# Multi-Modal Spatial Querying[*]

**Max J. Egenhofer**

National Center for Geographic Information and Analysis, Department of Spatial Information Science and Engineering, and Department of Computer Science, University of Maine, Orono, ME 04469-5711, USA, max@spatial.maine.edu

## Abstract

People who use multiple channels at the same time communicate more successfully about spatial problems than those who rely exclusively on either voice or pictures. To achieve a similarly successful interaction between a person and a geographic information system (GIS), we use two concurrent communication channels—graphics and speech—to construct a multi-modal spatial query language in which users interact with a geographic database by drawing sketches of the desired configuration, while simultaneously talking about the spatial objects and the spatial relations drawn. Through the combined use of graphics and sketch, more intuitive and more precise specifications of spatial queries are possible. The key to this interaction is the exploitation of complementary or redundant information present in both graphical and verbal descriptions of the same spatial scenes. A multiple-resolution model of spatial relations is used to capture the essential aspects of a sketch and its corresponding verbal description. The model stresses topological properties, such as containment and neighborhood, and considers metrical properties, such as distances and directions, as refinements where necessary. This model enables the retrieval of similar, not only exact, matches between a spatial query and a geographic database. Such new methods of multi-modal spatial querying and spatial similarity retrieval will empower experts as well as novice users to perform easier spatial searches, ultimately providing new user communities access to spatial databases.

## 1. Introduction

Today's methods of interacting with geographic databases are largely non-spatial, as they require their users to deal with geographic data primarily through alphanumeric command languages. Currently, spatial querying is done by typing a command in some spatial query language, such as an extended version of SQL (Ingram and Phillips 1987; Herring *et al.* 1988; Egenhofer 1994), or by selecting the same or a similar syntax through a forms interface or from pull-down menus (Egenhofer 1990; Calcinelli and Mainguenaud 1994; Aufaure-Portier 1995). Such spatial querying is a tedious process, because it often requires extensive training in the use of the particular query language. A more serious disadvantage of such textual spatial querying is that it forces users to translate a spatial image they may have in their minds about the situation they are interested in, into a non-spatial language. Graphical user interfaces provide only little improvement for such query languages, because they use the same type of syntax and grammar as the typed languages, and they only release users from remembering the particular syntax (Egenhofer 1992). The problems with communicating a user's request to a spatial database through conventional spatial query languages become most apparent when several users have to work together and have to understand their intentions. Verbal descriptions of spatial situations are frequently ambiguous and may easily lead to misinterpretations, particularly in multi-language working groups. Traditional spatial query languages have serious limitations when geographic concepts are used that are vague, imprecise, little understood, or not standardized. As an example, take the notion of the spatial predicate "cross" whose semantics may vary depending on the context in which it is used, the meaning of

---

[*] This work was partially supported by Rome Laboratories under grant number F30602-95-1-0042. Max Egenhofer's research is further supported by grants from the National Science Foundation under grant number No. SES 88-10917 for the NCGIA and grant number IRI-9309230; the Scientific and Environmental Division of the North Atlantic Treaty Organization; Intergraph Corporation; Environmental Systems Research Institute Inc.; Space Imaging, Inc., and 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.

M. Egenhofer, Multi-Modal Spatial Querying, in: J.-M. Kraak and M. Molenaar (eds.), Seventh International Symposium on Spatial Data Handling, Delft, The Netherlands, Taylor & Francis, London, pp. 785-799.

the objects the predicate relates to, and the topology and the metric of the particular configuration. These drawbacks make current spatial query languages error-prone and difficult to use.

Sketch-and-Talk is an innovative spatial query language that uses simultaneously graphic and voice input. It is made possible by the advent of pen-based computers and is a response to the increased interest in Mobile Computing (Imielinski and Korth 1995). Such a multi-modal spatial query language will allow users to interact more intuitively with spatial data than traditional GIS query languages and GIS user interfaces do, because it supports more directly human spatial thinking and familiar human interaction techniques through the combination of graphical (“sketch”) and voice (“talk”) input.

Sketch-and-Talk aims at retrieving from a spatial database those configurations that match a set of constraints specified. In this process, spatial relations play a significant role (Frank 1982; Pizano *et al.* 1989; Egenhofer 1992; Papadias and Sellis 1994) as they often specify the principal constraints about the data to be retrieved. Sketch-and-Talk addresses the particular need for modeling semantics of visual information, which participants at a recent NSF-ARPA workshop identified as essential for the success of Visual Information Management Systems.

> “*… currently a major bottleneck are, however, the techniques to introduce and manage the semantics in these [visual information] systems. … The queries require assignment of semantics to data.*” (Jain and Pentland 1995)

Sketch-and-Talk allows users to choose their favorite interaction mode to compose a spatial query and the query processor integrates the two representations into a canonical form. The redundancy in the voice and graphics mode is critical information for the information system’s query processor as it contributes to solving *spatial incompleteness and spatial ambiguities* that may exist in one mode through specifications in the other mode; detecting *spatial contradictions* among the different modes, therefore, saving precious time when trying to process a query that the user did not specify precisely enough; and exploiting *spatial consensus* among the different modes to determine more reliable input.

Sketch-andTalk is motivated by findings in Human-Computer Interaction where a combined speech and gesture interaction mode for the manipulation of graphic images was preferred over a pure gesture or a pure speech interaction (Hauptmann 1989), as well as results from cognitive science and linguistics, where researchers argued that natural language descriptions of spatial configurations absorb detail, while graphics are generally overspecified (Talmy 1983). For example, a natural-language description like, “the road that enters the park” addresses the salient properties of the configuration—the principal topological concept of getting into without specifying such details as where the road starts or how often it crosses the park’s boundary—while it abstracts away other geometric characteristics such as the shape of the objects, their sizes, and the orientation among them. This leaves ample freedom for topological and metric variations and these geometric details remain unspecified when processing such a description as a query against a spatial database. As an alternative to natural-language descriptions, sketches may constrain more precisely a particular spatial situation, but they introduce additional spatial properties that the user did not necessarily intend to specify, such as the objects’ shapes, their relative sizes, and directions. When parsing a sketch it becomes impossible to distinguish explicit from implicit spatial constraints. Through the combined use of sketch and voice queries, Sketch-and-Talk balances between underspecifications and overspecifications by giving different priorities to the details of a drawing and a corresponding natural-language description.

This paper discusses the design of Sketch-and-Talk. After a review of related work (Section 2), we give in Section 3 a guided tour through a simple interaction with Sketch-and-Talk. Section 4 presents the model used for integrating verbal and sketched spatial relations, and Section 5 introduces how Sketch-and-Talk queries get processed. Conclusions in Section 6.

## 2. Related Work

Sketch-and-Talk is complementary to other spatial search mechanisms (Kuhn 1992) such as spatial data mining to discover interesting spatial configurations (Koperski and Han 1995), and browsing to explore a spatial data set for selection (Clementini *et al.* 1990). Sketch-and-Talk supports spatial similarity retrieval (Chang and Lee 1991) and complements methods of content-based image retrieval (Hirata and Kato 1992; Faloutsos *et al.* 1994; Flickner *et al.* 1995; Ogle and Stonebraker 1995) and specifying spatial consistency constraints graphically (Pizano *et al.* 1989). Although Sketch-and-Talk employs pen-based interaction with spatial data, it differs from sketch-based CAD languages that employ sketching to construct new configurations (Borning 1986) or update spatial information systems (White 1988; Kuhn 1990), including 3-D drawing systems (Deering 1995). Geometric construction and updating specify a particular configuration and require precise geometric positioning, whereas sketching a spatial query describes a prototypical configuration from which valid query results may deviate by relaxing spatial constraints (“best fit”).

Sketching for querying was used in Query by Visual Example (Hirata and Kato 1992) and Query by Image Content (Faloutsos *et al.* 1994), which are targeted for content-based image retrieval. While the interaction mode of these query languages is similar to some of the basics of Sketch-and-Talk—in both cases users draw an approximate spatial configuration of what to retrieve—scope, modalities, and sketch interpretation are considerably different. Sketches for content-based image retrieval assume that the user draws something that matches quite closely the target and that all relations are intended as drawn. Their query processors accommodate primarily metric variations (Del Bimbo *et al.* 1994) and they are very sensitive to variations in sizes, orientations, and shapes. On the other hand, Sketch-and-Talk assumes that the user’s sketch and the targets may vary, at times considerably, as long as they match in the most important criteria.

The semantics of spatial relations, essential for the interpretation and processing of a Sketch-and-Talk query, has received attention primarily in the arena of linguistics and artificial intelligence. Clark (1973) suggested a strong correspondence between Perceptual Space, which humans use to perceive the space around them, and Linguistic Space, which is used by language to represent the perceived space. This correspondence has been widely used in subsequent research for eliciting natural language descriptors of scenes. Talmy’s (1983) seminal paper on “How Language Structures Space” establishes the link between prototypical spatial configurations and the use of natural language predicates. According to Talmy, at the fine-structural level of conceptual organization, language shows greater affinity with topology than with metric spaces. Work on metaphors (Lakoff and Johnson 1980) and people’s conceptualizations of spatial relations (Grimaud 1988; Japkowicz and Wiebe 1991) are fundamental to the study of human cognition of spatial relations. We consider them to be complimentary to our approach and do not explicitly use these ideas of conceptualization in our work, because our formalism is focused on semantics that can be captured from the geometric configuration of spatial relations. Positive results obtained from our earlier work with human subject testing (Mark and Egenhofer 1994b) justify our assumption for taking this approach. Our study also relates to Rosch’s (1978) general theory of human categorization whereby prototypical cases are used and objects are specified in terms of their distances from these prototypes. We are, however, dealing with spatial relations and not objects. In this respect, we are building on Herskovits’s (1986) work, which used Rosch’s method of prototypical categorization and applied it to spatial relations.

## 3. A Sample Scenario

The following scenario provides a rough outline of the interaction a user may perform when sketching and simultaneously talking about a query. Sketch-and-Talk uses a touch-sensitive input device—ideally a touch screen with a pen, such as Apple’s Newton, or alternatively a tablet with a pen—plus a microphone to record the user’s voice (Figure 1).

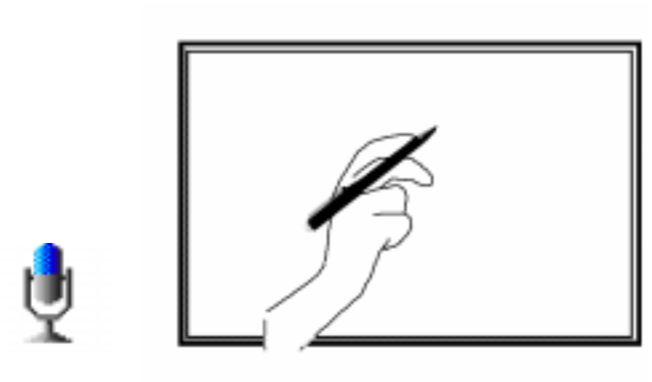

**Figure 1**: Sketch-and-Talk with a pen-based user interface and a microphone.

The user employs a pen to sketch an example of what she wants to find in the database. In this particular case, the user is interested in land parcels with a wooded area on the Penobscot River. While the user starts describing her request verbally (“***I’m interested in a land parcel …***”), she draws the boundary of a land parcel (Figure 2). Sketch-and-Talk parses that the object drawn is a land parcel and adds the label to the sketch.

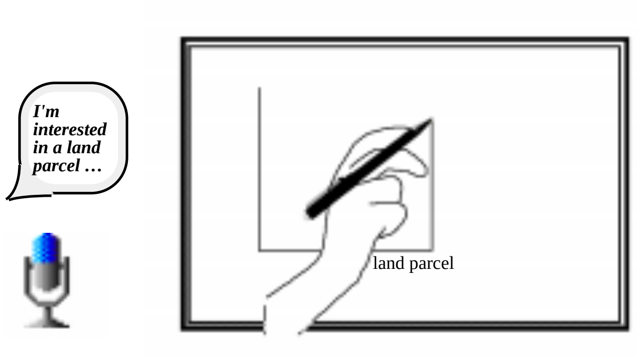

**Figure 2**: The user draws the outline of a land parcel, while describing that the object of interest is a land parcel.

The user continues the query by specifying that the land parcel should be *“**… with some forest on it …**”* and simultaneously drawing part of the forest’s boundary to indicate its location with respect to the land parcel (Figure 3). Again, the new object drawn is labeled with its type.

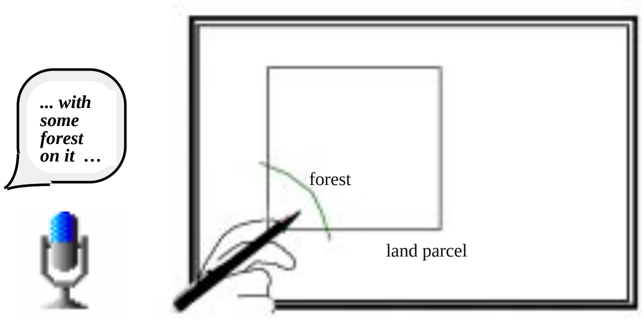

**Figure 3**: The user adds the boundary of a forest such that it intersects with the land parcel.

M. Egenhofer, Multi-Modal Spatial Querying, in: J.-M. Kraak and M. Molenaar (eds.), Seventh International Symposium on Spatial Data Handling, Delft, The Netherlands, Taylor & Francis, London, pp. 785-799.

Since it is unclear on which side of the line the forest is located, the user gestures with the pen by filling the interior of the forest (Figure 4).

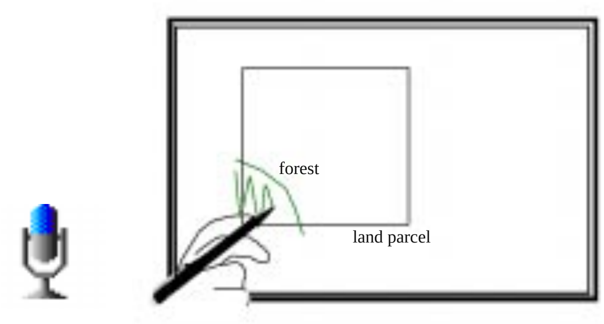

**Figure 4**: To determine the location of the forest (i.e., on which side of the boundary it lies), the user fills the scribbles into the forest's interior.

Finally the user draws a line that crosses the land parcel, while completing the constraint that the land parcel must be "***… located on the Penobscot River***" (Figure 5).

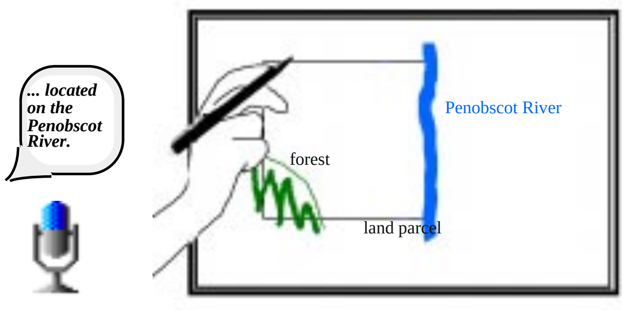

**Figure 5**: The user adds the location of Penobscot River such that runs along the land parcel, but does not intersect with the forest.

## 4. Modeling Spatial Relations for Multi-Modal Querying

A critical component for the success of Sketch-and-Talk is the integration of what the user drew and what she described verbally. This integration must occur at a semantic level such that the meanings of the sketch and the verbal description can be compared. Our approach to this integration is based on the use of a canonical representation for Sketch-and-Talk queries, which is compatible with natural-language spatial relations as well as graphics. Sketch-and-Talk bases its analysis of spatial relations on the premise *Topology Matters, Metric Refines* (Egenhofer and Mark 1995b). The remainder of this section reviews the pieces of the model used as a symbolic representation of spatial relations.

### 4.1 9-Intersection for Topological Relations

We base the analysis of spatial relations on the 9-intersection, which is a comprehensive model for binary topological spatial relations. It applies to objects of type area, line, and point (Egenhofer and Herring 1990; Egenhofer and Franzosa 1991) and characterizes the topological relation between two point sets, $A$ and $B$, by the set intersections of $A$'s interior ($A^\circ$), boundary ($\partial A$), and exterior ($A$ ) with the interior, boundary and exterior of $B$ (Equation 1). With each of these nine intersections being empty ($\varnothing$) or non-empty ($\neg\varnothing$), the model has 512 possible topological relations between two point sets, some of which cannot be realized. For two simple regions

M. Egenhofer, Multi-Modal Spatial Querying, in: J.-M. Kraak and M. Molenaar (eds.), Seventh International Symposium on Spatial Data Handling, Delft, The Netherlands, Taylor & Francis, London, pp. 785-799.

without holes embedded in $R^2$, the categorization shows eight distinct topological relations. They have been called *disjoint*, *meet*, *equal*, *overlap*, *inside*, *contains*, *covers*, and *coveredBy*. For two simple lines (non-branching, no self-intersections) embedded in $R^2$, 33 different topological relations can be realized with the 9-intersection, and for a line and a region, 19 different situations are found (Egenhofer and Herring 1990).

$$I(A,B)=\begin{pmatrix} A^\circ \cap B^\circ & A^\circ \cap \partial B & A^\circ \cap B^- \\ \partial A \cap B^\circ & \partial A \cap \partial B & \partial A \cap B^- \\ A^- \cap B^\circ & A^- \cap \partial B & A^- \cap B^- \end{pmatrix} \tag{1}$$

We use the 9-intersection relations as the key for analyzing spatial relations sketched, because it captures topological relations at a coarse level and, therefore, is an appropriate candidate for grouping sketches into classes of similar relations (Figure 6). By mapping the sketched relations onto 9-intersection relations, we capture the most salient features of a sketch in a form that is independent of orientations and sizes. This abstraction is critical to translate a sketched configuration into a database query.

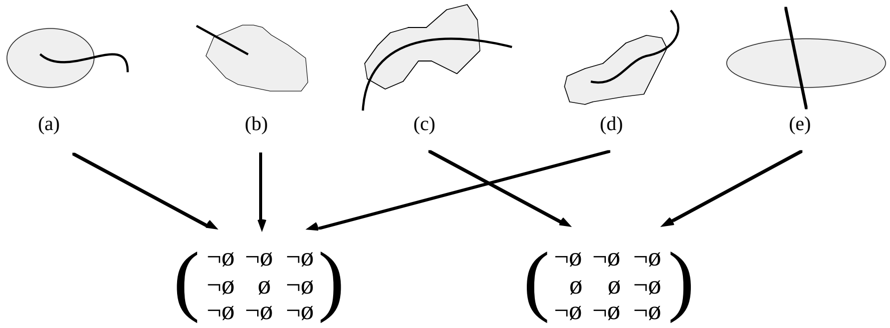

**Figure 6**: The 9-intersection as a categorization of the spatial relations sketched: (a), (b), and (d) map onto the same 9-intersection, while (c) and (d) map onto a different 9-intersection.

### 4.2 Component Invariants for Detailed Topological Relations

More detailed distinctions about topological relations are possible if further criteria are employed to evaluate the non-empty intersections. In order to establish topological relation equivalence between two regions (i.e., to decide whether or not two pairs of objects have the same topological relations), it is sufficient to describe such invariants for the components (or separations) of the boundary-boundary intersection, since the other intersections can be inferred from them (Egenhofer and Franzosa 1995). The necessary invariants to consider are:

- The *sequence* of components counted along the boundaries (Figure 7a).
- The *dimension* of each component (Figure 7b).
- The *type* of boundary-boundary component intersection—*touching* if the boundary enters and leaves the intersection from the same part, or *crossing* if the boundary enters from a different part than it leaves (Figure 7c).
- The *complement relationship*, i.e., whether a component is a next to a bounded or unbounded exterior (Figure 7d).

Detailed topological relations between two regions are expressed by the *component invariant table* for non-empty boundary-boundary sequences, which lists the sequence of boundary-boundary components and each component's dimension, type, and complement relationship (Egenhofer *et al.* 1994; Egenhofer and Franzosa 1995).

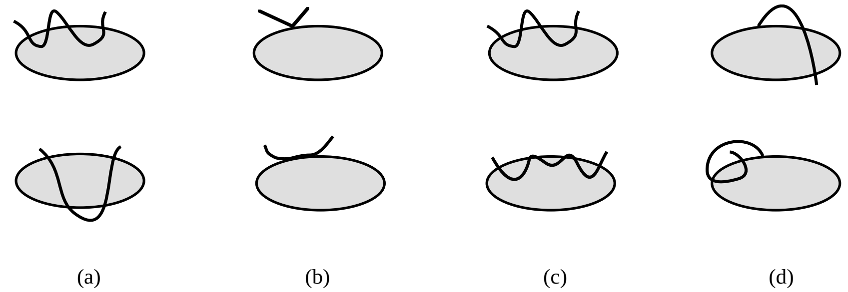

**Figure 7**: Pairs of topological relations that distinguish by component invariants: (a) different boundary sequences, (b) different component dimensions, (c) different component types, and (d) different complement relationships.

We use the component invariants as the key for analyzing the intentional complexity of the spatial relations sketched. The component invariants capture complexity of spatial relations. If a user draws a sketch with a high level of complexity, then we assume that this complexity was intended and that it provides the lower bound of what should be retrieved; therefore, a configuration in a spatial database with the same 9-intersection relation, but lower-rated component invariants, would not qualify as a match. On the other hand, a sketch of a low-complexity spatial relation may indicate that more complex configurations under the same 9-intersection category should be considered as well.

### 4.3 Conceptual Neighborhoods for Similarity of Topological Relations

Similarity among topological relations is described in terms of the *conceptual neighborhood graph*, which links those relations that are most similar to each other (Figure 8). It is based on the computational model of determining for each relation those relations with the least number of differences in the 9-intersection matrices (Egenhofer and Al-Taha 1992; Egenhofer and Mark 1995a).

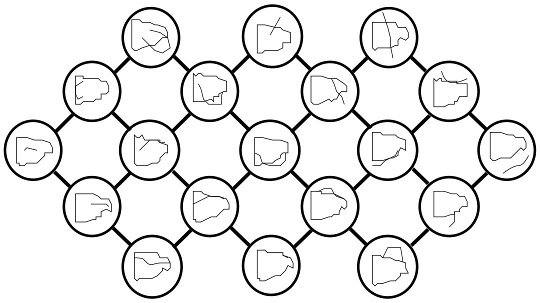

**Figure 8**: Conceptual neighborhood graph of the 19 line-region relations.

We use the conceptual neighborhoods of 9-intersection relations as the key to the semantics of natural-language spatial relations. In our previous work we found that topology is the primary part of the definition of natural-language spatial relations and that the 9-intersection is an appropriate model for a wide range of natural-language spatial relations (Mark and Egenhofer 1994a), applicable to different natural languages (Mark and Egenhofer 1995). For this investigation we have available a corpus of 1500 drawings for sixty different English-language spatial relations. An analyses of selected sketches found topological agreements in the way subjects referred to spatial relations. Figure 9 shows four examples of natural-language spatial relations and their mappings onto the conceptual neighborhoods of corresponding line-region relations (i.e., the set of highlighted relations).

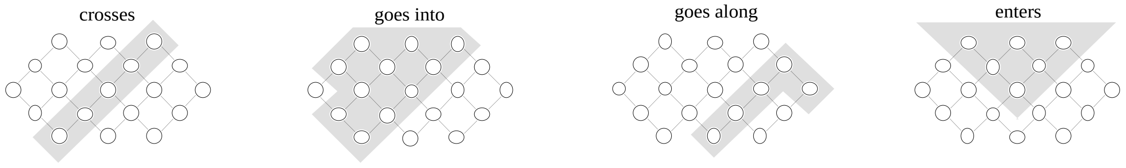

**Figure 9**: Conceptual neighborhoods of the topological relations for the natural-language terms (a) crosses, (b) goes into, (c) goes along, and (d) enters.

### 4.4 Quantitative Refinements for Metric Details

Occasionally, topology *per se* is insufficient to characterize the essence of spatial relations. For instance, in order to capture the semantics of the spatial relation between Interstate I-95 and New Hampshire requires the consideration of some metric properties in addition to topological concern—I-95 divides New Hampshire into a very small area to the East of I-95 and a larger piece to the West. We apply such measures about areas, lengths, and directions as refinements of the topological properties. For line-region relations we use the following set of metric properties:

- the ratio by which a line’s interior divides a region’s interior (Figure 10a);
- how much of a line’s interior is inside a region’s interior (Figure 10b);
- how much of a line’s interior coincides with a region’s boundary (Figure 10c);
- the ratio between the distance from a line to and region’s boundary, and the line’s total length (Figure 10d);
- the ratio between the area made up by an equidistant enlargement (reduction) of the region (also known as a *buffer zone*) and the actual area (Figure 10e); and
- the cardinal direction between the objects, measured in a qualitative orientation scheme (Frank 1992) such as orthogonal half planes with a neutral zone (Figure 10f).

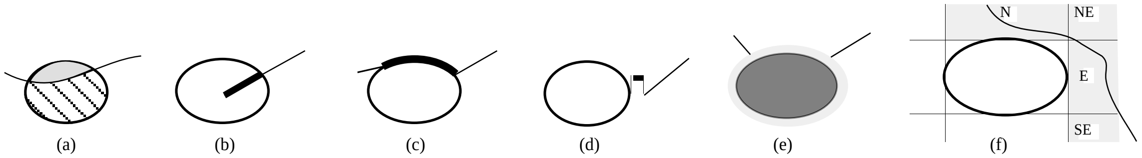

**Figure 10**: Quantitative refinements of topological relations due to (a) the area ratio, (b) the line-interior ratio, (c) the line-boundary ratio, (d) the line’s closeness, (e) the region’s closeness, and (f) the orientation.

We use the quantitative refinements of the 9-intersection relations as the key to formalizing detailed geometric constraints about natural-language spatial relations and sketched spatial relations. Some natural-language spatial relations may depend heavily on metric properties. For example, *Northeast* describes a relation where the road is outside of the park, but it is equally important that one object has a particular orientation with respect to the other (Figure 10e). More complex configurations include such natural-language descriptions as *along*, which qualifies for configurations of varying topology in combination with certain metric constraints about the line-boundary ratio (Figure 10 c), the line’s closeness (Figure 10d), and the region’s closeness (Figure 10e). Qualitative refinements are also critical to interpret sketched configurations that deviate considerably from the prototypical configuration. A particularly short line-closeness measure (Figure 10d), for instance, may indicate that the user would accept as an answer a configuration with a slightly different topology. The qualitative metric properties would indicate which of the configurations targeted through the conceptual neighborhoods would qualify.

M. Egenhofer, Multi-Modal Spatial Querying, in: J.-M. Kraak and M. Molenaar (eds.), Seventh International Symposium on Spatial Data Handling, Delft, The Netherlands, Taylor & Francis, London, pp. 785-799.

## 5. Processing Stages of a Sketch-and-Talk Query

### 5.1 Recording

The user draws a query and simultaneously talks about what she is drawing. Sketch-and-Talk records graphics, voice, and their temporal concurrencies by time-stamping them at regular intervals (Figure 11). We analyze the users' multi-modal interactions and determine which spatial concepts they prefer to draw and which ones they rather describe verbally.

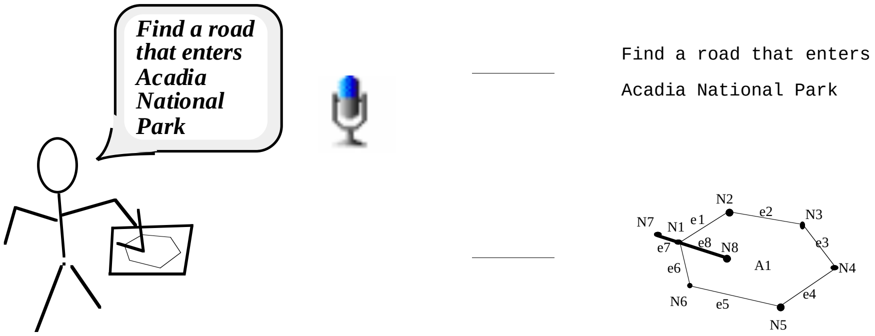

**Figure 11:** Recording the spoken and sketched queries and translating them into symbolic representations of unstructured text and a topological vector model, respectively.

### 5.2 Parsing

Sketch-and-Talk parses the graphical and verbal sentences, and translates them into a compatible format (Figure 12). We build a semantic network for the spoken spatial query and translate its natural-language spatial relations into their corresponding conceptual neighborhoods of 9-intersection relations, building on a library of mappings from English-language terms onto 9-intersections and quantitative refinements. For the sketched query, we extract the 9-intersections, components invariants, and quantitative refinements from the topological data model.

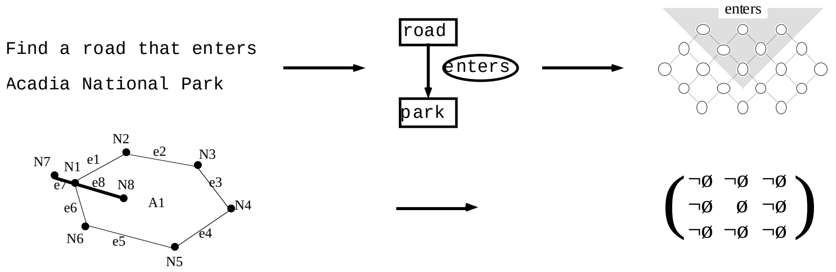

**Figure 12:** Translating text and graphical query into a compatible representation.

### 5.3 Integration

Sketch-and-Talk integrates the two representations by cross-referencing in the two sentences the corresponding relations and adding, as available from the verbal description, the object types and their values to the objects sketched. Where necessary, Sketch-and-Talk consolidates the integrated spatial relations resolving incompleteness, ambiguities, and contradictions (Figure 13). We exploit the unifying format of the 9-intersection (and where necessary component invariants and quantitative refinements) to link the query parts expressed in different modes. If the same relations are expressed both verbally and graphically, we use this information to consolidate the query. This includes the completion of otherwise incomplete query parts, and the resolution of ambiguities in one representation by using information derived from the other interaction mode. Finally, we exploit consensus by putting higher preferences on those aspects of a spatial query that were well specified both verbally and graphically.

M. Egenhofer, Multi-Modal Spatial Querying, in: J.-M. Kraak and M. Molenaar (eds.), Seventh International Symposium on Spatial Data Handling, Delft, The Netherlands, Taylor & Francis, London, pp. 785-799.

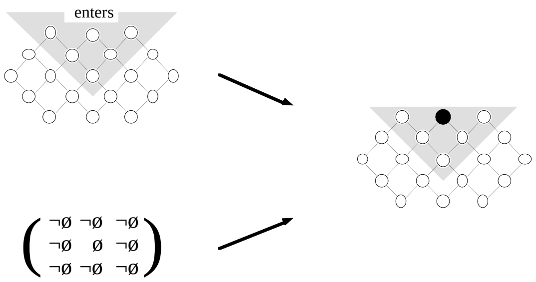

**Figure 13**: Integrating the spatial relation described verbally and sketched.

### 5.4 Query Processing

Sketch-and-Talk develops a query processing plan, translates it into the database's query language (e.g., SQL-3 MM), and retrieves the scenes that match the query. Retrieval and presentation of the results should be in a prioritized order such that the scenes most similar to the sketch are presented first to the user. We use the integrated representation of spoken and sketched query and iteratively relax the spatial constraints to generate a set of database queries of decreasing specificity. We start with executing the most specific queries first and build for each configuration found a similarity measure to assess how closely it matches the query asked. If the set of answers is small, further relaxations may be necessary. In a post-processing step, the answers are sorted according to the similarity measures so that the best matches are presented first to the user.

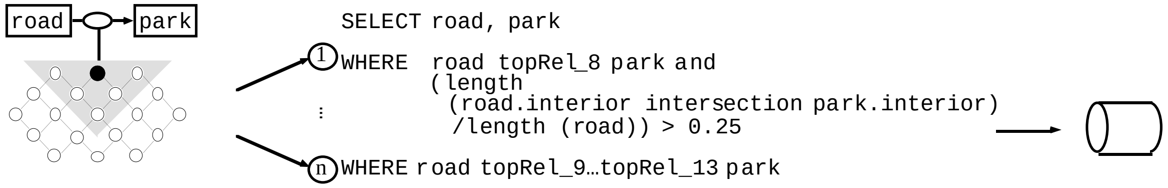

```
SELECT road, park
WHERE  road topRel_8 park and
      ( length
        (road.interior intersection park.interior)
        /length (road)) > 0.25
WHERE road topRel_9…topRel_13 park
```

**Figure 14:** Developing a prioritized query processing plan from the integrated representation and processing these queries against a spatial database.

## 6. Conclusions

The combination of graphical and voice input enables a new style of spatial querying. The two interaction modes are complementary, and when used together they provide more power than either of them alone. The conceptual design of Sketch-and-Talk will be tested through a prototype implementation, and we will evaluate Sketch-and-Talk with user tests.

## 7. References

M.-A. Aufaure-Portier (1995) Definition of a Visual Language for GIS. in: T. Nyerges, D. Mark, R. Laurini, and M. Egenhofer (Ed.), *Cognitive Aspects of Human-Computer Interaction for Geographic Information Systems*. pp. 163-178, Kluwer, Dordrecht.

A. Borning (1986) Defining Constraints Graphically. in: *Human Factors in Computing Systems, CHI `86*, Boston, MA, pp. 137-143.

D. Calcinelli and M. Mainguenaud (1994) Cigales, a Visual Query Language for a Geographical Information System: the User Interface. *Journal of Visual Languages and Computing* 5: 113-132.

C. Chang and S. Lee (1991) Retrieval of Similar Pictures on Pictorial Databases. *Pattern Recognition* 24(7): 675-680.

H. Clark (1973) Space, Time, Semantics, and the Child. in: T. Moore (Ed.), *Cognitive Development and the Acquisition of Language*. pp. 27-63, Academic Press, New York, NY.

E. Clementini, A. D'Atri, and P. D. Felice (1990) Browsing in Geographic Databases: An Object-Oriented Approach. in: *Workshop on Visual Languages*, Skokie, IL, pp. 125-131.

M. Deering (1995) HoloSketch: A Virtual Reality Sketching/Animation Tool. *ACM Transactions on Computer-Human Interaction* 2(3): 220-238.

A. Del Bimbo, P. Pala, and S. Santini (1994) Visual Image Retrieval by Elastic Deformation of Object Sketches. in: *IEEE Simposium on Visual Languages*, St. Louis, MO, pp. 216-223.

M. Egenhofer (1990) Interaction with Geographic Information Systems via Spatial Queries. *Journal of Visual Languages and Computing* 1(4): 389-413.

M. Egenhofer (1992) Why not SQL! *International Journal of Geographical Information Systems* 6(2): 71-85.

M. Egenhofer (1994) Spatial SQL: A Query and Presentation Language. *IEEE Transactions on Knowledge and Data Engineering* 6(1): 86-95.

M. Egenhofer and K. Al-Taha (1992) Reasoning About Gradual Changes of Topological Relationships. in: A. Frank, I. Campari, and U. Formentini (Ed.), *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*. *Lecture Notes in Computer Science* 639, pp. 196-219, Springer-Verlag, Pisa.

M. Egenhofer, E. Clementini, and P. di Felice (1994) Evaluating Inconsistencies Among Multiple Representations. in: T. Waugh and R. Healey (Ed.), *Sixth International Symposium on Spatial Data Handling*, Edinburgh, Scotland, pp. 901-920.

M. Egenhofer and R. Franzosa (1991) Point-Set Topological Spatial Relations. *International Journal of Geographical Information Systems* 5(2): 161-174.

M. Egenhofer and R. Franzosa (1995) On the Equivalence of Topological Relations. *International Journal of Geographical Information Systems* 9(2): 133-152.

M. Egenhofer and J. Herring (1990) A Mathematical Framework for the Definition of Topological Relationships. in: K. Brassel and H. Kishimoto (Ed.), *Fourth International Symposium on Spatial Data Handling*, Zurich, Switzerland, pp. 803-813.

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 (Ed.), *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.

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 (1982) MAPQUERY—Database Query Language for Retrieval of Geometric Data and its Graphical Representation. *ACM Computer Graphics* 16(3): 199-207.

A. Frank (1992) Qualitative Spatial Reasoning about Distances and Directions in Geographic Space. *Journal of Visual Languages and Computing* 3(4): 343-371.

M. Grimaud (1988) Toponyms, Prepositions, and Cognitive Maps in English and French. *Journal of the American Society of Geolinguistics* 14: 54-76.

A. Hauptmann (1989) Speech and Gestures for Graphic Image Manipulation. in: K. Bice and C. Lewis (Ed.), *CHI '89*, Austin, TX, pp. 241-245.

J. Herring, R. Larsen, and J. Shivakumar (1988) Extensions to the SQL Language to Support Spatial Analysis in a Topological Data Base. in: *GIS/LIS '88*, San Antonio, TX, pp. 741-750.

A. Herskovits (1986) *Language and Spatial Cognition—An Interdisciplinary Study of the Prepositions in English*. Cambridge University Press, Cambridge, MA.

K. Hirata and T. Kato (1992) Query by Visual Example—Content-Based Image Retrieval. in: A. Pirotte, C. Delobel, and G. Gottlob (Ed.), *Advances in Database Technology—EDBT '92, 3rd International Conference on Extending Database Technology*. *Lecture Notes in Computer Science* 580, pp. 56-71, Springer-Verlag, Vienna, Austria.

T. Imielinski and H. Korth (1995) MOBIDATA Workshop Report. *MOBIDATA: An Interactive Journal of Mobile Computing* 1(2): http://www.cs.rutgers.edu/~badri/journal/contents12.html.

K. Ingram and W. Phillips (1987) Geographic Information Processing Using a SQL-Based Query Language. in: N. R. Chrisman (Ed.), *AUTO-CARTO 8, Eighth International Symposium on Computer-Assisted Cartography*, Baltimore, MD, pp. 326-335.

R. Jain and A. Pentland (1995) *Workshop Report: NSF-ARPA Workshop on Visual Information Management Systems*. Technical Report available from http://www.virage.com/vir-res/reports/.

N. Japkowicz and J. Wiebe (1991) A System for Translating Locative Prepositions from English into French. in: *29th Annual Meeting of the Association for Computational Linguistics*, UC Berkeley, pp. 153-160.

K. Koperski and J. Han (1995) Discovery of Spatial Association Rules in Geographic Information Databases. in: M. Egenhofer and J. Herring (Ed.), *Advances in Spatial Databases—4th International Symposium, SSD '95, Portland, ME*. *Lecture Notes in Computer Science* 951, pp. 47-66, Springer-Verlag, Berlin.

W. Kuhn (1990) From Constructing Towards Editing Geometry. in: *ACSM-ASPRS Annual Convention*, Denver, CO, pp. 153-164.

W. Kuhn (1992) Paradigms of GIS Use. in: D. Cowen (Ed.), *Fifth International Symposium on Spatial Data Handling*, Charleston, SC, pp. 91-103.

G. Lakoff and M. Johnson (1980) *Metaphors We Live By*. University of Chicago Press, Chicago, IL.

D. Mark and M. Egenhofer (1994a) Calibrating the Meanings of Spatial Predicates from Natural Language: Line-Region Relations. in: T. Waugh and R. Healey (Ed.), *Sixth International Symposium on Spatial Data Handling*, Edinburgh, Scotland, pp. 538-553.

D. Mark and M. Egenhofer (1994b) Modeling Spatial Relations Between Lines and Regions: Combining Formal Mathematical Models and Human Subjects Testing. *Cartography and Geographic Information Systems* 21(3): 195-212.

D. Mark and M. Egenhofer (1995) Topology of Prototypical Spatial Relations Between Lines and Regions in English and Spanish. in: D. Peuquet (Ed.), *Autocarto 12*, Charlotte, NC, pp. 245-254.

V. Ogle and M. Stonebraker (1995) Chabot: Retrieval from a Relational Database of Images. *IEEE Computer* 28(9): 40-48.

D. Papadias and T. Sellis (1994) Qualitative Representation of Spatial Knowledge in Two-Dimensional Space. *VLDB Journal* 3(4): 479-516.

A. Pizano, A. Klinger, and A. Cardenas (1989) Specification of Spatial Integrity Constraints in Pictorial Databases. *Computer* 22(12): 59-71.

E. Rosch (1978) Principles of Categorization. in: E. Rosch and B. Lloyd (Ed.), *Cognition and Categorization*. pp. Erlbaum, Hillsdale, NJ.

L. Talmy (1983) How Language Structures Space. in: H. Pick and L. Acredolo (Ed.), *Spatial Orientation: Theory, Research, and Application*. pp. 225-282, Plenum Press, New York, NY.

R. M. White (1988) Applying Direct Manipulation to Geometric Construction Systems. in: N. Magnenat and D. Thalmann (Ed.), *New Trends in Computer Graphics: Computer Graphics International '88*. pp. 446-455, Geneva, Switzerland.