

A Flexible Image Database System for Content-Based Retrieval*

Andrew P. Berman and Linda G. Shapiro†

Department of Computer Science and Engineering, University of Washington, Seattle Washington 98195-2350

E-mail: aberman@cs.washington.edu, shapiro@cs.washington.edu

There is a growing need for the ability to query image databases based on similarity of image content rather than strict keyword search. As distance computations can be expensive, there is a need for indexing systems and algorithms that can eliminate candidate images without performing distance calculations. As user needs may change from session to session, there is also a need for runtime creation of distance measures. In this paper, we present FIDS, “flexible image database system.” FIDS allows the user to query the database based on complex combinations of dozens of predefined distance measures. Using an indexing scheme and algorithms based on the triangle inequality, FIDS can often return matches to the query image without directly comparing the query image to more than a small percentage of the database. This paper describes the technical contributions of the FIDS approach to content-based image retrieval. © 1999 Academic Press

1. INTRODUCTION

There is a growing interest in image databases that can be queried based on image content rather than just with keywords. Such queries are based on the use of distance measure-scoring functions that rate the similarity of two images based on predefined criteria. There are several challenges that arise from this approach. The first challenge is that of flexibility. The user’s definition of similarity may change from session to session, creating a need for runtime creation of distance measures. Another challenge is that of speed. Distance measure computation requires accessing either the images being compared or precomputed associated data. Furthermore, distance computation can be somewhat expensive. Thus, a system that must compute the distance from the query image to each image in a large database may exhibit prohibitively unsatisfactory performance. For certain distance measures and data sets, indexing or clustering schemes can be used to reduce the number of direct comparisons. Yet standard clustering or indexing schemes may not be efficient for some combinations of data sets and distance measures.

We have designed and implemented a prototype database system that allows the user a great deal of flexibility in run-

time distance measure creation. The system can also reduce the number of direct distance measure calculations between a given query object and the database elements. FIDS, or “flexible image database system,” has been tested on a database of 37,000 images and is being upgraded to handle 100,000 images. FIDS allows the user to find approximate matches to query images using complex combinations of dozens of predefined distance measures. FIDS can often return results without directly comparing the query image to more than a small percentage of the original database.

There are algorithms in the literature based on the triangle inequality that can reduce the number of distance measure calculations [4, 6, 3, 10, 2] in object retrieval. These methods have the advantage of being applicable to any distance measure that satisfies the triangle inequality. FIDS uses algorithms based on the triangle inequality that are extensions to the methods in the literature.

In certain application domains, object *features* can be pre-computed and stored. A system can take advantage of this pre-computation to speed up object searches at runtime. However, the cost of feature comparison by itself can still be prohibitive. In this paper, we do not distinguish between “image comparisons” and “feature comparisons.” Although FIDS does precompute features of images in its database, it does not currently store those features and does not use them directly. Thus, when we state that a direct comparison is eliminated, the implication is that the feature comparison is also eliminated.

This paper describes the technical contributions of the FIDS approach to content-based image retrieval. Section 2 discusses related work in image retrieval. Section 3 discusses the use of multiple distance measures in user queries and the operations that can be used to combine them. Section 4 describes the FIDS indexing scheme, which employs the triangle inequality to rule out large portions of the image database from direct comparison with the query, and Section 5 presents the FIDS interface. Section 6 gives experimental results of the basic retrieval algorithm, while Section 7 discusses and presents results for several different methods for selecting keys to be used in indexing. Section 8 describes an additional data structure, the triangle-trie, and presents experimental results on using it in conjunction with the basic FIDS retrieval methods to speed up the search even more.

* This research was partially supported by the National Science Foundation under Grant IRI-9711771.

† Corresponding author.

2. RELATED LITERATURE

The related work falls into two categories: (1) systems that attempt to provide a general image retrieval capability and (2) systems that provide a single retrieval technique, often for a particular application. In general image retrieval systems, one of the first efforts was the work of Chang [13] in which retrieval of images was achieved through attribute matching, spatial relation matching, structural (contour) matching, and similarity matching using various application-specific similarity measures. Kato [23, 27] developed an experimental database system called ART-MUSEUM that was intended to be an electronic art gallery. The system included a visual interface by which a user could enter a hand-drawn sketch, a monochrome photo or xerox copy, or a full color image of a painting in order to retrieve matching images from the database. The QBIC (query by image content) system developed at IBM Almaden [33] became the first commercial product. The original QBIC allowed retrieval of images by color, texture, and the shape of image objects or regions. The system is constantly updated to add new retrieval methods. Another commercial product, VIR, developed by Virage, Inc. [20], allows retrieval based on color, composition, texture, and structure measures. Virage has collaborated with Compaq to produce the *AV Photo Finder*, a website that allows retrieval of images based on color, composition, structure, texture, and keywords. They catalog over 10 million images.

Pentland's group at MIT developed the Photobook system [34], a set of interactive tools for browsing and searching images and image sequences that allows queries by appearance, shape, and texture. Appearance refers to the technique of matching with eigenimages developed by Turk and Pentland [43], shape is based on the work of Pentland and Sclaroff [37], and texture matching comes from the work of Picard and Liu [36].

Most single-purpose systems have also concentrated on one of the four main classes of matching: shape, color, texture, and composition matching. In shape matching, most of the 2D contour matching techniques developed for computer vision apply, although several researchers have particularly targeted their work toward image databases. Grosky and Mehrotra [19, 31] used index trees to access a database of 2D object contour models of industrial parts. Califano and Mohan [11] developed a related indexing method that uses multidimensional global invariants of tuples of local interest features as indexes that vote for object models in the database. Del Bimbo [8] retrieved images containing specified 2D shapes using an elastic matching technique.

The spatial and/or frequency distributions of colors or gray tones in an image are often used for retrieval. All of the general purpose systems contain one or more distance measures for color or "appearance-based" retrieval. Jacobs, Finkelstein, and Salesin [24] used a multiresolution wavelet decomposition to represent images for rapid matching and retrieval. Color indexing has been thoroughly analyzed by Swain and Ballard [41].

Texture classification and segmentation algorithms have been heavily studied in the computer vision/pattern recognition com-

munity [42]. In the image database community, Picard and Minka [35] developed a texture-based vision annotation system that learns from user feedback. Kelly and Cannon [28, 29] at Los Alamos National Laboratory used a global signature to characterize images and a signature distance function to compare them. Jain, Murthy, and Chen [26] have developed a methodology for comparing various texture similarity measures.

Most of the above techniques are global in that they compare some global feature or features of an image, such as signatures, histograms, eigenimages, and wavelet representations. Structural techniques break the image into a set of extracted entities and produce a description in terms of these entities, their properties, and their interrelationships [40]. Descriptions are encoded as string, tree, or graph structures, and relational matching is performed to determine the distance between a query and a stored description. Chang, Shi, and Yan [14] encoded the spatial relations among picture objects as 2D strings and provided an iconic indexing technique using these strings to retrieve images. In 3D object recognition work performed at the University of Washington, Costa and Shapiro developed an accumulator-based relational indexing technique that uses subgraphs of a structural description of the image to vote for 3D view-class object models [39].

Del Bimbo [7, 9] used both spatial and temporal relations to select image sequences that match a query. Bach, Paul, and Jain developed a feature-based approach to retrieval of face images [1]. The distance between a query image and an image from the database was a weighted sum of the differences between pairs of corresponding query image features and database image features.

Much of the older work in content-based retrieval only gave illustrations of system performance without any real evaluation. In newer work, Jain and Vailaya [25] describe a system for content-based retrieval of trademark images using a combination of color and shape features. A thorough set of experiments was run to judge the accuracy, stability, and speed of the approach. Accuracy was defined to mean that the most similar image to the query image should be in the top set of returned images. Stability or robustness meant that the procedures should not break down under various conditions; the conditions tested were arbitrary rotations, size variations, and addition of random noise. Another interesting and thorough study was the recent work of Minka and Picard [32]. The Four Eyes system, which extends their texture work, represents a general supervised learning approach to image retrieval.

More recently, Sclaroff *et al.* [38] used a k-d tree algorithm to index images collected on the Web. Using a relevance feedback mechanism, the user attempts to progressively home in on the desired images. They also combined their results with keyword searches. Netra [30] allows searches based on simple combinations of color, texture, shape, and position of shapes extracted from images offline.

Because of the shortcomings of computer vision segmentation algorithms, very little of the older work has involved actual




FIG. 1. Result of query with a color-based distance measure.

object identification. Forsyth *et al.* [18] recently proposed a new approach that involves a sequence of grouping activities that satisfy increasingly strict constraints. As an example, Fleck, Forsyth, and Bregler [16] developed a procedure for finding naked people in color images, for possible use in products that restrict access to web sites with pornographic material. The process has two levels of filtering. The first filter looks for large areas (30% of the image) of skin color. Regions that pass this test go on to a geometric analyzer looking for elongated regions that can be grouped into certain spatial relationships that are typical of limbs connected to torsos. This work is a first step in the recognition of abstract concepts in difficult, real images. Another promising approach is that of Carson *et al.* [12] who segment images into *blob* regions with annotated color and texture properties, enabling retrieval of images that contain similar regions.

3. THE USE OF MULTIPLE DISTANCE MEASURES

A *distance measure* is a function that computes and returns a value corresponding to the similarity between two objects according to some predefined criteria. For example, Fig. 1 illustrates the result of a query that used a color-histogram distance measure to retrieve images similar to the leftmost image. Figure 2 illustrates the result of a query that used a texture-based distance measure to retrieve images for the same query image as in Fig. 1.

Two fundamental classes of distance measures are those measures based on color properties of images and those measures

based on image texture properties. There are many different specific color measures and texture measures in use and an infinite number of distinct possible measures. Other types of measures may be based on features such as the existence of particular objects within the image, the relative position of various image features or literally any quality that one can ascribe to an image.

3.1. The Need for Multiple Measures

One can create an innumerable number of distance measures for images based on any set of features and an arbitrary scoring mechanism. We cannot program all these distance measures in advance. This difficulty motivates the idea of giving the user a predefined set of base distance measures that he or she can combine to create more complex measures. The potential space of creatable distance measures increases as we increase the number of predefined measures. For example, Haering, Myles, and da Vitoria Lobo [22] demonstrated a neural network with 43 distance measures as input that was trained to detect deciduous trees.

Apart from having many different distance measures, we believe that there is a need for multiple variants of the basic distance measures. A piece of software representing a distance measure also represents many decisions made by the programmer. The user of the distance measure may disagree with, not understand, or not even know about these decisions. For example, there are innumerable ways to define bin size and placement for the underlying color histogram of a color-histogram-based distance




FIG. 2. Result of query with a texture-based distance measure.

measure. A user who does not wish to distinguish pink from red will be unhappy with a color measure that does distinguish the two colors. Giving the user multiple color histograms from which to choose increases the potential utility of the system for such a user.

The problems of color measurement only increase when dealing with texture distance measures. As in color-histogram measures, there may be many decisions made by the programmer that control the behavior of the texture measure. These decisions are opaque to the user, yet they affect the performance of the measure for the user. The difficulties are compounded in that the semantics of texture are more obtuse than the semantics of color. Thus, just giving the user multiple textures from which to choose is not a complete strategy, since he or she will have no method (other than trial and error) by which to decide what measures to use.

A more subtle problem with distance measures is that they force the user to make concrete decisions about similarity when the user may not desire to do so. The user may wish to find images that are “similar” without actually specifying any particular measure. For example, consider the user who wants the set of all closest matches to a query image over all possible color-histogram-based distance measures. This set may actually be quite small, yet reliance on any single distance measure will leave out some matches. In this case, such a fuzzy similarity can be approximated by querying the database with several different color distance measures.

3.2. Methods of Combining Distance Measures

Given a database system with several distance measures, one has to decide in what ways the user shall be allowed to combine the distance measures. There are three separate, possibly conflicting objectives. The first objective is to provide many choices to the user. The second objective is to provide an interface that allows the user to understand his choices. The third objective is to support database search of the resultant queries in an efficient manner.

Systems such as QBIC [17] and Virage [21] offer the user the ability to take weighted combinations of color, texture, shape, and position measures, combined with keyword search. But what

of the user who wishes to formulate queries such as “match on colors, unless the texture and shape are both very close” or “two out of three of color, texture, and shape must match”? These queries cannot be expressed as a weighted sum of individual distance measures. In order to expand users’ searching vocabulary, more complex combinations of distance measures must be offered.

To deal with these problems to some extent, we proposed [5] the following set of operations to enable more expressive queries ($d_1 \dots d_n$ represent distance measures):

- Addition: $d = d_1 + d_2$
- Weighting: $d = cd_1$
- Max: $d = \text{Max}(d_1, d_2, \dots, d_n)$
- Min: $d = \text{Min}(d_1, d_2, \dots, d_n)$.

These operations are all invariant under inequality. That is,

$$x_1 \leq y_1, x_2 \leq y_2 \Rightarrow x_1 + x_2 \leq y_1 + y_2$$

$$x \leq y, c \geq 0 \Rightarrow cx \leq cy$$

$$x_1 \leq y_1, x_2 \leq y_2 \Rightarrow \text{Min}(x_1, x_2) \leq \text{Min}(y_1, y_2)$$

$$x_1 \leq y_1, x_2 \leq y_2 \Rightarrow \text{Max}(x_1, x_2) \leq \text{Max}(y_1, y_2).$$

We later show how to take advantage of this invariance to apply triangle-inequality-based pruning algorithms to distance measures that are combined together using the above operations.

Weighting and addition are already commonly used in both commercial and research systems [21, 17]. The reason for Max and Min is that they enable queries that we expect to be useful. Max, for example, enables tight searches: “I want a match on color and texture and position and shape.” Min enables more speculative searches similar to those required by the data mining community: “I want a match on color or texture or position or shape.” Figures 3, 4, and 5 demonstrate the utility of combining distance measures. The query image is a particularly difficult one for the distance measures in the system. In Fig. 3, a color distance measure returns unsatisfactory results. A texture measure proves no better as shown in Fig. 4. Yet a distance measure based on the sum of the color and texture measures returns a viable match to the query in Fig. 5.

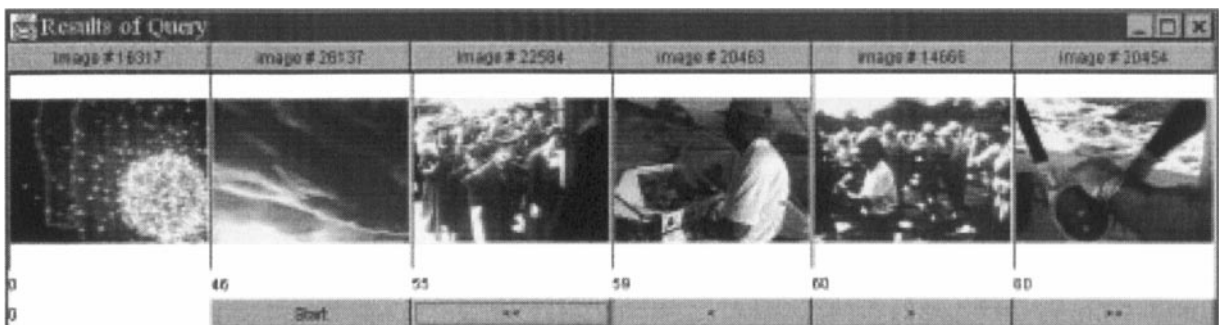


FIG. 3. Poor result of a query with a color-based distance measure.




FIG. 4. Poor result of a query with a texture-based distance measure.

Fagin [15] has also proposed extending multimedia queries to boolean combinations, using Min and Max to implement them. We experimented with taking powers of distance measures [6] but decided that the functionality was too unintuitive with little apparent gain in utility.

3.3. Query Strategies: Threshold and Best-Match

Given an image database \mathcal{S} , a query image Q , and a distance measure d , there are two main types of searches. One can request all images $I \in \mathcal{S}$ such that the distance $d(Q, I)$ is not greater than some given threshold value t . We label this task to be a *threshold* style query. The second method is to find the image or images in \mathcal{S} which minimize $d(Q, I)$. We label this task to be a *best-match* style query. Note that we can naturally extend best-match queries to find the n best matches.

3.4. How to Query the Database

User understanding of the distance measures is a problem with any content-based retrieval system. Our proposed set of operations adds a great deal of complexity to the system. There is a need for an additional layer to bridge the gap between user understanding and system capabilities. Possibilities include example-based learning and natural language translation. It may be that different image database domains will require different interfaces.

Standard comparison techniques usually involve comparing features which have been computed over entire images. How-

ever, one can consider the utility of comparing corresponding sections of images. For example, the user may wish to find images whose centers match the center area of a query image. Or the user may care about color in the top half of an image, but care about texture in the lower half. We can give the user some control over the locations which the user cares about by using *gridded* distance measures. The idea is to break up each image into a grid of rectangles of the same proportions as the original image. Given image I , let $I_{r,c,x,y}$ be the (x, y) rectangle of I which has been broken up into a (r, c) grid of rectangles. Given a distance measure d , let $d_{r,c,x,y}(I, Q)$ be defined as $d(I_{r,c,x,y}, Q_{r,c,x,y})$. That is, $d_{r,c,x,y}$ “pretends” that the chosen rectangles are the entire images to be compared.

Using gridded distance measures gives the user the capability to create queries with some positional semantics. For example, a user could express the following queries with combinations of gridded distance measures:

- Match the lower left corner by color.
- Match the top half by color and texture and the bottom by texture.
- Match the center by color and all of the image by texture.

Assuming a 3×3 grid, color measure C , and texture measure T , this last query would be expressed in our formulation by $\text{Max}(C_{3,3,1,1}, T)$. The distance measure $C_{3,3,1,1}$ compares the $(1, 1)$ or central rectangle of the query Q to the central rectangle of the database image I as illustrated in Fig. 7. The texture




FIG. 5. Improved result of query with a combination distance measure.




FIG. 6. Looking for water and trees, but not necessarily fall tree colors.

measure T compares all of Q to all of I . The Max operation implements the equivalent of a logical AND operator.

Our system employs such gridded measures. We have found it useful in retrieving matches to images that have different qualities in different areas of the image. In Fig. 6, we are looking for matches to a scene with water in the foreground and trees in the background. The picture was taken in fall as can be discerned by the autumn coloring of the leaves. However, we do not wish to restrict the returned set to autumn images. Therefore, we choose to match on texture in all four quadrants of the image, but only match on color in the bottom half of the image. Not including the query image itself, two of the returned images have water in the lower half and trees in the upper half. However, the trees in the returned images do not share the fall coloring of the query image. Given that T was our texture measure and C was our color measure, the formula for the composite distance measure used in this search was $2(\text{Max}(T_{2,2,0,0}, T_{2,2,0,1}, T_{2,2,1,0}, T_{2,2,1,1})) + \text{Max}(C_{2,2,0,1}, C_{2,2,1,1})$. This formula is ample illustration of the need for good user interfaces. It is not something a user would wish to type.

One could imagine extending gridded distance measures to arbitrary shapes. One can further imagine extending this to nonconstant shapes such as those calculated by segmentation. Here are some other examples using these extensions:




FIG. 7. Gridded distance measures are defined on a particular grid rectangle of the query and the corresponding grid rectangle of the image being compared. The highlighted grid rectangle is the (1, 1) rectangle of the (3, 3) grid. Distance measures with subscript (3, 3, 1, 1) reference this grid rectangle.

- Match the largest segments by color and the next largest segments by texture.
- Match the center circular area by color and the largest segments by color and texture.

4. INDEXING WITH THE TRIANGLE INEQUALITY

There are several schemes in the literature [2–4, 6, 10] that take advantage of the triangle inequality to reduce the number of direct comparisons in a threshold style database search. The intuition behind all the schemes is that the distance between two objects cannot be less than the difference in their distances to any other object.

The indexing scheme and algorithm described here, defined informally as the *bare-bones triangle inequality algorithm*, outputs a value for each database image corresponding to a lower bound on the distance between that image and the query image. This set of values can be used in several different ways. One can use the lower bounds to discard images that are shown to be too far from the query image to be a potential match. Alternatively, one can sequence the images in increasing order of their calculated lower bounds. Experimental evidence suggests that the first images in such a sequence are the ones most likely to be the best matches to the query. In our examples below, we assume that a threshold is applied to the lower bounds on the images, rather than an ordering.

The bare-bones triangle inequality algorithm is probabilistic in nature. The lower bounds are guaranteed to be correct, but the ordering of the images based on their lower bounds is not guaranteed to be the same as the ordering of the images based on their true distances to the query image. To improve accuracy, one can add an additional step in which the database images with the smallest lower bounds are compared directly to the query image to obtain their true distances. Our experiments have shown that good results can be obtained with direct comparison of only a tiny fraction of the original database to the query.

4.1. Indexing with a Single Distance Measure

Let I represent a database object, Q represent a query object, K represent an arbitrary fixed object known as a *key*, and d represent some distance measure that is a metric. As d is a

pseudo-metric, the two triangle inequalities, $d(I, Q) + d(Q, K) \geq d(I, K)$ and $d(I, Q) + d(I, K) \geq d(Q, K)$, must be true. We can combine them to form the inequality which places a lower bound on $d(I, Q)$,

$$d(I, Q) \geq |d(I, K) - d(Q, K)|. \quad (1)$$

Thus, by comparing the database and query objects to a third key object, a lower bound on the distance between the two objects can be obtained. We define $l(d, K, I, Q) = |d(I, K) - d(Q, K)|$ to be equal to this lower bound on $d(I, Q)$. We further shorten $l(d, K, I, Q)$ to $l(d, K)$ when there is no confusion as to the identity of I and Q .

Burke and Keller [10] first proposed the idea of using the triangle inequality to reduce comparisons. This idea was used by Uhlmann [44] to create vantage-point trees, where each node in a tree corresponds to a carefully chosen key. The subtree rooted at that node was partitioned according to the distance of the leaf elements to the key. Berman [4] and Baeza-Yates *et al.* [2] separately refined Burke and Keller's algorithm by creating a single set of keys to use for all the objects in the database. The single key method differs from vantage-point trees in that it reduces the total number of key comparisons at the expense of increasing the number of fast operations. Barros *et al.* [3] successfully used a single set of keys and the triangle inequality in a real image database. What follows is a description of the general algorithm used with a single set of keys:

Equation (1) can be extended naturally by substituting a set of keys $\mathcal{K} = (K_1, \dots, K_M)$ for K as

$$d(I, Q) \geq \max_{1 \leq s \leq M} |d(I, K_s) - d(Q, K_s)|. \quad (2)$$

We can see that this inequality is valid by noting that $d(I, Q) \geq |d(I, K_s) - d(Q, K_s)|$ for all values of s . We define $l'(d, \mathcal{K}, I, Q)$ to be equal to the lower bound on $d(I, Q)$ found by using Eq. (2). As before, we shorten $l'(d, \mathcal{K}, I, Q)$ to $l'(d, \mathcal{K})$ where possible.

Consider a large set of database objects, $\mathcal{S} = \{I_1, \dots, I_n\}$, and a much smaller set of key objects, $\mathcal{K} = \{K_1, \dots, K_m\}$. Precalculate $d(I_s, K_t)$ for all $\{1 \leq s \leq n\} \times \{1 \leq t \leq m\}$. Now consider a request to find all database objects I_s such that $d(I_s, Q) \leq t$ for some query image Q and threshold value t . We can calculate lower bounds on $\{d(I_1, Q), \dots, d(I_n, Q)\}$ by calculating $\{d(Q, K_1), \dots, d(Q, K_m)\}$ and repeatedly using Eq. (2). If we prove that t is less than $d(I_s, Q)$, then we eliminate I_s from our list of possible matches to Q . After the elimination phase, we may search linearly through the uneliminated objects, comparing each to Q in the standard fashion. This algorithm involves $m + u$ distance measure calculations, and $O(mn)$ simple (constant cost) operations where u is the number of uneliminated objects. The hope is that $m + u$ is sufficiently smaller than n to result in an overall time savings. The pseudocode for setting up the index structure, calculating a lower bound for a query and an image, and retrieving images for a query with a single distance measure is given in Fig. 8.

```

procedure Index_Single(ImageSet[], KeySet[], d())
{
for X := 0 to SizeOf(ImageSet)-1
  for Y := 0 to SizeOf(KeySet)-1
    DistanceMatrix[X][Y] := d(ImageSet[X], KeySet[Y]);
}

procedure Calc_Lower_Bound(Query, i, d(), DistanceMatrix)
‘Query’ is the Query image
‘i’ is the database Image’s index
‘d()’ is the distance function
‘DistanceMatrix’ is the precalculated
  image-key distance matrix for distance function ‘d()’
{
for X := 0 to SizeOf(KeySet)-1
  QKDist[X] = d(Query, KeySet[X]);
  Distance := 0;
for Y := 0 to SizeOf(KeySet)-1
  if (Abs(QKDist[Y] - DistanceMatrix[i][Y]) > Distance)
    then Distance = Abs(QKDist[Y] - DistanceMatrix[i][Y]);
}

procedure Retrieve_Single_With_Threshold(Query, t, d())
{
  R = null;
for Y := 0 to SizeOf(ImageSet)-1
  if Calc_Lower_Bound(Query, Y, d(), DistanceMatrix) ≤ t
    then R = R ∪ ImageSet[Y];
return R;
}

```

FIG. 8. Pseudo-code for the indexing and retrieval procedures using the triangle-inequality algorithm with a single distance measure.

EXAMPLE OF INDEXING WITH A SINGLE DISTANCE MEASURE. Let our sample database be an image database composed of the images $\mathcal{S} = (I_1, \dots, I_6)$. Our keys are images $\mathcal{K} = (K_1, K_2)$. To initialize the database for distance measure d , we calculate $d(I_s, K_j)$ for all s, j as shown in Table 1. Now suppose we wish to find all images I_s in our database such that $d(I_s, Q) \leq 2$ for some query object Q . We calculate $d(K_1, Q) = 3$ and $d(K_2, Q) = 5$. We subtract 3 from each element in the first column in Table 1 and subtract 5 from each element of the second column. We then place the absolute values of the results in Table 2.

By examining the values of $l'(d, \mathcal{K}, I_s, Q)$ for $1 \leq s \leq 6$, we see that only I_2 and I_3 can possibly be within a distance of

TABLE 1
Sample Database and Stored Distances

Image	$d(I, K_1)$	$d(I, K_2)$
I_1	2	8
I_2	4	4
I_3	1	5
I_4	6	9
I_5	4	1
I_6	7	3

2 to query Q . Thus, only $d(I_2, Q)$ and $d(I_3, Q)$ need to be calculated to determine all close matches to Q . The efficiency of the algorithm is highly dependent on the selection of keys, the relative expense of distance measure calculation, and the statistical behavior of the distance measure over the set of database objects.

Figure 9 shows a real example from our system. In this case, we used a simple color measure to find close matches to the leftmost image. The possible threshold range for this color measure is 0 for an exact match to 1000 for no color match at all. Using a threshold of 100, we eliminated all but 19 out of 37,748 images as potential matches to the query.

The Berman [4] and Baeza-Yates [2] papers also introduced the idea of combining all the precalculated distances into a *trie*. Use of this data structure can, in some circumstances, reduce the number of simple calculations to below $O(nm)$. Berman showed an $O(u \log n + n^\epsilon)$ expected upper bound on simple calculations for the string-matching problem using the Hamming distance on random binary strings, where ϵ is a complicated function always less than one. We discuss this data structure, the *triangle trie*, in Section 9.

4.2. Indexing with Multiple Distance Measures

We extended the above scheme to work with combinations of distance measures [6]. The intuition is that lower bounds on the distance between two objects for distance measures d_1 and d_2 can often be used to calculate a lower bound between the objects for distance measure d when d can be calculated as a combination of d_1 and d_2 .

TABLE 2
Calculating Minimum Distances of Each Image in a Database to Query Image q by Use of the Triangle Inequality

Image	$l(d, K_1)$	$l(d, K_2)$	$l'(d, K)$
I_1	1	3	3
I_2	1	1	1
I_3	2	0	2
I_4	3	4	4
I_5	1	4	4
I_6	4	2	4

Let $D = \{d_1, \dots, d_p\}$ be a set of distance measures. These distance measures will be known as the *base* distance measures. Let $\mathcal{K}' = \{\mathcal{K}_1, \dots, \mathcal{K}_p\}$ be a sequence of sets of keys, one set of keys for each distance measure. Note that each set may have a different number of keys and that the sets may or may not intersect. Let $L(D, \mathcal{K}', I, Q)$ be the set of lower bounds $l'(d_s, \mathcal{K}_s, I, Q)$ calculated from Eq. (2) for each pair $(d_s \in D, \mathcal{K}_s \in \mathcal{K}')$, $1 \leq s \leq p$.

Now consider a new distance measure d' that is of the form

$$d'(I, Q) = f(d_1(I, Q), \dots, d_p(I, Q)),$$

where f is monotonically nondecreasing in its parameters. For example, f might describe a weighted sum of the base measures, or even combinations of minimums and maximums of sets of the base measures. Since $l'(d_s, \mathcal{K}_s, I, Q) \leq d_s(I, Q)$ for all s , substituting $l'(d_s, \mathcal{K}_s, I, Q)$ for each instance of $d_s(I, Q)$ gives us

$$d'(I, Q) \geq f(l'(d_1, \mathcal{K}_1, I, Q), \dots, l'(d_p, \mathcal{K}_p, I, Q)).$$

Thus we can calculate a lower bound on $d'(I, Q)$ given lower bounds on the base distance measures. We can then either order the database images based on these lower bounds or threshold out database images as candidates for matches to the query image. Here we note that the operations on distance measures described earlier—addition, weighting, max, and min—can be combined to form monotonically nondecreasing functions. We also note that the power function $f(x) = cx^e$, $c \geq 0$, $e \geq 0$, is monotonically nondecreasing.




FIG. 9. An example from our system using a simple color measure. The query image is on the far left, potential matches to the query.


```

procedure Index_Multiple(ImageSet[],KeySets[],d())
{
for X := 0 to SizeOf(d)-1
  for Y := 0 to SizeOf(ImageSet)-1
    for Z := 0 to SizeOf(KeySet[X])
      DistanceMatrix [X][Y][Z] =  $d_X(\text{ImageSet}[Y],\text{KeySet}[X][Z]);$ 
}

procedure Calc_Lower_Bound_Multiple(Query,i,f(),DistanceMatrix)
‘Query’ is the Query image
‘i’ is the database Image’s index
‘f()’ is the composite distance measure, as in  $d(Q, I) = f(d_1(Q, I), d_2(Q, I), \dots)$ 
{
for X := 0 to NumDistanceMeasures
  L[X] = Calc_Lower_Bound(Query,I,d_X(),DistanceMatrix);
return f(L[1],...,L[NumDistanceMeasures]);
}

procedure Retrieve_Multiple_With_Threshold(Query,t,f())
{
R = null;
for Y := 0 to SizeOf(ImageSet)-1
  if Calc_Lower_Bound_Multiple(Query,Y,f(),DistanceMatrix)  $\leq t$ 
    then R = R  $\cup$  ImageSet[Y];
return R;
}

```

FIG. 10. Pseudo-code for the indexing and retrieval procedures using the triangle-inequality algorithm with multiple distance measures.

The pseudocode for setting up the index structure, calculating a lower bound for a query and an image, and retrieving images for a query with multiple distance measures is given in Fig. 10.

EXAMPLE OF INDEXING WITH MULTIPLE DISTANCE MEASURES. Let our database be a set of images (I_1, \dots, I_6) , with two base distance measures (d_1, d_2) and two sets of keys, $\mathcal{K}_1 = (K_{11}, K_{12})$ and $\mathcal{K}_2 = (K_{21}, K_{22})$. We precalculate $d_s(I_t, K_{su})$ over all s, t, u to obtain Table 3. Now suppose we have query Q and distance

TABLE 3
Sample Database and Stored Distances with Multiple Distance Measures

Image	$d_1(K_{11}, I)$	$d_1(K_{12}, I)$	$d_2(K_{21}, I)$	$d_2(K_{22}, I)$
I_1	2	8	5	15
I_2	4	4	3	6
I_3	1	5	12	9
I_4	6	9	10	8
I_5	4	1	2	8
I_6	7	3	15	15

TABLE 4
Calculating Lower Bounds on $d' = d_1 + 3d_2$ by Use of Triangle Inequality

Image	$l(d_1, K_{11})$	$l(d_1, K_{12})$	$l(d_2, K_{21})$	$l(d_2, K_{22})$	$l'(d', \mathcal{K}')$
I_1	1	3	3	7	$3 + 3 * 7 = 24$
I_2	1	1	1	2	$1 + 3 * 2 = 10$
I_3	2	0	10	1	$2 + 3 * 10 = 32$
I_4	3	4	8	0	$4 + 3 * 8 = 28$
I_5	1	4	0	0	$4 + 3 * 0 = 4$
I_6	4	2	13	7	$4 + 3 * 13 = 43$

measure $d'(X, Y) = d_1(X, Y) + 3d_2(X, Y)$. We wish to find all objects I in the database such that $d'(I, Q) \leq 10$. We calculate $d_1(K_{11}, Q) = 3$, $d_1(K_{12}, Q) = 5$, $d_2(K_{21}, Q) = 3$, and $d_2(K_{22}, Q) = 8$. Taking the absolute differences between these values and the values in Table 3, we produce the $l(d_s, K_{su})$ values over s, u and combine them to calculate $l'(d_1, \mathcal{K}_1)$ and $l'(d_2, \mathcal{K}_2)$. We then combine these results to produce the $l'(d', \mathcal{K}')$ values. The $l(d_s, K_{su})$ and $l'(d', \mathcal{K}')$ values are shown in Table 4. In this case, $l'(d', I_5, Q, \mathcal{K}') \leq 10$ and $l'(d', I_2, Q, \mathcal{K}') \leq 10$. Thus, I_2 and I_5 are returned as possible matches, with the remaining images eliminated.

We can modify the algorithm to return the best match. In this case, the images are returned in increasing order of their $l'(d', \mathcal{K}')$ values as $(I_5, I_2, I_1, I_4, I_3, I_6)$. Direct comparisons could then be made from the query image to some prefix of this set to validate the best image.

Figure 11 shows a real example from our system with multiple distance measures. The query is the leftmost image. We used a measure representing a weighted sum of a color measure and a texture measure, with an appropriate threshold. We eliminated all but 290 out of 37,748 images as potential matches to the query. This example also illustrates a best-match example, as the returned images are ordered by their calculated distances to the query.

5. FAST IMAGE DATABASE SYSTEM: A PROTOTYPE FOR TESTING

FIDS, the fast image database system, is a prototype content-based image retrieval system, which currently has over 37,000 images. It contains a number of distance measures based on color, texture, and feature detection. Position matching is implemented by “gridding” the distance measures and performing distance calculations on corresponding sub-rectangles.

The FIDS interface is shown in Figs. 12 and 13. There are four windows. The LIST window is a list of available distance measures. When the user clicks on a distance measure, a brief description is offered. The user can select distance measures from this window and copy them to either the BUILD window or the QUERY window. The BUILD window contains a smaller




FIG. 11. Results of a query using a combination of a color measure and a texture measure. The query is on the far left.

set of distance measures. One creates a new distance measure by performing four activities.

1. Selecting two or more measures in the BUILD window,
2. Selecting one of the three operations “SUM,” “MIN,” or “MAX,”
3. Entering the appropriate weights
4. Entering a name for the new distance measure,
5. Pressing the “build” button.

This new distance measure is then added to the LIST window.

The QUERY window contains the query image and the chosen distance measure. The user runs the query by pressing the

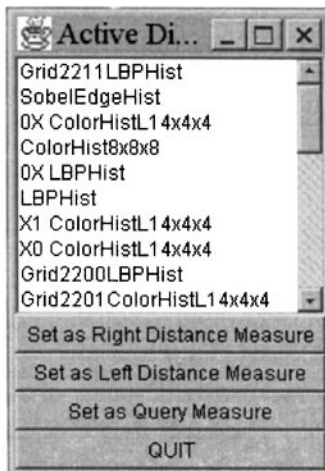
run-query button. The RESULTS window contains the 16 images judged to be the closest by the system. In the full FIDS system, there are three ways to get results. The first way is to order the images by the lower bounds calculated from the triangle inequality. This method is the fastest as no direct comparisons of the query to the database are made. In the second method, the user selects how many images to verify. An image is verified if its true distance to the query is less than the triangle-inequality-derived lower bounds of all the remaining images. The algorithm calculates the true distances to the returned images in lower-bound-derived order until the required number of images have been verified. Optionally, one can set an upper limit on the number of images that need to be verified. The third method is to perform direct comparisons on all the images. In Berman and Shapiro [6], verification with an upper limit of a few percentages of the database almost always returned the correct images.

The user can click on any of the images in the RESULTS window to move that image to the QUERY window. This provides an efficient browsing mechanism for groups of similar images.


6. DISTANCE MEASURES IN FIDS

Experiments conducted using the FIDS system were performed with various subsets of the following distance measures. They are labeled in the experiments using the labels in boldface below:

- **Color.** The color histogram distance measure was first published by Swain and Ballard [41]. Consider RGB space as a three-dimensional cube with the axes labeled red, green, and blue. We quantize this cube by breaking it up into a set of subcubes. For our experiments, we created a $4 \times 4 \times 4$ set of subcubes of equal size. Given an image I , we label each subcube with the fraction of I that has any of the colors contained in that subcube. The distance between two images is the sum of the absolute differences between the values in their corresponding subcubes, known more formally as the L_1 distance. For some experiments, we divided the cube into an $8 \times 8 \times 8$ set of subcubes. To prevent confusion, we note the dimensions of the subcubes in some cases. If no dimensions are given, the $4 \times 4 \times 4$ set was used.



(a) List of available distance measures



(b) New distance measure creation window

FIG. 12. The LIST and BUILD windows for the Fast Image Database System.

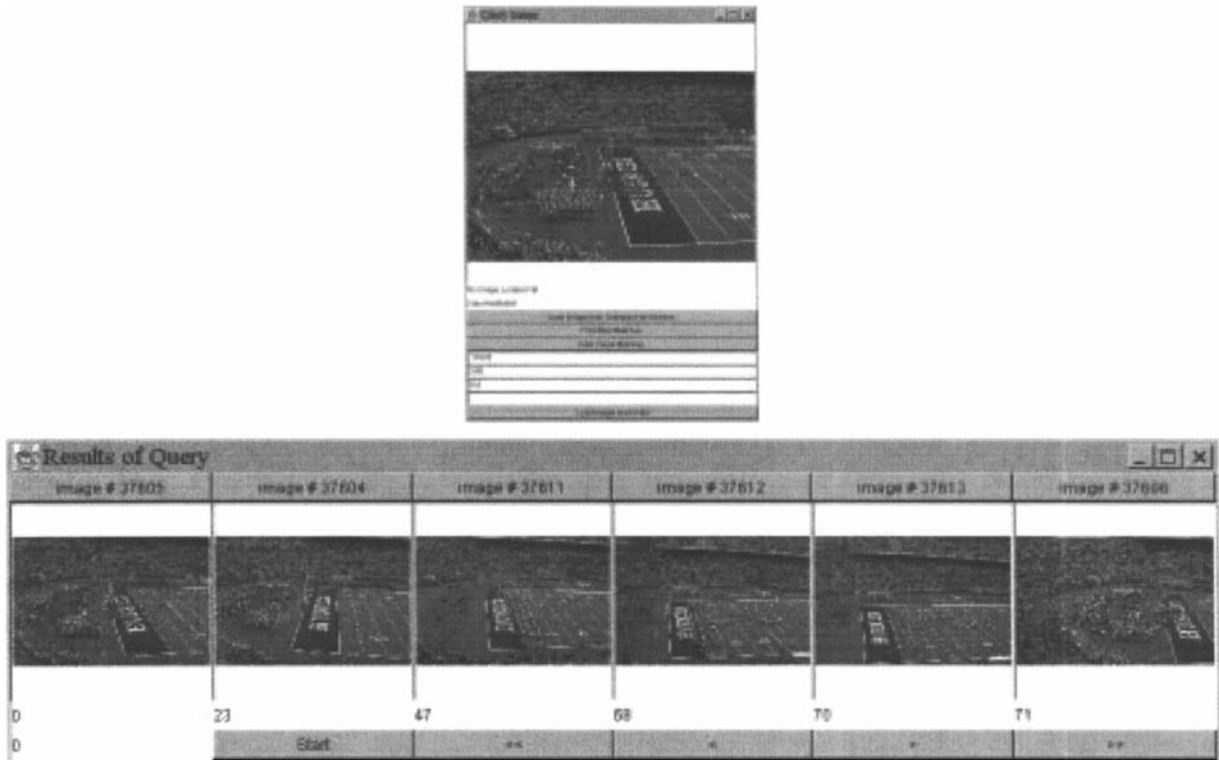


FIG. 13. The QUERY and RESULTS windows for the Fast Image Database System.

- **LBP.** The local binary partition texture measure is a standard texture measure with very good performance [45]. For each pixel P , the eight neighbors are examined to see if their intensity is greater than P 's intensity. The result is encoded as an eight-bit binary number. A histogram of these numbers is created for each image. Two images are compared by taking the L_1 distance between their corresponding histograms.

- **Sobel.** The Sobel edge detector was run over the greyscale version of each image. A histogram of the values of the resultant matrix was calculated. The distance between two images was the L_1 distance of the histograms. Note that in the standard Sobel edge detector, a threshold value is used to create a binary edge image. We did not use a threshold value, but simply used the results of the Sobel convolution.

- **Wavelet.** This measure is based on the wavelet decomposition distance measure developed by Jacobs, *et al.* [24]. A Haar wavelet decomposition of the images is calculated, resulting in a set of coefficients. These coefficients are then weighted and a distance is calculated between two images based on the difference between corresponding coefficients. As the wavelet decomposition is only defined for images of size $2^n \times 2^n$, we resized all the images to 32×32 before calculating the coefficients.

- **Flesh.** We implemented a flesh detector based on the work of Fleck and Forsyth [16]. For each image, we calculated the percentage of pixels that contained flesh according to the detector. The distance between two images is simply the difference in this

percentage. Technically, this is equivalent to a 2-bin histogram with an L_1 distance calculation.

- **Grid color, grid LBP, grid flesh, grid wavelet, grid Sobel.** We implemented a positionally dependent version of each of the distance measures described above. Rather than comparing whole images, gridded versions of distance measures only compare chosen pieces of the images. For these experiments, the gridded versions of the distance measures compare the lower left quarter of each image.

- **Horizontal color and horizontal LBP.** For these two distance measures, each image was split into three equal-sized horizontal pieces. Two images were compared by averaging the color or LBP distance between the corresponding pieces.

- **Vertical color and vertical LBP.** These distance measures are similar to the horizontal measures above, except that the images were split vertically.

- **SUM(d_1, d_2, \dots).** A distance measure with the suffix SUM represents a composite measure that is the sum of the enclosed distance measures. That is, the distance between the images in question is computed for each enclosed measure, and the sum of these distances is returned.

- **MIN(d_1, d_2, \dots).** A distance measure with the suffix MIN represents a composite measure that is the minimum of the enclosed distance measures. That is, the distance between the images in question is computed for each enclosed measure and the minimum value is returned.

- $\text{MAX}(d_1, d_2, \dots)$. A distance measure with the suffix MAX represents a composite measure that is the maximum of the enclosed distance measures. That is, the distance between the images in question is computed for each enclosed measure and the maximum value is returned.

In some experiments, one or more base distance measures were combined. An equal weighting of distance measures was always used in those cases.

7. PERFORMANCE OF THE TRIANGLE INEQUALITY ALGORITHM

We first measured the speed of the bare-bones triangle inequality algorithm. The process by which a system using this algorithm returns a set of matches to a query can be broken into four steps:

- *Step 1.* The system extracts relevant features from the query image. In our system, this step takes from a fiftieth to a quarter of a second, depending on the distance measure. Table 5 shows this range in the second column.
- *Step 2.* The system calculates the distance from the query image to each of the key images. For the basic distance measures on our system, this step takes from about a microsecond per image to more than four-fifths of a millisecond per image. The third column of Table 5 shows the values.
- *Step 3.* The system calculates the lower bound distances from the query image to each of the database images.
- *Step 4.* The system returns the images with the smallest lower bound distances calculated in the previous step.

The third step above is the deciding factor on throughput. The timing of the other steps is relatively stable across database sizes, although the number of keys in step two which are necessary for adequate performance will tend to increase as the number of images in the database increases.

We measured a time of approximately 4 ms to perform the third step on 1000 images if 35 keys are used. This represents a throughput of well over 250,000 images per second, once the

TABLE 5

Feature Extraction and Distance Calculation Time for Representative Distance Measures Used in FIDS on a Pentium Pro 200-mHz PC

Distance measure	Feature extraction time in seconds	Distance calculation time in ms per 1000	Number of distance calculations per second
Color ($4 \times 4 \times 4$ RGB)	0.12	46.00	2174
Color ($8 \times 8 \times 8$ RGB)	0.13	37.00	2669
LBP	0.04	28.00	3623
Flesh	0.25	0.12	833333
Sobel	0.20	4.10	24937
Wavelet	0.02	863.00	115

feature extraction and query-key comparisons are finished. By comparison, the final column of Table 5 gives the throughput per second for each distance measure given a system that stores the features for its images.

7.1. The Accuracy of the Triangle-Inequality Algorithm

To evaluate the efficiency of using the triangle inequality to return close matches to queries, we performed experiments using both single and composite distance measures. Given a query and distance measure, the system returns all the images in the database ordered by calculated lowest bounds on their distances to the query. We use terms such as *lower-bound sequence* or *returned ordering* to refer to the system's output. On the other hand, there is the *true sequence* or *true ordering*, which consists of all the images in the database ordered by their true distance to the query. The hope is that the images which are at the front of the true ordering are also at or near the front of the lower-bound sequence. Measuring the placement of close matches was the main goal of these experiments.

In order to test the quality of the returned ordering, some ground truth was required. To achieve this ground truth, we examined the database by hand and selected 51 pairs of similar images. For each pair, we queried the system with one of the images and measured the position of the other *target* image in the returned sequence. We then calculated the true position of the target image and compared the returned positions to the true positions.

We measured the fraction of matches that were returned as part of the first 25 images. This corresponds to a scenario where a user might ask for several images to be returned for closer examination. We further measured the fraction of matches that were returned within the first 400 matches, corresponding to a scenario in which the user or system is willing to do more work to find a closer match. We also measured the fraction of images that were returned in their precise position. As we selected matching pairs of test images by hand, we had to deal with the possibility that the test image pairs did not truly match with respect to the distance measures on the system. Our experiments were designed to measure the ability of the system to find the closest images to queries, so we excluded an image pair from tests on a distance measure if it was discovered that the target image was not one of the five actual closest images to the query image for that distance measure. Our database contained 37,748 images, and we used 35 keys for each base measure.

7.2. Results for the Bare-Bones Triangle Inequality Algorithm

Tables 6 and 7 show the results of experiments using the full set of valid matches. Depending on the distance measure, from 50% to 100% of all the matches were within the first 25 images returned by the system. In almost every case, over 90% of the matches were within the first 400 returned images. In a second set of similar experiments, we simulated a user with a more restricted definition of closeness. For each distance measure,

TABLE 6
Summary Performance of Distance Measures Using Full Set of Validated Images

Distance measure class	100% in first 400	90–99% in first 400	80–89% in first 400
Single measures	3	2	3
AND’ed measures	7	12	1
OR’ed measures	16	1	3
SUM’ed measures	10	10	0
Totals	36	25	7

Note. There were no distance measures in which less than 80% of the target images were returned in the first 400 images out of 37,748.

the query-target pairs were sorted in increasing order of query-target distance. Only the first half of these lists were used to compute these results. The results were markedly better for the experiments using the closer half of the valid matches. Tables 8 and 9 show the results of experiments with the restricted set of valid matches. In this case, only nine distance measures out of the 68 tested measures had less than 80% of the target images returned within the first 25 images. Furthermore, in 64 out of 68 distance measures, every single one of the matches were returned within the first 400 images.

8. KEY SELECTION

We begin this section with a discussion of what makes a good single key. Later, we discuss the choice of keys in combination. We make the simplifying assumption that all distances are within the range of 0 to 1 inclusive.

8.1. Good Keys for Threshold Style Queries

Consider database image I , query image Q , key image K , distance function d . We say that key K separates Q from I for value v if $|d(I, K) - d(Q, K)| > v$. Suppose that $d(I, Q) > t$ for some threshold t . The triangle inequality implies that the value $|d(I, K) - d(Q, K)|$ can range from 0 to $d(I, Q)$. Key K will eliminate image I as a candidate match to Q only if it

TABLE 7

Summary Performance of Distance Measures Using Full Set of Validated Images, Showing How Many Distance Measures of Each Type Returned the Target Images in the First 25 Images

Distance measure class	100% in first 25	90–99% in first 25	80–89% in first 25	70–79% in first 25	50–69% in first 25
Single measures	1	0	1	4	2
AND’ed measures	2	2	10	6	0
OR’ed measures	6	0	0	12	2
SUM’ed measures	4	6	10	0	0
Totals	13	8	21	22	4

TABLE 8
Summary Performance of Distance Measures Using Closer Pairs of Validated Images

Distance measure class	100% in first 400	90–99% in first 400	80–89% in first 400	70–79% in first 400
Single measures	6	0	2	0
AND’ed measures	20	0	0	0
OR’ed measures	18	1	0	1
SUM’ed measures	20	0	0	0
Totals	64	1	2	1

Note. Almost every distance measure returned all targets within the first 400 images.

separates I from Q for value t . The purpose of the algorithm is to eliminate as many nonmatching candidate images as possible through key comparison. Thus, a good key will eliminate more candidate images than a poor key. The concept of separation described above motivates the following discussion.

Given a set of database images \mathcal{S} , distance measure d , and key K , we can compute a density function f on $d(I, K)$, $I \in \mathcal{S}$. Since we do not know the queries in advance, we make the simplifying assumption that the queries are taken from the database images and ignore exact matches in our searches. Given threshold t , we can calculate the fraction of images that K will separate from a random query by looking at this density function. For example, if all of the area of the density function lies in a narrow range $(x, x + e)$, $e < t$, as shown in Fig. 14a, then K will never separate any query from any image in the database. If the density function has a uniform distribution, as shown in Fig. 14b, then for $0 < t < 1/2$, $P(K \text{ separates } I \text{ from } Q) = (1 - t)^2$. If the density function is multipolar, with n equally sized narrow spikes separated by distance greater than t , as shown in Fig. 14c, then $P(K \text{ separates } I \text{ from } Q) = (n - 1)/n$. If the density function has a Gaussian shape, as shown in Fig. 14d, then, roughly speaking, greater standard deviations will indicate greater average separation of images by the key.

The issue gets more complicated when choosing several keys. Using keys K_1 and K_2 will be no better than just using K_1 if they both separate the same images from queries. The question

TABLE 9

Summary Performance of Distance Measures Using Closer Pairs of Validated Images, Showing How Many Distance Measures of Each Type Returned the Target Images in the First 25 Images

Distance measure class	100% in first 25	90–99% in first 25	80–89% in first 25	70–79% in first 25	50–69% in first 25
Single measures	2	1	2	1	2
AND’ed measures	13	5	0	2	0
OR’ed measures	6	0	10	2	2
SUM’ed measures	16	3	1	0	0
Totals	37	9	13	5	4

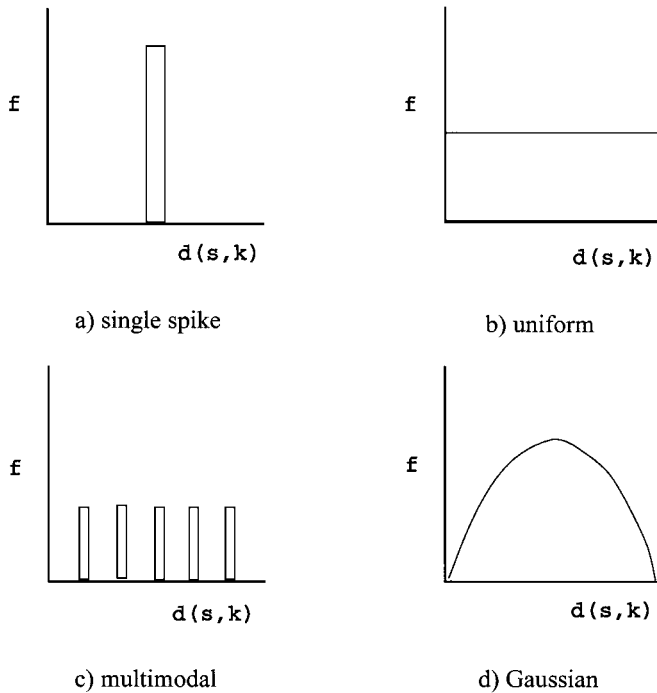


FIG. 14. The shape of the density functions determines the performance of the keys.

of whether or not two keys separate the same images is computationally expensive to answer in the general case, but can be approximately answered by sampling. One can also use the fact that very similar keys will separate many of the same images and thus try to avoid keys that are too close together. For example, in a clusterable database, keys should come from different clusters. Indeed, the key selection algorithms with the best results make use of clustering and ensuring that different keys separate different images.

8.2. Good Keys for Best-Match Queries

Given images I_1 and I_2 , query Q and key K , assuming that $d(I_1, Q) < d(I_2, Q)$, key K orders I_1 and I_2 correctly if $l(d, K, I_1, Q) < l(d, K, I_2, Q)$. We can extend this definition naturally to sets of keys and multiple distance measures. Although our analyses were for *threshold* queries, the results were very good for *best-match* queries as well. Further analysis of keys optimized for *best-match* queries is an open area of research.

8.3. Algorithms for Key Selection

We examined five different algorithms for key selection: random keys, choosing keys by examining the variance of the density function, ranking by testing thresholding efficiency, a greedy thresholding algorithm, and a clustering algorithm. The algorithms assume a database \mathcal{S} and a set of candidate keys.

RANDOM. Our prototype image database system currently uses a set of (up to 35) keys chosen randomly and uniformly from the database itself. The triangle inequality algorithms give

excellent performance compared to linear search even with random keys, so this is a natural benchmark against which to test the other algorithms.

VARIANCE. Taking a subset \mathcal{S}' of our database \mathcal{S} , we calculated the density function of $d(I, K)$, $I \in \mathcal{S}'$, for each candidate key K . We selected those candidate keys which had the density functions with the greatest variance.

SEPARATION. We examined our database by hand to find pairs of images that we judged to be approximate matches. The average distance between these pairs was calculated. This value t represented a potential threshold value that one might use in a query to find approximate matches. We then selected those candidate keys K which maximized $P(|d(I_1, K) - d(I_2, K)| > t)$ over all pairs $I_1, I_2 \in \mathcal{S}'$, where \mathcal{S}' was a subset of our database \mathcal{S} .

GREEDY. **VARIANCE** and **SEPARATION** may choose several keys which separate the same pairs of images. We thus modified **SEPARATION** to keep track of which pairs of images were separated by each key. The first key selected was the same as that selected by **SEPARATION**. The performances of the remaining keys were then recalculated to discount pairs of images already separated by the first key. This process was continued for subsequent keys until a preset number of keys was selected.

CLUSTER. We used a simple clustering algorithm on the database. We selected the two database images K_1 and K_2 that were furthest apart and used them as initial seeds for clustering. These two images were placed into our set of keys, and the remaining images were assigned to clusters based on their distances to the key images. We then found the image that was furthest from the current set of keys, added it to the set, and reclustered the database on the updated set of keys. We continued this process until the correct number of keys were selected. Note that this algorithm differs from the others in that the selected keys came from the database itself.

8.4. Experiments

For our experiments, we collected two sets of images, one with 600 members, and one with 800 members. From each set, 100 images were chosen arbitrarily to be candidate key images. The remaining 500 and 700 images became the test database. The five algorithms were run on the candidate images to choose sets of 1 to 9 keys.

We queried the database against itself, testing the system's performance using the keys chosen by the key selection algorithms. To eliminate exact matches, we temporarily removed each query image from the database. To test the performance of the keys on a *best-match* search, we determined the best match to the query and calculated its position in the ordering of the lower bounds. To test the performance of the keys on a *threshold* search, we counted the number of images separated from the query by a given threshold value. This threshold value was determined off-line by calculating the average distance between pairs of images known to be similar. For the **RANDOM** key selection algorithm, we ran the tests 10 times and averaged the results.

TABLE 10
Best Algorithms for *Best-Match* with Nine Keys
on a Database of 500 Images

A	B	C	D	E	F	G
Color histogram	GREEDY	0.9%	CLUSTER	1.1%	1.6%	1.8
LBP texture	CLUSTER	1.2%	GREEDY	1.6%	2.3%	1.9
Horizontal color	GREEDY	3.6%	CLUSTER	3.7%	4.6%	1.3
Vertical color	GREEDY	2.5%	RANDOM	3.5%	3.5%	1.4
Horizontal texture	CLUSTER	2.4%	GREEDY	3.7%	4.8%	2.0
Vertical texture	CLUSTER	2.1%	GREEDY	2.2%	3.9%	1.9
Min (Color, Texture)	GREEDY	1.0%	CLUSTER	1.7%	2.2%	2.2
Color + Texture	GREEDY	1.0%	CLUSTER	1.1%	1.7%	1.7

Note. Column headings: A, distance measure; B, best key selection algorithm; C, average rank of best match using best algorithm; D, second best key selection algorithm; E, average rank of best match using second best algorithm; F, average rank of best match using randomized key selection; G, ratio of performance of randomized algorithm to best algorithm (F/C).

8.5. Results of Key Selection Tests

We discuss the performance of the various key selection algorithms, first for *best-match* and then for *threshold*. As the rankings of the algorithms did not change much as a function of number of keys, we only show the results for nine keys, the maximum number tested. As the performance of the algorithms on the two databases was very similar, we only show tables for the larger database.

8.5.1. Performance of key selection algorithms for Best-match. As is shown in Table 10, CLUSTER provided the best keys for the texture measures, while GREEDY provided the best keys for the color measures and the combination color/texture measures. The second best algorithm was also always GREEDY or CLUSTER, except for the *vertical texture* measure in the larger database, which had RANDOM as the second best.

Columns C, E, and F of the table show the average rank of the best match using the appropriate key selection algorithm. For example, a 2% would mean that the true best match was ranked in the top 2% of the returned images. Column G represents the ratio of the number of images that would be examined using the RANDOM keys to the number of images that would be examined using the best discovered keys. Thus, in the first row, the best keys returned the closest match in the top 0.9% of the images. For the 700 image database, this translates to the top 6 or 7 images. The random keys returned the closest match in the top 11 images. If this database was a representative sample of a 700,000 image database, then the number of images needed to be directly compared would be approximate 630 and 1120, respectively. On average, there was a 42% reduction in the number of images examined using the best discovered keys, compared to using the random keys.

The overall performance of the algorithms was excellent. In the worst case for the database of 700 images, *Horizontal Color*, the closest match was ranked in the top 3.6%—that meant that only approximately 24 images had to be compared directly after

pruning to find the best match. The database of 500 images had slightly worse performance with the average ranking of the best image ranging from 1.2% to 5.1%, again with the worst performance found in *Horizontal Color*.

8.5.2. Performance of key selection algorithms for Threshold. As Table 11 shows, the GREEDY key selection method was the clear winner for *threshold*, yielding the best performance for every distance measure. There was no clear second place algorithm—RANDOM, VARIANCE, and CLUSTER all appeared in second place for several measures. The GREEDY keys reduced the number of images that had to be directly compared by the RANDOM keys by 16% to 44%.

The second thing to note in Table 11 is the wide range of performance between distance measures. The triangle inequality algorithm thresholded 98.5% of the images for the *Color Histogram* distance measure, yet it only thresholded 54.6% of the images for the *Vertical Texture* distance measure. It is difficult to compare across distance measures since the distribution of distance values across pairs of images varies greatly from measure to measure. Especially interesting was the fact that using five keys for *Vertical Texture* resulted in a 53% thresholding. Thus the additional four keys only eliminated an additional 2% of the database. In [2], Baeza-Yates *et al.*, demonstrated how, given a random model for database objects and keys, a logarithmic number of keys should threshold almost all of the database. For our experiments to have supported this, the addition of four more keys would have had to increase the thresholding from 53% to about 70%. That this did not occur demonstrates that traditional models of randomness do not really apply to sets of real images.

8.6. Analysis

The performance of keys in image retrieval is intimately tied to the statistical behavior of the distance measures over the image set. At present, we have a limited understanding of this

TABLE 11
Best Algorithms for *Threshold* with Nine Keys
on Database of 700 Images

A	B	C	D	E	F	G
Color histogram	GREEDY	98.5%	RANDOM	97.4%	97.4%	1.7
LBP texture	GREEDY	83.2%	CLUSTER	81.2%	77.8%	1.3
Horizontal color	GREEDY	94.4%	RANDOM	90.8%	90.8%	1.6
Vertical color	GREEDY	93.2%	VARIANCE	90.0%	89.8%	1.5
Horizontal texture	GREEDY	55.3%	CLUSTER	49.9%	47.5%	1.2
Vertical texture	GREEDY	54.6%	VARIANCE	52.7%	47.2%	1.2
Min (Color, Texture)	GREEDY	97.8%	CLUSTER	96.4%	96.1%	1.8
Color + Texture	GREEDY	94.5%	CLUSTER	91.5%	91.2%	1.6

Note. Column headings: A, distance measure; B, best key selection algorithm; C, average percentage of database eliminated using best algorithm; D, second best key selection algorithm; E, average percentage of database eliminated using second best algorithm; F, average percentage of database eliminated using randomized key selection; G, ratio of performance of randomized algorithm to best algorithm $(100\% - F)/(100\% - C)$.

behavior; this limits the sophistication of our key selection algorithms. Thus, more research into the behavior of the distance measures is called for. A more complete set of distance measures will also be used in our future tests. Distance measures have been proposed for color, texture, shape [8], object presence [18], and object spatial relationships [7]. We would like to include representatives of each type of measure in our tests.

In our work, we selected keys from the database itself. The space of possible keys is huge—it is the space of possible images. We would like to take advantage of this freedom in some tractable manner. For example, it may be possible to construct artificial images that are excellent keys for either a specific database, or even for large image domains. Furthermore, our analysis contained the assumption that the query domain was similar to the database. This is not necessarily the case.

Even if we restrict our candidate keys to some random subset of n images, the number of possible subsets of m keys is exponential in m . There is no guarantee that there is not some difficult to find set of keys that will prune the database far more than any other set. It may be that heuristics like those traditionally used for NP-complete algorithms may be applicable for key selection.

Finally, there has been no published work on the proper number of keys to use for a database of a given size. There is a trade-off between the elimination power of a set of keys and the execution time required to compare the query to the key set. Some queries may require more keys than other queries for good performance.

Of the algorithms tested, CLUSTER and GREEDY clearly gave the best results. The improvement over random key selection was up to a factor of two. As random key selection reduces *best-match* searches to just a few percentages of the database, the use of random keys may be perfectly acceptable for smaller databases.

One thing to note is that, given n sample database images and m candidate keys, CLUSTER and GREEDY take $O(n^2)$ and $O(mn^2)$ time, respectively. Thus, for very large databases, one might consider sampling the database. We used sampling in GREEDY and not in CLUSTER, yet GREEDY was essentially as good as CLUSTER in some cases and better than CLUSTER in the rest. It is promising that relatively simple algorithms were able to increase performance to the extent shown in this paper.

9. THE TRIANGLE TRIE

Although much faster than direct comparisons, the basic triangle-inequality algorithm described above has a running time of $O(nk)$, where n is the number of images and k is the number of keys. Running time may become unacceptable for very large databases with a large number of keys. Therefore, we take advantage of a data structure called the *triangle trie* [4] to reduce the number of operations.

A *triangle trie*, also called a *Really Fixed Query Tree* [2], is a data structure developed for approximate-match searching both




FIG. 15. Triangle trie with two levels.

by Baeza-Yates *et al.* and by Berman in independent research efforts. A single triangle trie is associated with a distance measure, a set of key images, and a set of database elements. It is a form of *trie*, which is a tree in which the edges leading from the root to a leaf “spell out” the index of the leaf. The leaves of the tree contain the database elements. Each internal edge in the tree is associated with a nonnegative number. Each level of the tree is associated with a single key. The path from the root of the tree to a database element in a leaf represents the distances from that database element to each of the keys.

Figure 15 illustrates a triangle trie with four elements (W, X, Y, Z), and two keys (J, K). The distance from W to J is 3 and the distance from W to K is 1. This is expressed in the trie by the path from the root to the leaf containing W .

Construction of the trie is straightforward. Compute the distances from the keys to the database elements. Starting with an empty trie, insert the database elements one at a time using the vector of its key distances. Create nodes as necessary until every element is in the trie. Formally, let $S = (x_1, \dots, x_n)$ be our set of objects in the database. Let key_1, \dots, key_j be another set of objects, known as “key objects.” For each x_i in S compute the vector $v_i = (d(x_i, key_1), d(x_i, key_2), \dots, d(x_i, key_j))$. Then combine the vectors v_1, \dots, v_n into a *trie*, with x_i being placed on the leaf reached by following the path represented by v_i .

9.1. Pruning the Search with a Triangle Trie

Suppose we are given query q and threshold integer t and wish to find all objects in our database with a distance from q of not more than t . Now, consider a node p at level l with a value of C . Every object at leaves descendant from p has a distance of C from the key object key_l . Thus, if $|C - d(q, key_l)|$ is greater than t , then we know from the triangle inequality that $d(q, s')$ is greater than t for all objects s' which are descendants of p . Thus, we can safely prune the search at node p .

The algorithm for searching the database using the triangle trie is straightforward. Compute the distances from q to each key: $d(q, key_1), \dots, d(q, key_j)$. Perform a depth-first search of the trie. If there is a node p at level l with value C such that $|c - d(q, key_l)| > t$, then prune the search at node p . When a leaf is reached, measure the distance from q to every object in the leaf and return those objects i for which $d(q, i)$ is less than or equal to t .

Searching the trie: An example. We continue with the example. Suppose we wish to search our database for a close match to object V where our maximum allowed distance to V is 1. We compute v_V by calculating $d(V, key_1)$ and $d(V, key_2)$. We discover that $v_V = (3, 8)$. We then perform our depth-first search. At the top level, we only search nodes with a value within 3 ± 1 , so both nodes at this level are searched. At the second level, we only search children of those nodes with a value within 8 ± 1 . At this level Y and Z are returned as potential matches, while X and W are eliminated. The final step is to compute $d(V, Y)$ and $d(V, Z)$. The algorithm does not need to compute $d(V, X)$ and $d(V, W)$.

There may be cases where the distance measure is not integer valued, or where the distance measure has such a wide variance that any resultant triangle trie would have a very quick fan-out. In these cases, it may be necessary to map the calculated distances to a smaller set of values. We call this process *binning* the distances. We have conducted experiments with triangle tries of various bin sizes and depths to better understand the effect of binning on performance.

9.2. Two-Stage Pruning with a Triangle Trie

The breadth of a triangle trie expands with its depth, up to a maximum breadth equal to the number of database elements. The value of a pruning step is directly related to the number of leaves of the pruned subtree. Thus, as the breadth increases, the performance of the tree-pruning algorithm decreases, until it is unfavorable when compared to directly calculating lower bounds for each database object. On the other hand, the efficiency of the triangle inequality algorithm increases with the number of keys used. Our work suggests that by using a relatively short trie and by storing additional key distances in the leaves, we can obtain the best of both worlds. Our two-stage algorithm works as follows: Given database images I_1, \dots, I_n , keys K_1, \dots, K_m , and distance measure d , we create a triangle trie T of depth T_{depth} , where $T_{\text{depth}} < m$. For each stored image I_i , we reference I_i in the trie along with $d(I_i, K_j)$ for all K_1, \dots, K_m . Given query Q , we perform our search of the trie as described above. Once completed, we calculate lower bounds on the returned images using all the keys. This further reduces the size of the returned set.

9.3. Using Triangle Tries with Composite Measures

A triangle trie is designed to enable thresholded database searches for a single distance measure. However, it is possible to use multiple triangle tries to enable thresholded database searches over a composite measure. The intuition is to figure out for each distance measure, what threshold should be used for the corresponding triangle trie. Searches are done on the individual tries and the returned images are either merged or intersected, depending on the particular operation. This operation substitutes for the first stage of the two-stage pruning algorithm. The second stage proceeds as before.

To see how this works, consider the following example: Define $R(T, Q, t)$ as the set of images returned from a search on trie T

with threshold t . Now consider the composite distance measure $d(I, Q) = \text{Min}(d_1(I, Q), d_2(I, Q))$. Assume the threshold used is t . Let T_1 and T_2 represent the tries associated with d_1 and d_2 , respectively. Since $d(I, Q) \leq t$ whenever $d_1(I, Q) \leq t$ or $d_2(I, Q) \leq t$, we must find all images where $d_1(I, Q) \leq t$ or $d_2(I, Q) \leq t$. Thus, we calculate $R(T_1, Q, t)$ and $R(T_2, Q, t)$ and merge the results. Call this resultant set S' . This set consists of all images that have a possibility of being within distance t to Q by distance measure d . We then prune S' with the triangle inequality on the composite function d .

The objective in using the triangle trie is to reduce the number of images for which we have to compute the triangle inequality with the full set of keys. Therefore, when using multiple triangle tries, our objective should be to return as small as possible set of images that need to be further pruned. We have developed algorithms for each of the operations—Min, Max, Sum, and Weight—that reduce the size of the returned set. We describe the various algorithms for binary operations and then show how to combine them for a more complex composite distance function.

9.3.1. The Max function. Given distance functions d_1 and d_2 , associated triangle tries T_1 and T_2 , query Q , and threshold t , suppose we wish to find all images I such that $d(I, Q) \leq t$, where $d(I, Q) = \text{Max}(d_1(I, Q), d_2(I, Q))$. For $d(I, Q) \leq t$ to be true, both $d_1(I, Q)$ and $d_2(I, Q)$ must also be true. Thus, the algorithm for the Max function is to calculate $R(T_1, Q, t) \cap R(T_2, Q, t)$ by searching on T_1 and T_2 and taking the intersection of the results. Second stage pruning is then applied to the intersection set.

Note that we could perform second stage pruning on $R(T_1, Q, t)$ and $R(T_2, Q, t)$ separately and take the intersection of the results. Assume the sizes of the two sets are n and m with k keys and an intersection of $w \leq \text{Min}(n, m)$. Intersecting before key comparisons results in an intersection step time of $O(n + m)$, followed by a key comparison step of time $O(kw)$. Performing key comparison first results in a key comparison time of $O(k(n + m))$ followed by an intersection step of unknown time (all elements could have been eliminated). Which method is faster in any particular case would probably rely heavily on details of the system.

9.3.2. The Min function. Suppose $d = \text{Min}(d_1, d_2)$. If image I has the property that $d(I, Q) \leq t$ either $d_1(I, Q) \leq t$ or $d_2(I, Q) \leq t$ must be true. Thus, I must be in $R(T_1, Q, t) \cup R(T_2, Q, t)$. To find potential approximate matches to Q in this case, we compute the union of the two R functions.

9.3.3. The Addition function. Suppose $d = d_1 + d_2$. Suppose image I has the property that $d(I, Q) \leq t$. Certainly $d(I, Q) \leq t \Rightarrow d_1(I, Q) \leq t, d_2(I, Q) \leq t$. Thus, if image I has the property that $d(I, Q) \leq t$, then as in the Min function above, I must be in $R(T_1, Q, t) \cap R(T_2, Q, t)$. This leads to an algorithm where we calculate $R(T_1, Q, t)$ and $R(T_2, Q, t)$ and take their intersection.

There are other algorithmic possibilities. For example, $s \in R(T_1, Q, t) \cap R(T_2, Q, t) \Rightarrow s \in R(T_1, Q, t)$ implies that we could just calculate $R(T_1, Q, t)$ and not bother to calculate

$R(T_2, Q, t)$. This may be faster in some cases. Similarly we could just compute $R(T_2, Q, t)$. We could compute both and return the smaller set. We also can reduce the thresholds for the individual sets, as follows: Suppose that $d_1(I, Q) > v$ for a given image I and some arbitrary value v . Then $d(I, Q) \leq t$ implies that $d_2(I, Q) \leq t - v$. Thus, $d(I, Q) \leq t \Rightarrow d_1(I, Q) \leq v, d_2(I, Q) \leq t - v$ for any v . We therefore can claim that I must be in $R(T_1, Q, v) \cup R(T_2, Q, t - v)$ for any legitimate $0 \leq v \leq t$. Thus, to find potential approximate matches in this case, we can pick some value for v and compute the union of the two R functions with the modified thresholds.

It is an open question as to how to efficiently decide the best value for v . Choosing $v = 0$ or $v = t$ has the advantage of eliminating the search of one trie entirely, as well as the consequent merging of results. Yet there is evidence that halving a threshold more than halves the returned results. In our system we simply choose $v = t/2$.

9.3.4. The weight function. Suppose $d = Cd_1$ for some positive constant C . Then $d(I, Q) \leq t$ implies $d_1(I, Q) \leq t/C$. In this case we find candidates for approximate matches to Q by calculating $R(T_1, Q, t/C)$.

9.4. Experiments on Triangle Tries

We tested the performance of the triangle trie in image searches. We used the basic distance measures described earlier and gridded versions of the measures. Currently, our database contains 37,748 images. We used only 20,000 of these images in our triangle-trie experiments to reduce disk accesses, which might otherwise have disrupted the time measurements.

Trie searches are performed given both a query and a threshold value. No images that are provably more distant to the query than the threshold value are returned from a trie search. The efficiency of the trie search is dependent on the chosen threshold value. Thus, we had to determine appropriate threshold values for each distance measure. We examined the distances between pairs of closely matching images in our database to determine appropriate threshold values. These values are listed in Table 12. The threshold values for the gridded distance measures were chosen to be the same as those for the nongridded distance measures. This consistency afforded an opportunity to see if there were marked differences in performance between the gridded measures and regular measures.

TABLE 12

Threshold Values Chosen for the Triangle-Trie Experiments

Distance measure	Threshold value
Color ($4 \times 4 \times 4$ RGB)	200
LBP	90
Wavelet	15
Sobel	10
Flesh	10

TABLE 13
Characteristics of the Best Tries Found for Each Distance Measure, Rated by the Speed Factor

Measure	Speed factor	Bin size	Trie depth	Time	Returned matches
Sobel	47.8	10	6	0.4	318.6
Gridded Sobel	38.6	30	15	0.7	343.0
Gridded wavelets	22.3	20	8	1.2	598.6
Gridded color	15.01	130	11	1.8	882.4
Color	8.8	80	8	4.0	1273.0
Wavelets	3.8	40	4	11.8	2343.6
Gridded flesh	3.2	30	4	7.0	4588.2
LBPHist	2.8	30	7	20.1	5740.7
Flesh	2.6	80	11	7.4	5980.2
Gridded LBPHist	2.4	80	8	12.1	5465.0

Note: Time is in milliseconds for a search of 20,000 images and is just for the trie search itself. The speed factor is the ratio of time taken for a search of 20,000 images using the triangle inequality algorithm without the triangle trie to a search using the triangle trie as a precursor to the triangle inequality algorithm.

We created a set of triangle tries for each distance measure, varying the tries by depth and bin size. The depths ranged from 1 to 11, and the bin sizes ranged from 10 to 130 stepping by 10. All the distance measures except the Haar Wavelets had a range of 0 to 1000. The Haar Wavelets had a theoretical range greater than 1000, but no distances greater than 1000 were calculated during the experiment. Experiments were not conducted for a given depth and bin size if the created trie would contain more than 20,000 nodes. This limit was chosen due to memory constraints.

There are 35 key images in our database. Any subset of them can be used as the keys for the tries. We created five random orderings of the keys. For each triple of 10 distance measures, 11 depths, and 13 bin sizes, we created five tries using the appropriate five prefixes of the random key orderings. For each such trie, we tested five images, resulting in a total of 25 tests for each triple of distance measure, depth, and bin size. This resulted in a total of 33,295 experiments, not including experiments that were abandoned due to too large a trie being created. Each experiment was repeated 250 times to improve the accuracy of the measured time. We recorded the size of the created tries, the number of matches returned for a search, the number of internal trie nodes accessed, the number of leaf trie nodes accessed, and the time it took to walk the trie and return the matches.

9.5. Triangle-Trie Test Results with Single Distance Measures

In Section 6, we estimated a time of approximately 4 ms to calculate the lower bounds on 1000 images using the bare-bones triangle inequality algorithm, not including key calculations and feature extraction. We used this value to calculate the speed-improvement factor gained by using the two-stage pruning algorithm instead of the bare-bones algorithm.

Table 13 shows the results of the tries that had the highest speed improvement factors. The gridded version of the local-

TABLE 14
Characteristics of Tries Found for Each Distance Measure
with the Median Speed Factor

Measure	Speed factor	Bin size	Trie depth	Time	Returned matches
Sobel	14.6	90	8	1.2	1014.0
Gridded Sobel	14.6	40	6	1.3	992.2
Gridded color	6.8	90	4	3.1	1962.0
Gridded wavelets	5.2	20	3	4.2	2564.4
Color	5.0	20	4	6.8	1826.5
Gridded flesh	2.2	60	6	10.5	5703.2
Flesh	1.8	130	8	10.3	7668.6
Gridded LBPHist	1.8	100	8	13.3	6250.8
Wavelets	1.6	130	5	18.7	6542.4
LBPHist	1.4	130	9	16.5	8466.4

Note. Time is in milliseconds for a search of 20,000 images and is just for the trie search itself. The speed factor is the ratio of time taken for a search of 20,000 images using the triangle inequality algorithm without the triangle trie to a search using the triangle trie as a precursor to the triangle inequality algorithm.

binary-partition texture measure had the minimum improvement, being about twice as fast with a trie as without one. At the other end of the scale, the Sobel-edge-histogram texture measure was almost 50 times as fast with a trie as without one. This represents a potential processing rate of almost 12 million images per second.

Note that there are many factors which affect trie performance. Even within a given trie depth, bin size, and distance measure, the speed of a trie search varied greatly with the keys chosen and the query image (see Tables 14 and 15). Trie search performance varies with the chosen threshold value—as the algorithm uses the threshold value to prune subtrees—reducing the threshold results in less of the tree being walked, yielding a time savings. In particular, the algorithm walks only one path with a threshold value of zero, while it walks the full trie with a threshold value of infinity.

TABLE 15
Averages of Trie Statistics (Each Column Individually Averaged)

Measure	Speed factor	Time	Returned matches	Internal nodes	Leaf nodes
Sobel	15.4	20.8	1679.5	17.3	4.9
Gridded Sobel	15.0	22.8	1514.8	28.1	6.4
Gridded color	6.6	61.6	3024.0	123.7	103.0
Gridded wavelets	6.5	63.4	3997.5	47.7	14.6
Color	4.8	77.3	3167.2	152.2	132.3
Gridded flesh	2.2	98.2	5954.8	9.4	1.8
Wavelets	1.7	231.9	5439.1	35.0	15.2
Flesh	1.7	115.8	7616.3	11.9	2.4
LBPHist	1.4	167.4	8060.9	350.7	257.2
Gridded LBPHist	1.6	163.8	7050.9	408.5	231.4

Note. The average is over all the tries in the experiments with depths ranging from 1 to 13, bin sizes ranging from 10 to 140, and no tries with more than 20,000 nodes included.

9.6. Triangle-Trie Results with Composite Distance Measures

The *two-stage pruning* algorithm, detailed in Section 5.2, is a method for using triangle-tries for composites of basic distance measures. Given a distance measure $d' = g(d_1, \dots, d_k)$, where the system has triangle tries for the distance measures d_1, \dots, d_k , the system computes appropriate threshold values for the basic distance measures and composites the returned potential matches dependent on the character of $g(\)$. The bare-bones triangle-inequality algorithm is then run on the returned potential matches.

The two major methods for compositing returned potential matches are intersection and union. As expected, intersection reduces the number of potential matches, while merging increases them. Suppose the returns from several triangle tries are merged, and the resultant set of images is the same size as the full database. In that case, no benefit is accrued from the use of the triangle trie. On the other hand, suppose the results from several tries are to be intersected, and a search of one of the tries returned only a small handful of potential matches. In that case, since the intersection of X and Y is contained within X , it may be worth directly proceeding to lower-bound calculation with that small returned handful and avoiding the cost of processing the remaining tries or of intersecting the results. Tables 16 and 17 show the average results of merging and intersecting various returned matches from the triangle tries for the basic distance measures.

9.7. Discussion of Results for Triangle Trie-Experiments

The Triangle-Trie offered improved performance on all of the tested distance measures. Some of the improvements were quite

TABLE 16
Average Size of Image Sets after Merging or Intersecting Image Sequences Returned from Triangle-Trie Searches, Using One or Two Distance Measures

Distance measures	Merged set average size	Intersected set average size
Color		3568
LBP		6102
Wavelet		3023
Sobel		427
Flesh		8090
Color flesh	10001	1656
Color LBP	8342	1327
Color Sobel	3889	106
Color wavelet	5570	1021
LBP flesh	11727	2464
LBP Sobel	6355	173
LBP wavelet	7949	1177
Wavelet Sobel	3382	68
Wavelet flesh	9771	1342
Sobel flesh	8366	151

Note: The rows labeled with single distance measures show the average sizes of the original sets.

TABLE 17

Average Size of Image Sets after Merging or Intersecting Image Sequences Returned from Triangle-Trie Searches, Using Three or Four Distance Measures

Distance measures	Merged set average set	Intersected set average size
Color LBP wavelet	9606	438
Color LBP Sobel	8537	46
Color LBP flesh	12964	653
Color wavelet Sobel	5854	31
Color wavelet flesh	11199	538
LBP wavelet Sobel	8169	35
LBP wavelet flesh	12812	581
LBP Sobel flesh	11894	64
Wavelet Sobel flesh	10010	31
Color Sobel flesh	10224	52
Color LBP wavelet Sobel	9781	16
Color LBP wavelet Flesh	13753	251
Color LBP Sobel flesh	13102	22
Color wavelet Sobel flesh	11398	18
LBP wavelet Sobel flesh	12960	17

significant and might enable otherwise difficult tasks such as nearest neighbor searches for large sets of images. The performance of triangle-inequality-based searches is dependent on a number of factors, including the nature of the image sets, the distance measures, the number and nature of the chosen keys, the chosen threshold, and the properties of the query images. Thus, it is impossible to categorically state that the performance boosts demonstrated here will translate to somebody else's image database. We have demonstrated, however, that the triangle-trie is worth considering in the case that a performance boost is needed.

10. SUMMARY

The overall goal of our research has been to create technology useful in a generalized system for content-based image retrieval. We believe that there are three main requirements for a useful content-based retrieval system. The system must search and retrieve images quickly. The system should allow as wide a selection of queries as possible. Finally, the system should be easy to use.

To enable fast searches, we have developed algorithms based on the triangle inequality and triangle-tries. The experiments in this paper provide evidence that these algorithms can be used for fast image retrieval in large databases. Our current system searches a database of 37,000 images in an average time of less than a second, depending on the query. We have provided a set of operations for combining distance measures in ways we believe to be flexible enough to allow users to fashion useful queries. Importantly, we have demonstrated that the triangle-inequality algorithms can be used with these combinations of distance measures. Thus the flexibility of our system was not obtained by sacrificing performance.

There are a number of open problems in the various data structures and algorithms we described. We have already mentioned some of them, like key selection, number of keys, trie depth, and bin size. More generally, the statistical behavior of distance measures over different sets of images influences the behavior of all the algorithms and thus needs to be explored.

Facilitating user understanding of a system such as ours is an open problem. We believe that user understanding will be an issue for any system that offers a reasonable degree of flexibility. Fundamentally, images are hard to describe. This is the flip-side of the phrase, "an image is worth a thousand words." And to whatever extent images are hard to describe, the relationship between two images can be even harder to describe. Yet it is precisely this relationship that the user must define when querying an image database based on content. There is a lot of room for fruitful research in the area of interfaces between users and content-based retrieval systems.

REFERENCES

1. J. R. Bach, S. Paul, and R. Jain, A visual information management system for the interactive retrieval of faces, *IEEE Trans. Knowledge Data Eng.* **5**(4), 1993, 619–628.
2. R. Baeza-Yates, W. Cunto, U. Manber, and S. Wu, Proximity matching using fixed-queries trees, in *Combinatorial Pattern Matching*, pp. 198–212, Springer-Verlag, New York/Berlin, 1994.
3. J. Barros, J. French, W. Martin, P. Kelley, and M. Cannon, Using the triangle inequality to reduce the number of comparisons required for similarity-based retrieval, in *IS&T/SPIE—Storage and Retrieval for Still Image and Video Databases, Jan. 1996*, Vol. IV.
4. A. Berman, *A New Data Structure for Fast Approximate Matching*, Technical Report 1994-03-02, Dept. of Computer Science, University of Washington, 1994.
5. A. Berman and L. G. Shapiro, A flexible image database system for content-based retrieval, in *17th International Conference on Pattern Recognition*, 1998.
6. A. P. Berman and L. G. Shapiro, Efficient image retrieval with multiple distance measures, in *Proceedings of the SPIE Conference on Storage and Retrieval for Image and Video Databases, February 1997*.
7. A. Del Bimbo, M. Campanai, and P. Nesi, 3d visual query language for image databases, *J. Visual Lang. Comput.* **3**, 1992.
8. A. Del Bimbo, P. Pala, and S. Santini, Visual image retrieval by elastic deformation of object sketches, in *IEEE Symposium on Visual Languages, 1994*, pp. 216–223.
9. A. Del Bimbo, E. Vicario, and D. Zingoni, Sequence retrieval by contents through spatio temporal indexing. In *IEEE Symposium on Visual Languages, 1993*, pp. 88–92.
10. W. A. Burkhard and R. M. Keller, Some approaches to best-match file searching, *Comm. ACM* **16**(4), 1973, 230–236.
11. A. Califano and R. Mohan, Multidimensional indexing for recognizing visual shapes, *IEEE Trans. Pattern Anal. Mach. Intell.* **4**, 1994, 373–392.
12. C. Carson, S. Belongie, H. Greenspan, and J. Malik, Region-based image querying, in *IEEE Workshop on Content-Based Access of Image and Video Libraries, 1997*, pp. 42–51.
13. S. K. Chang, J. Reuss, and B. H. McCormick, Design considerations of a pictorial database system, *Int. J. Policy Anal. Inf. Systems* **1**(2), 1978, 49–70.

14. S. K. Chang, Q. Y. Shi, and C. W. Yan, Iconic indexing by 2d strings, *IEEE Trans. Pattern Anal. Mach. Intell.* **9**, 1987, 413–427.
15. R. Fagin, Fuzzy queries in multimedia database systems, in *Proceedings of the 17th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 1998*, pp. 1–10.
16. M. M. Fleck, D. A. Forsyth, and C. Pregler, Finding naked people, in *Proceedings of the European Conference on Computer Vision*, pp. 593–602, Springer-Verlag, 1996.
17. M. Flickner, H. Sawhney, W. Niblack, J. Ashley, Qian-Huang, B. Dom, M. Gorkani, J. Hafner, D. Lee, D. Petkovic, D. Steele, and P. Yanker, Query by image and video content: The QBIC system, *Computer* **28**(9), 1995, 23–32.
18. D. A. Forsyth, J. Malik, M. M. Fleck, H. Greenspan, T. Leung, S. Belongie, C. Carson, and C. Bregler, Finding pictures of objects in large collections of images, in *Proceedings of the 2nd International Workshop on Object Representation in Computer Vision, April 1996*.
19. W. I. Grosky and R. Mehrotra, Index-based object recognition in pictorial data management, *Comput. Vision Graphics and Image Processing*, **52**, 1990, 416–436.
20. A. Gupta, *Visual Information Retrieval: A Virage Perspective*, Technical report, <http://www.virage.com/literature/wpaper.html>, 1995.
21. A. Gupta, *Visual Information Retrieval: A Virage Perspective*, Technical report, Virage, Inc. <http://www.virage.com/wpaper>, 1995–1997.
22. N. Haering, Z. Myles, and N. de Vitoria Lobo, Locating deciduous trees, in *Content-Based Access of Image and Video Libraries, June 1997*.
23. K. Hirata and T. Kato, Query by visual example, *Advances in Database Technology, 1992*, pp. 56–71.
24. C. E. Jacobs, A. Finkelstein, and D. H. Salesin, Fast multiresolution image querying, in *Computer Graphics Proceedings, Annual Conference Series, 1995*.
25. A. K. Jain and A. Vailaya, Image retrieval using color and shape, *Pattern Recognition* **29**, 1996, 1233–1244.
26. R. Jain, S. N. J. Murthy, and P. L. J. Chen, Similarity measures for image databases, in *Proceedings of SPIE Storage and Retrieval for Image Databases III, 1995*, pp. 58–65.
27. T. Kato, T. Kurita, N. Otsu, and K. Hirata, A sketch retrieval method for full color image database, in *11th International Conference on Pattern Recognition, 1992*, pp. 530–533.
28. P. M. Kelly and T. M. Cannon, Candid: Comparison algorithm for navigating digital image databases, in *Proceedings of the Seventh International Working Conference on Scientific and Statistical Database Management, 1994*, pp. 252–258.
29. P. M. Kelly and T. M. Cannon, Query by image example: The CANDID approach, in *SPIE Storage and Retrieval for Image and Video Databases III, 1995*, Vol. 2420, pp. 238–248.
30. W. Y. Ma and B. S. Manjunath, Netra: A toolbox for navigating large image databases, in *IEEE International Conference on Image Processing, 1997*.
31. R. Mehrotra, F. K. Kung, and W. Grosky, Industrial part recognition using a component-index, *Image and Vision Computing* **3**, 1990, 225–231.
32. T. P. Minka and R. W. Picard, *Interactive Learning Using a "Society of Models."* Technical Report 349, MIT Media Laboratory Perceptual Computing Section, 1995.
33. W. Niblack, The QBIC project: querying images by content using color, texture, and shape, Technical Report RJ 9203, IBM Research, 1993.
34. A. Pentland, R. W. Picard, and S. Sclaroff, *Photobook: Tools for Content-Based Manipulation of Image Databases*, Technical report, MIT Media Lab, 1993.
35. R. W. Picard and T. P. Minka, Vision texture for annotation, *Journal of Multimedia Systems* **3**, 1995, 3–14.
36. R. W. Picard and F. Liu, A new word ordering for image similarity, in *Proc. Conf. on Acoust., Speech, and Signal Proc., 1994*.
37. S. Sclaroff and A. Pentland, Object recognition and categorization using modal matching, in *Proc. of 2nd CAD-Based Vision Workshop, 1994*, pp. 258–265.
38. S. Sclaroff, L. Taycher, and M. La Cascia, Imagerover: A content-based image browser for the world wide web, in *Proc. IEEE Workshop on Content-based Access of Image and Video Libraries, June 1997*.
39. L. G. Shapiro and M. S. Costa, Scene analysis using appearance-based models and relational indexing, in *IEEE Symposium on Computer Vision, 1995*, pp. 103–108.
40. L. G. Shapiro and R. Haralick, Structural descriptions and inexact matching, *IEEE Trans. Pattern Anal. Mach. Intell.* **3**, 1981, 504–519.
41. M. Swain and D. Ballard, Color indexing, *International Journal of Computer Vision* **7**, 1991, 11–23.
42. M. Tuceryan and A. K. Jain, Texture analysis, *Handbook of Pattern Recognition and Computer Vision*, pp. 235–276, Academic Press, San Diego, CA, 1994.
43. M. Turk and A. Pentland, Eigenfaces for recognition, *Journal of Cognitive Neuroscience*, 1991.
44. J. Uhlmann, Satisfying general proximity/similarity queries with metric trees. *Information Processing Letters* **40**, 1991, 175–179.
45. L. Wang and D. C. He, Texture classification using texture spectrum, *Pattern Recognit. Lett.* **13**, 1990, 905–910.