceptual neighbors. Using this concept, statements like “meet is similar to overlap” or “meet is more similar to overlap than to contains” are made.

Topological similarity may have little meaning beyond simple, two-object scenes. Its real utility comes with a combination of equivalent concepts for distance and direction relations, and scenes with more than two objects.

3 Similarity of Distance Relations

Qualitative distance relations are difficult to define for general spatial objects. The concepts and terms used are highly subjective, and sensitive to the scale of the spatial data being considered. Irregular shapes introduce special cases that confound most defined distance relations. Figure 2 shows one possible distance relation scheme based on increasing buffer distances. Objects range from having no distance between them, *zero*, to being *very close*, to *close*, and then *far*. The actual definitions of these distances have been investigated elsewhere (Hong 1994; Hernández *et al.* 1995).

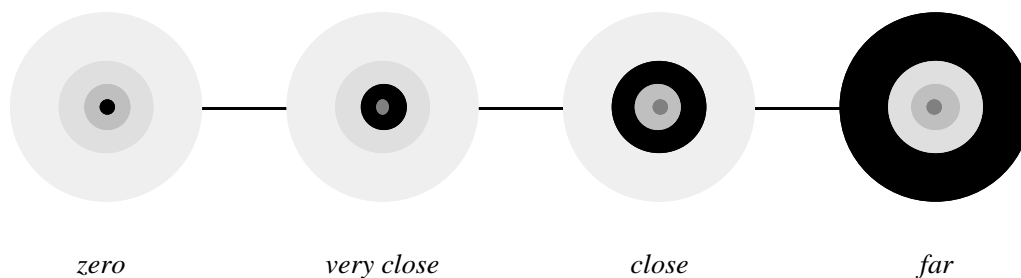


Figure 2: Conceptual neighborhoods of distance relations, and the symbols that represent them.

For such distances, conceptual neighborhoods are derived by imposing an order relation $<$ (less than) over the distance symbols, which corresponds to two objects gradually moving away from each other. Adjacent symbols are more similar than non adjacent symbols. For example, *very close* is more similar to *close* than to *far*, because $zero < very\ close < close < far$. Transitivity applies to this order relation, supporting such statements as *far* is greater than *very close*. This type of reasoning supports the determination of the difference in spatial relations between two scenes, which can guide the process of determining the number and type of gradual changes required to transform one scene into another, which forms the basis of the similarity assessment presented here.

4 Similarity of Direction Relations

Cardinal directions describe, qualitatively, the orientation between spatial objects (Frank 1991). Figure 3 shows a conceptual neighborhood diagram of a subset of the 169 possible spatial relations between rectangles (Chang and Lee 1991). Such schemes have been formally derived from Allen's interval relations applied to orthogonal projections (Sharma and Flewelling 1995). The subset shown depicts the relations between two equally-sized squares. Neighborhood diagrams for other shapes can be generated similarly, but they are more difficult to depict and decipher in the 2-dimensional plane. The gradual changes accommodated by these neighborhoods are translations of either object in any direction. The directions on the outside edge of the diagram are labeled with acronyms such as N for "north," SWS for "south-southwest," etc. Direction relation schemes are highly sensitive to the orientation of the objects in a scene, as well as the scene itself. Still, any direction scheme can be arranged into conceptual neighborhoods, which is the important concept here.

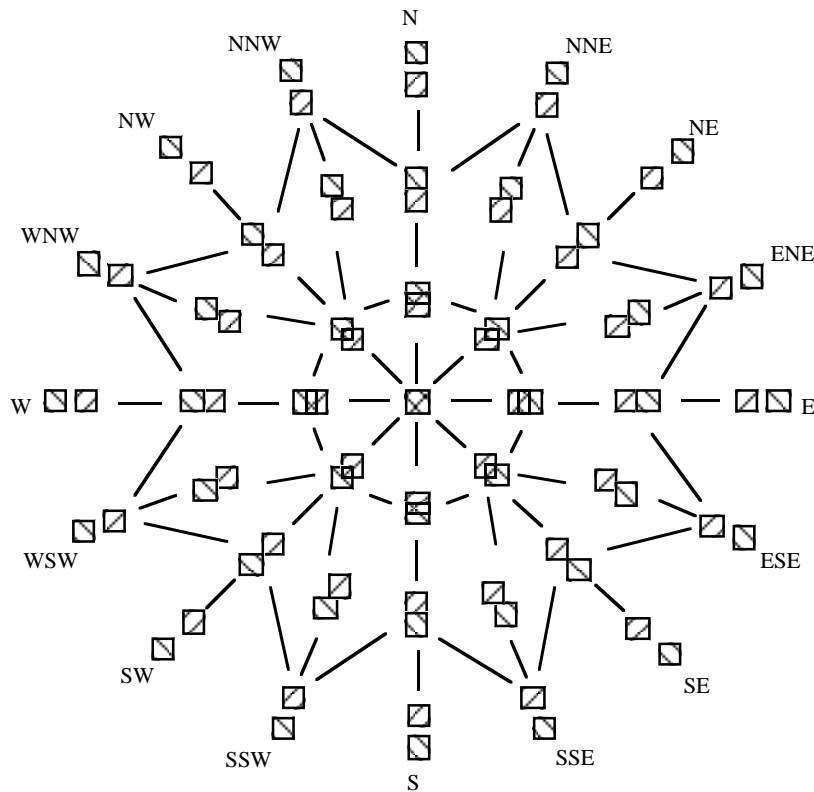


Figure 3: Conceptual neighborhoods of direction relations for same-size square objects.

Figure 3 illustrates how to produce, by gradual change, any direction relation given any other direction relation. It is assumed that each object can move in a continuous, smooth manner in any direction, but may not suddenly jump to a new location. (Note, however, that all relations can be represented by fixing one object, and moving the other.) The lines in the figure show the links between conceptual neighbors of direction relations. If two relations are not conceptual neighbors, then one cannot be derived from the other via gradual change without first producing one or more intermediate relations. The assumption made here is that the number of intermediate relations is proportional to the similarity of relations. The similarity of scenes of multiple relations is somehow proportional to the sum of the similarities of the individual relations.

The structure of the figure shows the interplay between direction and topological relations. There are four rings, each representing a topological relation. The center ring represents the

topological relation equal. The next ring out shows all direction relations for the topological relation overlap. Further out are the rings for meet and disjoint, in that order. Additional rings would represent associations between direction and distance relations, while preserving topology (disjoint) and the corresponding direction relation.

5 Assessing Spatial Similarity

The three models of topological relations, cardinal direction, and approximate distances form the basis for the assessment of spatial similarity of scenes. Given two scenes with an equal number of spatial objects, and different spatial relations, there exists a minimum set of gradual changes that will transform one scene into the other. The number of gradual changes in this process is considered to be proportional to the similarity of the two scenes. The more steps required, the less similar the scenes are.

5.1 Counting the number of different spatial relations

In order to determine the steps between two scenes, one may evaluate both scenes' binary spatial relations and count their differences. Figure 4 shows such an analysis. The only difference between Figure 4a and 4b is that one object, object C, has been moved.

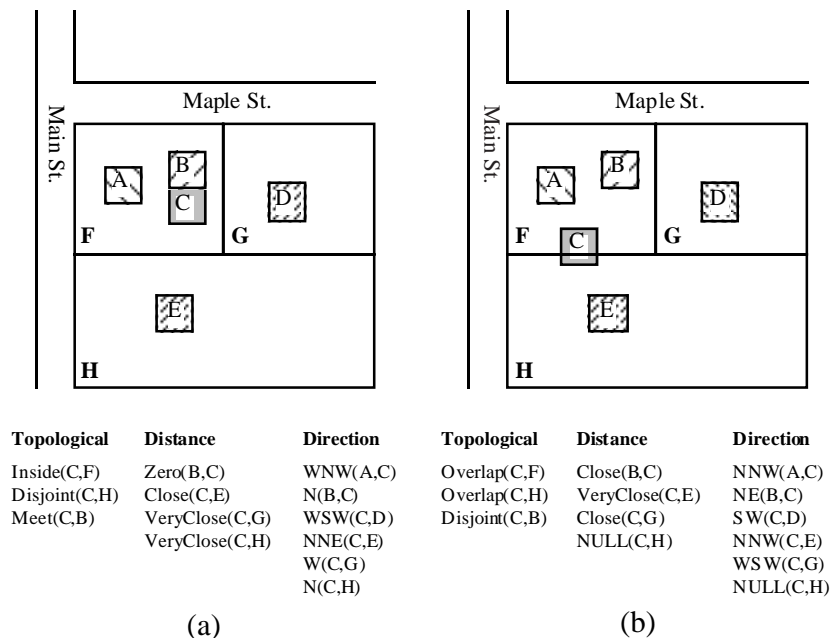


Figure 4: Two similar scenes and the differences in their databases.

For even this simple scene, such a slight change can have a dramatic impact on the set of spatial relations. The similarity of the two scenes could be assessed as the number of different relations, and the degree of difference between the relations, but it is unclear how the individual assessments would be combined, and what the priorities are for the different types of spatial relations. More important, since spatial relations are often highly correlated, it would be too simplistic to consider each of them in isolation.

5.2 Counting the gradual changes spatial relations

An alternative is to determine a set of gradual changes that affect the observed differences in the spatial scenes. For example, consider the scene in Figure 5a (a sub-scene of Figure 4a) and how it might be transformed into the scene in Figure 5b (a sub-scene of Figure 4b).



Figure 5: How is scene (a) transformed into scene (b)?

There are differences in distance, direction, and topological relations, and the goal is to determine the fewest number of gradual changes required to transform the start scene (Figure 5a) into the target scene (Figure 5b). In general, the set of gradual changes required, and the order in which they are applied, is unknown. The approach, then, is to systematically apply all possible gradual changes to a scene, producing a set of possible new scenes. Each new scene is the result of applying one gradual change. Each gradual change results from replacing one spatial relation, distance, direction, or topological, with one of its conceptual neighbors. For each conceptual neighbor for each type of relation, a new possible scene is produced. This process is recursively applied to all the derivative scenes, until the destination scene is among the new derivatives. The number of steps required in this process is the measure for similarity.

Figure 6 shows the results of applying this process once to the scene in Figure 5a. The start scene, Scene 0, is in the center of Figure 6. The scenes surrounding the start scene were produced by applying a single gradual change of some spatial relation. Scenes 1 through 6 result from a gradual change in a direction relation. Scenes 7 and 8 result from a gradual change in a distance relation. And scenes 9 through 12 result from a gradual change in a topological relation. Each scene results from a single gradual change, so each is said to be equally similar to the original scene.

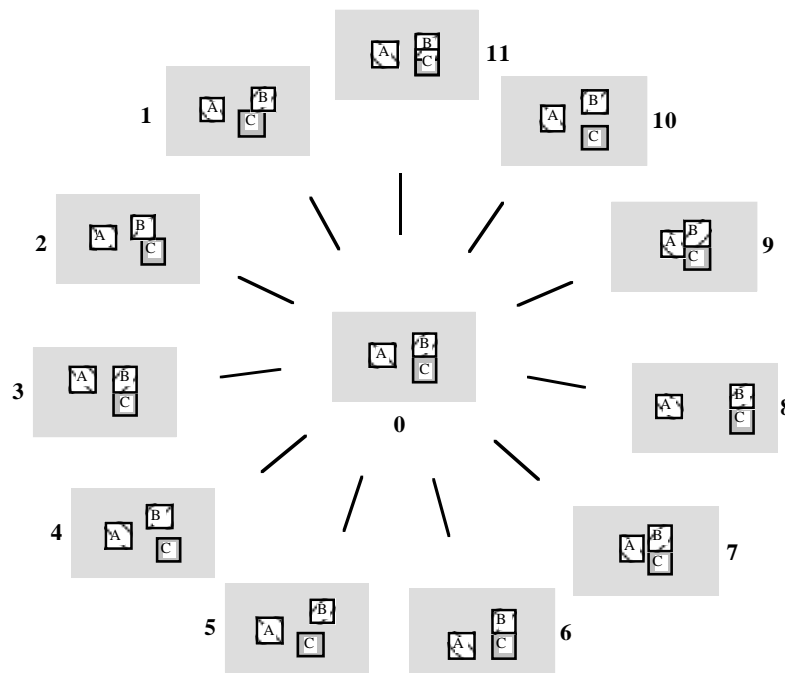


Figure 6: Derivative scenes resulting from applying a single gradual change to the center scene from Figure 5a.

This process continues by taking each of the derivative scenes from Figure 6 and again applying all possible gradual changes. It is limited so as not to reproduce any scene visited before. Eventually, a scene network is formed, which corresponds to a partially ordered set.

Figure 7 shows such a network for a two object scene, starting with the objects meeting, and one object strictly north of the other. All scenes at each level of the partially ordered set are considered equally similar to the original scene.

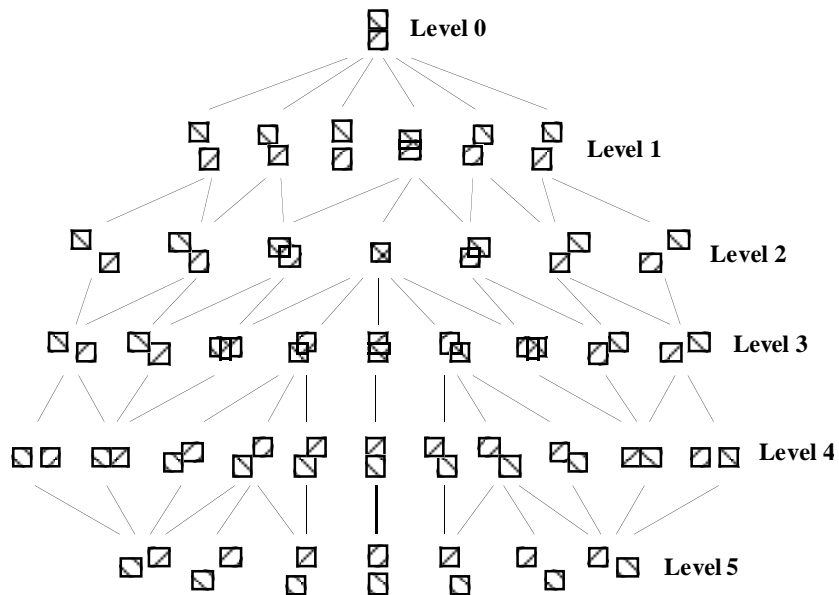


Figure 7: Similarity network.

This search process can be thought of as dividing the space of all possible scenes into regions of equal similarity. Figure 8 shows the conceptual neighborhood diagram of Figure 3 divided into 5 similarity levels. The lines between the levels are *similarity contours* in scene-space.

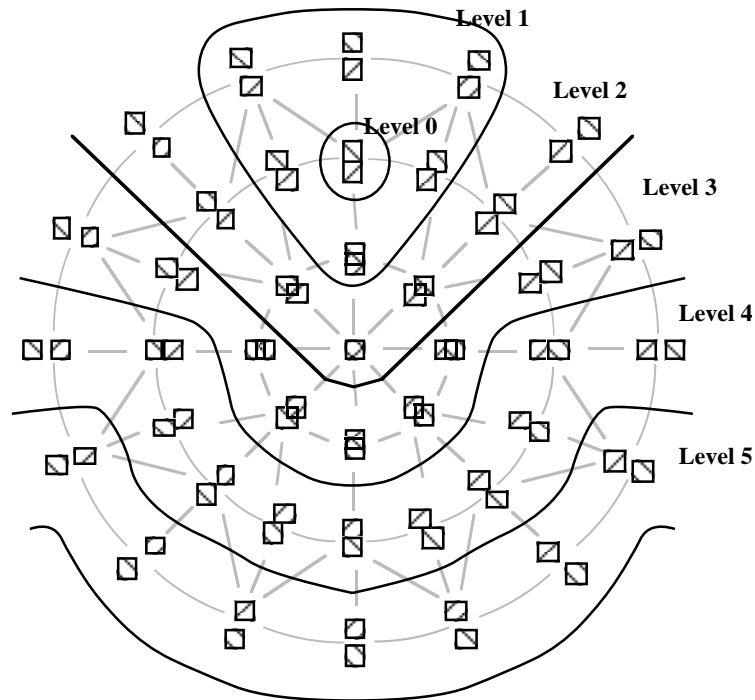


Figure 8: Similarity contours in scene-space indicating levels of equal similarity.

Levels of similarity can be used to directly assess the *relative* similarity of a set of scenes. One scene is compared against two or more others, which are ranked according to the similarity level at which they occur. In Figure 9, scene (a) is being compared against three other scenes (b-d). To produce the scene in Figure 9b from Figure 9a, one gradual change is required, so the scene is ranked at level 1. For the scene in Figure 9c, two gradual changes are needed, putting the scene at level 2. It takes four gradual changes to achieve Figure 9d, placing it at level 4. Given these level assessments, it can be said that scene 9b (Figure 9b) is more similar to scene 9a than scene 9c, and scene 9c is more similar to 9a than 9d.

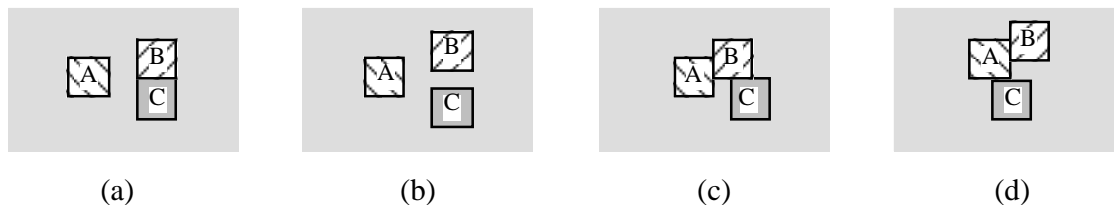


Figure 9: Relative similarity.

Within any similarity level there may be many scenes, and it would be useful to be able to establish a ranking within a level as well. Note that not all gradual changes will have an equal impact on a scene database. This impact, the number and type of changed spatial relations, could be used to rank the similarity of scenes within similarity level.

6 Conclusions and Future Work

This paper presents a method for assessing spatial similarity. It is based on the topological, directional, and distance relations in a spatial scene, and their conceptual neighborhoods. Similarity is determined by a process that gradually replaces spatial relations in a scene by their conceptual neighbors, in an attempt to construct one scene from another. Derivative scenes are ordered into levels by the number of required gradual changes. Scenes that require fewer changes, which are therefore placed at a lower level, are said to be more similar than scenes requiring more changes.

There are numerous aspects of scene similarity that have not been discussed in this paper. Most obvious is the need for testing whether the model chosen matches with human intuition, in the tradition of earlier evaluations and calibrations of natural-language spatial relations (Mark *et al.* 1995). Two formal aspects of this research are of immediate concern: the derivation of conceptual neighbors for the full set of 169 spatial relations, and a ranking scheme for scenes at a given level of similarity. The types of gradual changes need to be expanded to accommodate more realistic scenes. The ability to gradually change size, rotation, and add or remove objects, holes, and other features is needed. Eventually, this research will consider more complex scenes, with line-region relations, objects with holes, islands, and irregular shapes.

The similarity of the details of spatial relationships will also be investigated. For example, topological invariants beyond the emptiness/non-emptiness of the 9-intersection have been used to model the details of a complex topological relation (Egenhofer and Franzosa 1995). What are the conceptual neighbors of a set of topological invariants? Consider the *overlap* relation at the top of Figure 10. A possible set of conceptual neighbors could be those that introduce another basic topological invariant, such as a one-dimensional *meet*. The bottom of Figure 10 shows the different ways such a *meet* could be added to the simple *overlap*. Other deformations would have to be defined such that any set of topological invariants could be produced from any other set. A search similar to the one discussed here would then yield the smallest set of deformations required to transform one scene to another. Further refinements based on the amount of overlapping areas are complementary, taking over the role of distance relations for disjoint relations.

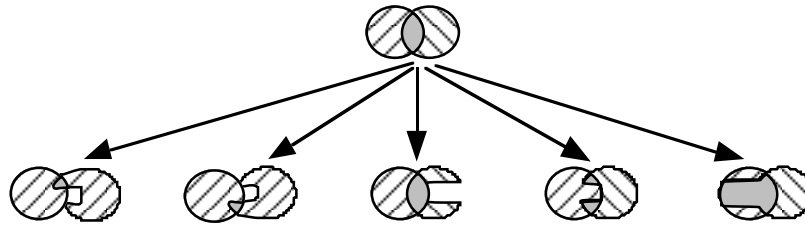


Figure 10: Derivative scenes resulting from adding a meet to a single-overlap scene.

A related topic is the generation of similarity queries, given some initial query. In Figure 11, a user has posed the query on the left side. This may or may not result in any database matches. The concepts of gradual change and conceptual neighborhoods can be used to modify this query. Each specification is relaxed to include a disjunction of its conceptual neighbors. Any combination of specifications describes, generally speaking, a similar scene. However, there may be some illegal combinations, and consistency constraints would have to be applied first.

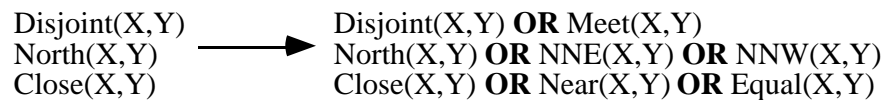


Figure 11: Fully relaxed similarity query.

Similarity assessments based on the relaxation of spatial relations is an promising alternative to the purely quantitative methods applied in image retrieval. Relation-based similarity stresses qualitative aspects, and allows for incrementally more detailed measures where necessary. Ultimately, such methods will need to be complemented by qualitative shape descriptions and analyses of the semantics of the objects involved in order to provide intelligent, computer-based spatial reasoning.

7 References

- C. Chang and S. Lee (1991) Retrieval of Similar Pictures on Pictorial Databases. *Pattern Recognition* 24(7): 675-680.
- A. del Bimbo, E. Vicario, and D. Zingoni (1994) A Spatial Logic for Symbolic Description of Image Content. *Journal of Visual Languages and Computing* 5: 267-286.
- M. Egenhofer and K. Al-Taha (1992) Reasoning About Gradual Changes of Topological Relationships. in: A. Frank, I. Campari, and U. Formentini (Eds.), *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space. Lecture Notes in Computer Science* 639, pp. 196-219, Springer-Verlag, New York.
- M. Egenhofer and R. Franzosa (1991) Point-Set Topological Spatial Relations. *International Journal of Geographical Information Systems* 5(2): 161-174.
- M. Egenhofer and D. Mark (1995a) Modeling Conceptual Neighbourhoods of Topological Line-Region Relations. *International Journal of Geographical Information Systems* 9(5): 555-565.
- M. Egenhofer and D. Mark (1995b) Naive Geography. in: A. Frank and W. Kuhn (Eds.), *Spatial Information Theory—A Theoretical Basis for GIS, International Conference COSIT '95, Semmering, Austria. Lecture Notes in Computer Science* 988, pp. 1-15, Springer-Verlag, Berlin.
- M. Egenhofer and R. Franzosa (1995). On the Equivalence of Topological Relations. *International Journal of Geographical Information Systems* 9(2): 133-152.
- C. Faloutsos, R. Barber, M. Flickner, J. Hafner, W. Niblack, D. Petrovic, and W. Equitz (1994) Efficient and Effective Querying by Image Content. *Journal of Intelligent Information Systems* 3: 231-262.

- M. Flickner, H. Sawhney, W. Niblack, J. Ashley, Q. Huang, B. Dom, M. Gorkani, J. Hafner, D. Lee, D. Petkovic, D. Steele, and P. Yanker (1995) Query by Image and Video Content: The QBIC System. *IEEE Computer* 28(9): 23-32.
- A. Frank (1991) Qualitative Spatial Reasoning about Cardinal Directions. in: D. Mark and D. White (Eds.), *Autocarto 10*, Baltimore, MD, pp. 148-167.
- C. Freksa (1992) Using Orientation Information for Qualitative Spatial Reasoning. in: A. Frank, I. Campari, and U. Formentini (Eds.), *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space. Lecture Notes in Computer Science* 639, pp. 162-178, Springer-Verlag, New York, NY.
- V. Gudivada and V. Raghavan (1995) Design and Evaluation of Algorithms for Image Retrieval by Spatial Similarity. *ACM Transactions on Information Systems* 13(2): 115-144.
- D. Hernández, E. Clementini, and P. di Felice (1995). Qualitative Distances. in: A. Frank and W. Kuhn (Eds.), *Spatial Information Theory—A Theoretical Basis for GIS, International Conference COSIT '95, Semmering, Austria. Lecture Notes in Computer Science* 988, pp. 45-57, Springer-Verlag, Berlin.
- J.-H. Hong (1994). *Qualitative Distance and Direction Reasoning in Geographic Space*, Ph.D. Dissertation, Department of Surveying Engineering, University of Maine.
- J. -H. Hong, M. Egenhofer, and A. Frank (1995). On the Robustness of Qualitative Distance- and Direction Reasoning. in: D. Peuquet (Ed.) *Autocarto 12*, Charlotte, NC, pp. 301-310.
- S. Lee, and F. Hsu (1990). 2D C-String: A New Spatial Knowledge Representation for Image Database Systems. *Pattern Recognition* 23(10): 1077-1087.
- D. Mark, D. Comas, M. Egenhofer, S. Freundschuh, M. Gould, and J. Nunes (1995) Evaluating and Refining Computational Models of Spatial Relations Through Cross-Linguistic Human-Subject Testing. in: A. Frank and W. Kuhn (Ed.), *Spatial Information Theory—A Theoretical Basis for GIS, International Conference COSIT '95, Semmering, Austria. Lecture Notes in Computer Science* 988, pp. 553-568, Springer-Verlag, Berlin.
- J. Sharma, J. and D. Flewelling (1995). Inferences from Combined Knowledge about Topology and Directions. in: M. Egenhofer and J. Herring (Eds.) *Advances in Spatial Databases—4th International Symposium, SSD '95, Portland, ME*. Berlin, Springer-Verlag. 951: 279-291.
- J. Sharma (1996) *Heterogeneous Spatial Reasoning in Geographic Databases*, Ph.D. Thesis, University of Maine (forthcoming).
- J. Uhlmann (1991) Satisfying General Proximity/Similarity Queries with Metric Trees. *Information Processing Letters* 40(4): 175-179.
- D. Unwin (1981) *Introductory Spatial Analysis*. Methuen & Co, New York, NY.