D. Fritsch, M. Englich & M. Sester, eds, 'IAPRS', Vol. 32/4, ISPRS Commission IV Symposium on GIS - Between Visions and Applications, Stuttgart, Germany.

SKETCH-BASED IMAGE RETRIEVAL IN AN INTEGRATED GIS ENVIRONMENT

Anthony Stefanidis and Peggy Agouris
Dept. of Spatial Information Science & Engineering
and the National Center for Geographic Information and Analysis
University of Maine
Orono, ME 04469-5711
tel: 207-5812180 fax: 207-5812206
{tony, peggy}@spatial.maine.edu
Commission III, Working Group IV/III.2

KEY WORDS: Image Retrieval, Spatial Information Management Systems, Shape Queries, Integrated Environments.

ABSTRACT

The intelligent retrieval of images from large databases is receiving increased attention in the computer vision community. In this paper we address this problem considering the particularities of image databases encountered in typical topographic applications. Specifically, we present the development of a spatial data management system which enables the use of sketch-based queries for the content-based retrieval of digital images from topographic databases. We discuss our overall strategy and present some relevant algorithmic and implementation aspects, as well as associated database design issues. The query input are user-provided sketches of the shape and spatial configuration of the object (or objects) which should appear in the images to be retrieved. The objective of the database search is to retrieve images which contain a configuration of objects sufficiently similar to the one specified in the query. The results are ranked according to statistical scores and the user can subsequently narrow or broaden his/her search according to the previously obtained results and the purpose of the search. Our approach combines the design of an integrated database with the development of the necessary matching tools. Our matching tool specifically is inspired by least-squares matching, and represents an extension of lsm to function with a variety of raster representations. The potential to eventually handle digital maps and general GIS datasets, in addition to images, makes this a solid step towards the integration of GIS and image databases.

1 INTRODUCTION

During the last few years, the intelligent retrieval of images from large databases is emerging as a major research direction within the computer vision community [Jain, 1992, Ogle & Stonebraker, 1995, Pickard & Minka, 1995; Gudivada & Raghavan, 1995; Guibas & Tomasi, 1996; Cohen & Guibas, 1996; Gupta et al., 1997]. The objective of intelligent retrieval is to find specific images in a large database using as query properties of these images. As a result of pioneering research efforts, some prototype systems have been reported, with notable examples being Virage [Gupta, 1995], Chabot [Ogle & Stonebraker, 1995], IBM's QBIC [Flickner et al., 1995], VisualSeek [Smith & Chang, 1995], ImageRover [Sclaroff et al., 1997], and PicHunter [Cox et al., 1997].

The common trend in these efforts is that they address the problem within the context of multimedia applications, and therefore they focus in general-use, multimedia-type image databases. A database representative of such applications includes for example images of sunsets, politicians, cartoon characters, snow-covered mountains, and wild animals. An objective of a query might be to retrieve the images of sunsets from the database. In such a scenario, low-level image properties (e.g. color and dominant pattern) are adequate for information retrieval, since the image members of the database display substantial differences in these properties and can be distinguished by them alone. For example, the images of sunsets would be characterized as images whose histogram shows high percentages of dark and red pixels, with the red pixels grouped together in a circle in the image above a dark region.

Topographic image databases are unique in the sense that they contain very large numbers of images (typically aerial and/or satellite) which represent striking similarity in terms of general low-level image properties. Therefore, general-purpose image retrieval approaches like the ones mentioned above are not sufficient for information retrieval in topographic image databases. Instead, what distinguishes images in a topographic database is the *shape* and *configuration* of the objects they contain. Up to this point, information retrieval in topographic image datasets is mostly supported by metadata, describing approximate image location, date of capture, scale etc. The support of metadata for image retrieval is compromised by two emerging trends:

- the increasing availability of multitemporal image representations of specific regions, essential for topographic database updating; in such images general metadata properties remain the same (besides, of course, the date parameter) while certain important objects within them evolve, and
- the need to make complex scene understanding decisions, based on the behavior (existence, modification, absence, relative positions) of objects within scenes, an increasingly important issue within modern integrated GIS environments.

Advancements in content-based image retrieval from topographic databases are a key issue in the move towards integrated geographic information systems. They allow the integration of images and GIS, enabling complex spatial analysis tasks.

D. Fritsch, M. Englich & M. Sester, eds, 'IAPRS', Vol. 32/4, ISPRS Commission IV Symposium on GIS - Between Visions and Applications, Stuttgart, Germany.

In this paper we present our approach for the retrieval of images from topographic image databases using query-by-sketch operations. Our approach is *progressive*, as we employ increasingly specific information to retrieve imagery. It is also *content-based*, with the term content referring to objects depicted in the images. We present the strategy and design considerations behind I.Q. (Image Query), our prototype system for image retrieval (section 2), and emphasize on system and database design (section 3), and some digital image analysis issues related to matching sketches to images for querying (section 4). It should be mentioned that while our research originates from topographic applications, the developed methodology can be applied to any type of imagery.

2 STRATEGY

A description of the envisioned operation environment for our system is shown in Fig. 1. A searchable topographic database comprises

- images (typically aerial or satellite),
- outline/object information for physical entities depicted in these images, and
- metadata.

Metadata of interest include the typical elements (e.g. sensor characteristics, date of capture, resolution) which describe and enhance the content and/or properties of common topographic data files (e.g. digital images, DEMs, maps in digital format).

The aim of our strategy is to take advantage of the intuitive method humans use to express spatial scenes, namely *sketching*. In accordance, the query interface of I.Q. intends to allow us to access individual members of this database by using as input parameters *metadata* information and a *sketch*. Our approach is designed to proceed as follows:

- an operator sketches a configuration of objects,
- he/she also provides additional metadata information, and
- the database is searched to yield the images which satisfy the given metadata information, and in which spatial configurations similar to the given sketch appear.

In this sense, a *query match* is an image which satisfies the given query parameters. Database searching is performed progressively. Metadata information is used to thin the pool of potential matches (*query by metadata*). Additionally, an analysis of shapes is performed to identify the best matches of the query, and to assign estimates of confidence to them (*query by shape*). The reason for this two-stage design is rather evident: metadata searches are computationally inexpensive, fast, and therefore optimal for thinning large search spaces. On the other hand, shape-based searches are in general computationally demanding, but allow us to move from global image properties - which are conveyed by metadata - to individual features (*content*) within images. Metadata information is further supporting shape-based queries, as it indirectly assigns some semantic connotations to the sought features. For example, a specific shape has different meaning when found in an aerial image of scale 1:5,000, than when found in a satellite image with 5-meter ground resolution.

The need for sketch-based queries within topographic databases is related to a simple fact: general color and texture properties are typically inadequate to differentiate aerial or satellite images. The variations caused by different depicted areas on histogram and other relevant global properties are only somewhat noticeable, especially when comparing images of rural to urban and coastal areas. But they are not sufficiently perceptible to be exploited in query processes within the context of topographic applications. Whereas the term “query-by-content” tends to be used for a wide variety of applications, the nature of our problem forces us to use it quite literally: perform searches in image databases according to objects depicted in them.

3 SYSTEM DESIGN

Outline and object extraction from digital images is a computationally expensive operation, and it is therefore performed not during the queries, but rather off-line when new images are introduced in the database. This results in a system architecture shown in Fig. 2, where dashed lines indicate off-line operations and solid lines indicate on-line operations. I.Q. is reached by Java servers through the Web, and accesses a library of features and metadata to retrieve the proper images.

The database design is shown in Fig. 3, and comprises the following elements:

<u>Image library :</u> contains one entry for every image of the database, and provides a link/pointer to the corresponding filename.

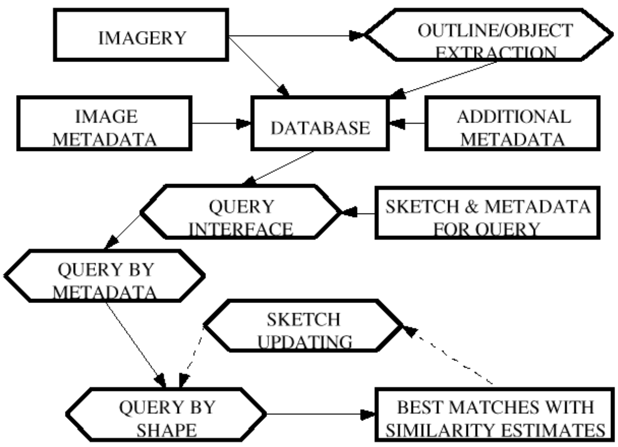

Fig. 1: Operation environment for I.Q.

D. Fritsch, M. Englich & M. Sester, eds, 'IAPRS', Vol. 32/4, ISPRS Commission IV Symposium on GIS - Between Visions and Applications, Stuttgart, Germany.

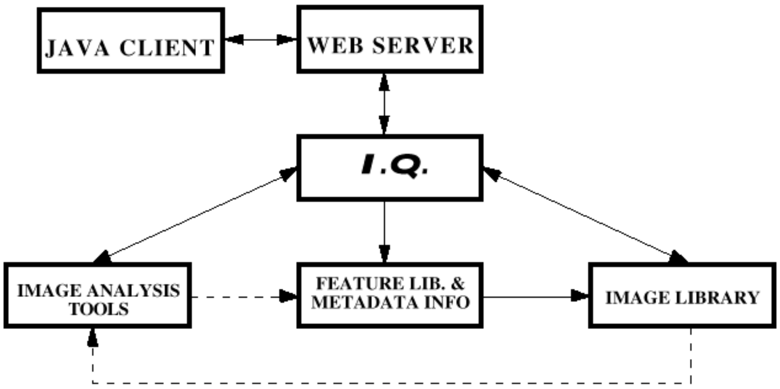

Fig. 2: System architecture

Metadata library : contains a listing of potential values for a set of attributes which describe general properties of the images. These attributes include date and time of acquisition, date and time of introduction in the database, scale/resolution, and location of the image (expressed in hierarchically arranged geographic entities like state, county, city). For more complex databases the attributes may be extended to incorporate information on the sensor and imagery type (e.g. b/w, color, pseudocolor).

Feature library : contains a set of distinct features (i.e. object shapes) and links to image files where such features appear. The role of the feature library is to provide the crucial link which allows us to reduce the search space of a query from a database to an abridged group of features. In order for the query to be efficient, this library needs to be optimal. The optimality criteria are two: the members of the library should be *exhaustive* (thus being able to adequately describe all possible input features), and the members of the library should be *independent* (avoiding unnecessary duplications). The two properties, when satisfied, are equivalent to an ideal library which is approaching a base spanning the space of shapes.

In addition to the above, the database may include a semantic library, which will contain semantic object information (e.g. object X is a hospital) and the corresponding links to image files.

Under this design, the *on-line* part of the query is performed in this manner:

- the user provides as input a set of metadata values and a sketch; the sketch depicts the object (or object configurations) which we wish to retrieve, and the metadata values describe the acceptable characteristics of the images in which the object(s) may appear,
- using the input values we identify an acceptable subset of the metadata library; and this provides links to features within the feature library and thus defines a feature subset,
- the input sketch is matched to the feature subset (instead of the complete feature library) using our on-line matching tool,
- acceptable matches give links to specific images (and locations within them) where objects similar to our input sketch do appear, and
- this information is returned to the user who then has the option to edit his/her query. If a semantic library is used, the above mentioned results will pass through another check (whereby the semantic properties of the detected object will be examined), and the query results would be provided after this added step.

In order to support the above sequence, the following sequence has to be performed *off-line* every time an image is introduced in the database:

- the user inputs the appropriate metadata information for this image,
- the metadata library is updated to add the new entry,
- objects/features are extracted from the input image using digital image analysis tools,
- these new features are compared to the existing complete feature library using our off-line matching tool, and library entries are updated accordingly to include links to the objects existing within the newly introduced image, and
- the links between metadata and feature libraries are updated to connect the metadata values of the new image to the features detected in it.

To extend the queries in configurations of objects, we employ the well-known concept of 9-intersection, describing the major topological relations between areal, linear, and point [Egenhofer & Franzosa, 1995]. Thus, we break the query for a spatial scene into two tasks. First, we identify where individual objects similar to the input exist in the database. This results in a list of links to image locations. Then, we identify these combinations of these locations which satisfy the given topological relationship.

For example, when attempting to retrieve raster files which depict an hexagonal building and a cross-like structure separated by an L shaped object, we identify the locations in the database where each of the three objects exists. Subsequently we analyze this information (which now comprises filenames and locations within these files) to come up with the combinations which satisfy the desired configuration. This formulation of scene queries (as combination of a space query and a topological check) permits easy and fast query edits. When attempting to edit a query by reconfiguring a given set of objects, we do so by re-examining their topological relations only, an operation simpler than searching the shape library.

D. Fritsch, M. Englich & M. Sester, eds, 'IAPRS', Vol. 32/4, ISPRS Commission IV Symposium on GIS - Between Visions and Applications, Stuttgart, Germany.

The rationale behind our database design becomes apparent when analyzing the meaning of the metadata and feature libraries, and their connection. Assuming that we have n distinct metadata values, the metadata space is an n-dimensional one. A point within this space corresponds to all images of the same area, captured at the same scale, at the same date, with similar sensor. When one or more of these parameters can accept less specific values we move to blobs within the n-dimensional metadata space. For example, photos of various scales of

matching algorithms are a modification of least squares matching (lsm), modified to function with both images and edge files. Least squares matching offers a robust method for establishing correspondences among image windows. Its mathematical background, based on least-squares principles, permits its successful extension for application in a multiple image matching scheme [Agouris & Schenk, 1996], or even for the establishment of correspondences in sets of 3-dimensional images [Maas et al., 1994].

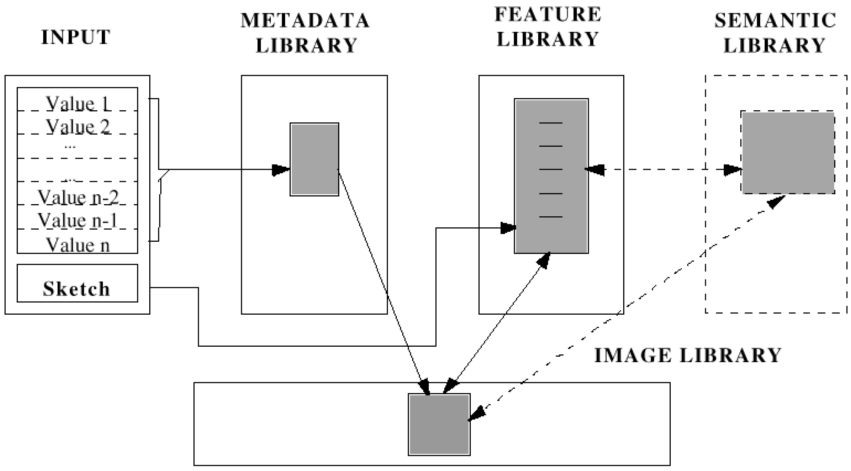

Fig. 3: Database design

a specific area taken on a specific date form a blob in the metadata space. This blob represents the scale space of the area at the time of data capture. When defining points (or blobs) of the metadata space we actually declare what kind of features appear at a specific area of the geographic space. When querying the database, we use the metadata information to narrow the area of interest, and then we perform the shape query against the shapes that we know (from the off-line process) to exist in our region of interest. In essence, our design reverses the traditional processes by which we identify objects in the geographic space. Instead of identifying an object and subsequently positioning it, we now identify a region and access a list of objects within this region.

It should be mentioned here that for implementation purposes, one can initially use as library any readily available set of shapes, for example a library like IPE, the Interactive Picture Environment [Schwartzkopf, 1995], which provides nearly two hundred illustrations, augmented to incorporate certain standard cartographic figures. The shape library can be augmented every time a new shape query does not find an adequate match in the library. In this case the query pattern becomes a member of the library, and is compared off-line against the database to acquire its link information (images and positions within them where a pattern similar to it appears).

4 MATCHING

The matching tools developed for the environment presented in the previous section need to handle both the off-line (population and update of the feature library) and on-line (comparison of query input to the feature library entries) matching parts of this project. Conceptually, our

Off-line Matching

The off-line matching process begins when a new image is entered into the image database. First the user inputs the relevant metadata (Fig. 4) associated with the new image and then I.Q. extracts the edges/objects from the image by applying a generic edge enhancing/thresholding filter to produce a final edge file consisting of only black (gray level = 0) and white (gray level = 255) pixels. This edge file is linked to the image with an identical filename but with different extension. Second, I.Q. begins to match all the features from the feature library to the edge file and records in the feature library which features found a match and the matching location within the image. This task is equivalent to matching structured object entries to the approximations of objects provided by common edge detection techniques.

Our matching tool is based on the analysis of dissimilarities between a template and an edge file. When a template is compared to an edge file, the template is divided into four quadrants, centered on the central pixel of the template feature. Each quadrant is matched separately to the image, and pixels vote to stay put (or move) according to their similarity (respectively dissimilarity). This resembles the comparison of gray values in least squares matching and the use of image gradients to identify shifts, rotations, and scalings. The sum of the votes for each pixel within the quadrant is recorded, together with move recommendations for the direction and magnitude of motion. Directions are checked in the +/- x and y directions. This means each quadrant votes to move in one of 5 options: left, right, up, down or stay put.

Each quadrant vote is then compared with the other quadrants and a weighted average of direction and

D. Fritsch, M. Englich & M. Sester, eds, 'IAPRS', Vol. 32/4, ISPRS Commission IV Symposium on GIS - Between Visions and Applications, Stuttgart, Germany.

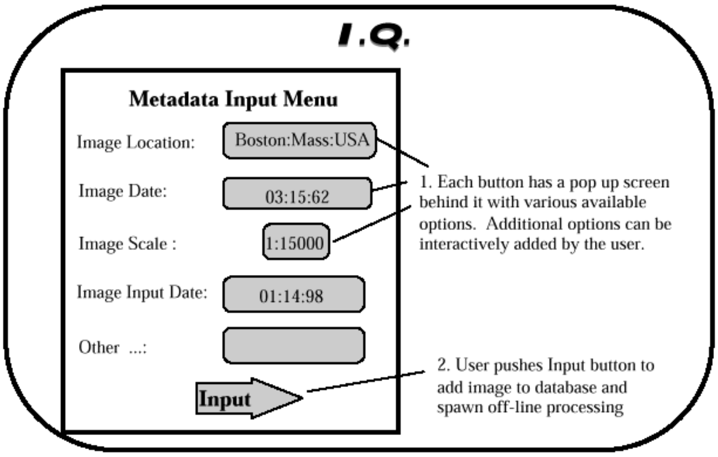

Fig. 4: Metadata Input Menu

distance allows us to reposition the template. A sample of *some* of the possible quadrant votes and template feature shifts can be found in Table 1. (Note: a vote of ✥ means this quadrant's vote could be in any of the four cardinal directions; currently I.Q. handles 40 cases of different combinations of quadrant votes).

The analysis of the quadrant shift directions and distances proceeds first through translation, then rotation then scale. The amount of translation is taken as the average distance calculated from the 4 quadrants. For rotation, the transformation of coordinates for the feature pixels are calculated using:

$$x' = x \cos \alpha + y \sin \alpha \qquad y' = y \cos \alpha - x \sin \alpha$$

where is the angle of rotation along the x axis. The rotation angle is taken in 15 degree increments, positive in the counter clockwise direction. If upon the next iteration the quadrant shifts determine a rotation angle in the opposite direction, then it is halved to 7.5 degrees. This halving of the rotation angle continues as long as the direction for rotation keeps alternating between positive and negative α or until the arbitrary limit of iterations (20 is a typical value) is met.

For scaling the feature, there are two approaches: independent axis and global. Independent axis scaling is only used when the user digitizes his own feature to match, therefore in the on-line matching process. This will allow for the feature to scale by different amounts in the x and y directions. For global scaling, it is used in the off-line matching process and may be used in the on-line matching process if the user chooses an existing feature from an image to search the feature library against. Global scaling assumes a constant scale change in all directions, which is usually the case when images are scanned at different resolutions or taken with different focal lengths or flying heights.

Once the library feature has settled onto a match its accuracy is determined by how many of its pixels continue to vote to move compared to those that vote to stay where they are. Each of the subregions in an image will have its best match position and match percentage recorded and sorted for each image in the database. The best match percentage for each of the images is used to determine which of the images get linked to a particular feature within the feature template library.

By breaking a feature template into four quadrants which can be manipulated independently from each other, we can handle extreme cases like occlusions. This is achieved by permitting two or three of the quadrants to adjust themselves, ignoring the part of the template (two or one quadrant respectively) which corresponds to the occluded region.

On-line Matching

The on-line matching process begins by sketching a feature to be matched against the feature library (which is in turn linked to the image database). This can be done in two ways. The user can select a feature off of one of the images in the database using the I.Q. interface and edit it, or, digitize from scratch his/her own feature outline. The I.Q. interface allows the user to turn on and off feature outline pixels inside the sketch-edit area (Fig. 5). When the user clicks on a white (255) pixel it will turn black (0) and vice versa.

If the user selects a feature from an existing image as the template feature, this object will have a nominal scale associated with it based on its scan resolution for satellite imagery and scan resolution combined with focal length and flying height for aerial images. The user can also provide a specific scale to the template feature if so desired. In this manner, matches will be prioritized such that the less the template feature has to scale in order to fit the images, the closer the feature is to its true match.

Once the template feature has settled onto a match, its accuracy is determined by how many of its pixels continue to vote to move compared to those that vote to stay where they are. Each library feature template will have its best match position and match percentage recorded and sorted. The best matching percentage is used to determine its ranking for sorting the returned images based on the original query (Fig. 6).

When the results of a query are returned (Fig. 6), the user can re-edit the template feature by clicking on its pixels or by changing some options concerning the metadata. The new query can then be run on a subset of the database, or on the entire database if the the metadata have also been modified. The subset of the database is a listing of images created during the first search that match to the query. For example, 50 images may be returned by the first query, with the top 5 being displayed on the query

D. Fritsch, M. Englich & M. Sester, eds, 'IAPRS', Vol. 32/4, ISPRS Commission IV Symposium on GIS - Between Visions and Applications, Stuttgart, Germany.

q1	q2	q3	q4	Shift
→	→	→	✥	→
←	←	✥	←	←
↑	✥	↑	↑	↑
✥	↓	↓	↓	↓
→	→	↑	↓	→
←	↑	↑	→	↑

q1	q2	q3	q4	Shift
→	←	←	→	scale in x larger
↑	↑	↓	↓	scale in y larger
←	→	→	←	scale in x smaller
↓	↓	↑	↑	scale in y smaller
←	↓	→	↑	rotate ↶
↓	→	↑	←	rotate ↷

Table 1: An Abridged Selection of Quadrant Votes and Feature Shifts

results screen. A cutoff matching percentage can be used to determine which images are included in this query subset of images. If the second or subsequent queries begin to return the sort of images the user is not interested in finding, he/she can use the undo button in (Fig. 6) to return to previous query parameters. Both the sketch and metadata query parameters are stored for each individual query and are available for recall.

Currently, I.Q. does not distinguish between how much a feature may have been scaled or rotated when determining its matching percentage. A test on how much a feature was manipulated in order to fit the image could and perhaps should be made and reflected in its ranking. A matching resulting from substantial manipulations should be ranked lower than one resulting from few. Also, another addition could be to further subdivide each of the quadrants into quadrants themselves. This will facilitate the template feature in warping itself to fit all the objects in the image. This further subdividing of quadrants into quadrants could conceivably go on to many levels of feature manipulation.

The above description covers only the basic concepts behind our matching approach. A more detailed description of the strategy behind our off-line matching may be found in [Agouris & Stefanidis, 1998].

5 EXPERIMENTAL RESULTS & COMMENTS

A prototype of the system described in this paper has been developed and is presently being implementationally enhanced and optimized. After extracting individual features from some of the images in the database, I.Q. was tested with some preliminary but meaningful results. Table 2 shows the magnitude of the matching percentages of extracted features f1, f2, ...,f5 with a sample of the images returned. To test the reliability of the matching technique, the extracted features were then modified (pixels added and/or re-arranged) and the query re-run on the same 30 image database. For example, modified features f3, f4 and f5 (renamed ff3, ff4, ff5) returned the same matched images with the same or lower matching percentage in 80% of the cases.

Considering the state-of-the-art in computer vision, the advantages of our method are mainly associated with its precise yet versatile theoretical foundation, and the overall system design. The matching core module used in I . Q . is able to properly function for comparing sketches/outlines to digital image windows or to other sketches. This offers the potential of extending this system to function within multimedia spatial databases, allowing queries to be performed simultaneously on images, maps, and GIS vector databases. This very important aspect anticipates the move towards integrated spatial information systems, and is the focus of our future work. It is expected to enable us to perform complex scene analysis tasks, by combining spatial and semantic content available over a variety of digital imagery and other, complementary, spatial datasets.

Due to both system design and the matching tool potential, our method is also scale independent. This is very important for query operations which are inherently multiresolutional, i.e. you have a sketch which depicts a feature or a combination of features at a certain resolution. By employing a multiresolutional strategy, these differences are directly accounted for, and the accuracy and overall performance potential of our technique is optimized.

Image→ Feature↓	I1	I2	I3	I4	I5
f1	100	23.4	28.1	18.8	12.5
f2	27.7	100	16.9	29.2	18.5
f3	6.4	18.6	100	8.7	19.3
f4	7.1	10.1	8.9	100	13.6
f5	7.6	12.4	18.6	15.9	100
ff3	6.2	14.1	54.2	10.3	14.6
ff4	6.1	7.9	11.7	48.6	13.6
ff5	7.3	14.1	17.2	15.1	37

Table 2: Match results (%)

Furthermore, the method is robust in terms of mathematical and statistical foundation, an advantage which can be best exploited for the objective interpretation of the query operation. Indeed, matching candidates can be arranged and ranked based on a statistical analysis of their similarity to the sought-after template. The potential for such objective interpretation is inherent in the method's set-up and is based on the least squares foundation of the employed matching module, something which is clearly extremely important for our project.

Last, but not least, this method presents a step forward in spatial database queries: we move from queries based on metadata or other global properties of images (e.g. dominant color) to queries using features within images. In doing so, we put more emphasis on the semantic information content of images, which is conceptually a more reliable, and accurate way for image retrieval. Furthermore, the use of user-provided sketches, as well as the ability of a human operator to manipulate the query process (e.g. by revising sketches, or changing metadata) brings the human operator into the process itself, taking advantage of his/her superior cognitive abilities.

D. Fritsch, M. Englich & M. Sester, eds, 'IAPRS', Vol. 32/4, ISPRS Commission IV Symposium on GIS - Between Visions and Applications, Stuttgart, Germany.

Acknowledgments

This work was partially supported by the National Science Foundation through grant number SBR-8810917 for the National Center for Geographic Information and Analysis, and through CAREER grant number IRI-9702233. We also acknowledge the work of Mr. James Carswell in the development of the I.Q. prototype.

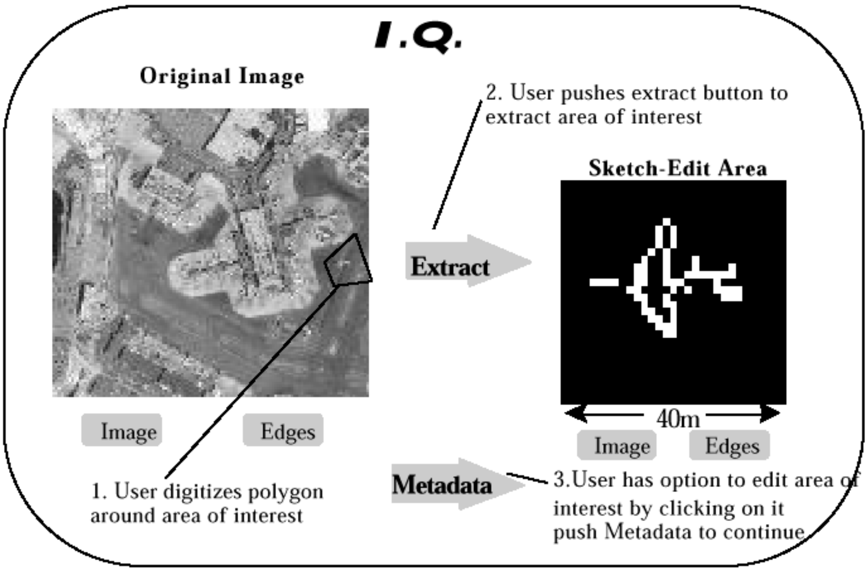

Fig. 5: Extracting Area of Interest

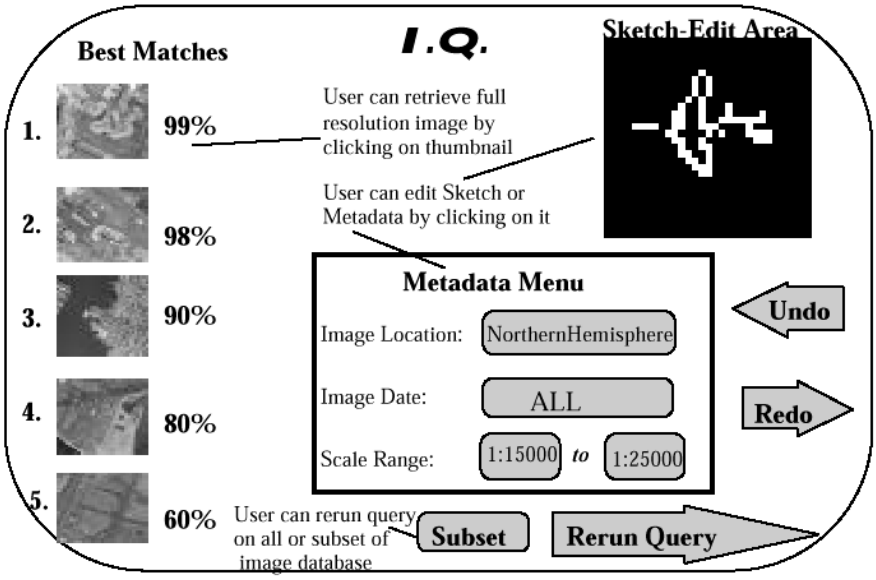

Fig. 6: Query Results

REFERENCES

Agouris P. & T. Schenk. Automated Aero-triangulation Using Multiple Image Multipoint Matching. *Photogrammetric Engineering & Remote Sensing*, Vol. 62, No. 6, pp. 703-710, 1996.

Agouris P. & A. Stefanidis. Intelligent Image Retrieval from Large Databases using Shape and Topology. *IEEE-ICIP98*, Chicago. Il, 1998. (in press)

Cohen S.D. & L.J. Guibas. Shape-Based Indexing and Retrieval; Some First Steps. in *Proceedings 1996 ARPA Image Understanding Workshop*, Vol. 2, pp. 1209-1212, 1996.

D. Fritsch, M. Englich & M. Sester, eds, 'IAPRS', Vol. 32/4, ISPRS Commission IV Symposium on GIS - Between Visions and Applications, Stuttgart, Germany.

Cox I.J., J. Ghosn, M.L. Miller, T. Papathomas & P. Yianilos. Hidden Annotation in Content Based Image Retrieval. *Proceedings IEEE Content-Based Access of Image and Video Libraries*, pp. 76-81, 1997.

Egenhofer M. & R. Franzosa. On the Equivalence of Topological Relations. *International Journal of Geographical Information Systems*, Vol. 9, No. 2, pp. 133-152, 1995.

Flickner M., H. Sawney, W. Niblack, J. Ashley, Q. Huang, B. Dom, M. Gorkaani, J. Hafner, D. Lee, D. Petkovic, D. Steele, & P. Yanker. Query by Image and Video Content: The QBIC System. *Computer*, September 1995, pp. 23-32, 1995.

Gruen A. & P. Agouris. Linear Feature Extraction by Least Squares Template Matching Constrained by Internal Shape Forces. in *International Archives of Photogrammetry & Remote Sensing*, Vol. 30, Part 3/1, pp. 316-323, 1994.

Gruen A, P. Agouris & H. Li. Linear Feature Extraction with Dynamic Programming and Globally Enforced Least Squares Matching. in *Automatic Extraction of Man-Made Objects from Aerial and Space Images*, (A. Gruen, O. Kuebler & P. Agouris eds.), pp. 83-94, Birkhaeuser Verlag, 1995.

Gudivada V. & V. Raghavan. Design and Evaluation of Algorithms for Image Retrieval by Spatial Similarity. *ACM Transactions on Information Systems*, Vol. 13, No. 2, pp. 115-144, 1995.

Gupta A. Visual Information Retrieval: A Virage Perspective. TR95-01, Virage, Inc., San Mateo, CA. 1995.

Gupta A., S. Santini, & R. Jain. In Search of Information in Visual Media, *Communications of the ACM*, Vol. 40, No. 12, pp. 34-42.

Guibas L.J. & C. Tomasi. Image Retrieval and Robot Vision Research at Stanford. in *Proceedings 1996 ARPA Image Understanding Workshop*, Vol. 1, pp. 101-108, 1996.

Jain R. *NSF Workshop on Visual Information Management Systems*. Redwood, CA, 1992.

Lindeberg T. *Scale-Space Theory in Computer Vision*. Kluwer Academic Publishers, Boston, 1994.

Maas H.-G., A. Stefanidis & A. Gruen. Feature Tracking in 3-D Fluid Tomography Sequences. in *ICIP-94*, Vol. I, pp. 530-534, 1994.

Mehrotra R. & J.E. Gary. Similar-Shape retrieval in Shape Data Management. *Computer*, pp. 57-62, September 1995.

Ogle V.E. & M. Stonebraker. Chabot: Retrieval from a Relational Database of Images. *Computer*, pp. 23-32, September 1995.

Pickard R.W. & T.P. Minka. Vision Texture for Annotation. *Multimedia Systems*, Vol3, No. 1, pp. 3-14, 1995.

Sclaroff S., L. Taycher & M. LaCascia. ImageRover: A Content-Based Image Crowser for thr World Wide Web. *Proceedings IEEE Content-Based Access of Image and Video Libraries*, pp. 2-9, 1997.

Schwartzkopf O. *Manual of Ipe*, Technical Report, Utrecht University. (ftp://ftp.cs.ruu.nl/pub/X11/Ipe/), Holland, 1995

Smith J.R. & S.-F. Chang. Searching for Images and Video on the World Wide Web. *Multimedia Systems*, Vol3, No. 1, pp. 3-14, 1995.