Image Retrieval in Multipoint Queries

Khanh Vu, Hao Cheng, Kien A. Hua

Department of Computer Science, University of Central Florida, Orlando, FL 32816-2362

Received 8 November 2007; revised 2 June 2008; accepted 12 June 2008

ABSTRACT: Traditional content-based image retrieval (CBIR) systems find relevant images close to an example image. This single-point model has been shown inadequate for complex queries built on high-level concepts. Recent CBIR systems allow users to use multiple examples to compose their queries. The multipoint model provides extended flexibility in identifying relevant sets of arbitrary shape that the previous approach is unable to formulate. However, the continuing use of conventional measures (e.g., L_p norms) to evaluate these queries undermines the advantages of the new system. From the results of recent studies, we show that two important inferences can be made. Specifically, they are the continuity of image representation and the nonhomogeneity of the feature space. These characteristics enable the precise identification of points that satisfy the constraints established in multipoint queries and for any orthogonal feature representations. Generally, the sets are convex hulls and can be described by a linear system of equations with constraints. We discuss how to solve the system and propose an indexing procedure to efficiently determine the exact sets. We evaluated the performance of the proposed technique against state-of-the-art methods on large sets of images. The results indicate that the new measure captures semantic image classes very well, and the superiority of our approach over the recent techniques is evident in simulated and realistic environments. © 2008 Wiley Periodicals, Inc. Int J Imaging Syst Technol, 18, 170–181, 2008; Published online in Wiley InterScience (www.interscience.wiley.com). DOI 10.1002/ima.20162

Key words: multimedia information retrieval

I. INTRODUCTION

The need to effectively handle growing volumes of multimedia data has attracted much research attention addressing a variety of issues in image indexing and retrieval (Chang and Kao, 1993; Niblack, 1993; Jacobs et al., 1995; Manjunath and Ma, 1996; Smith and Chang, 1996; Carson et al., 1999; Natsev, 1999; Huitao and Alexandros, 2000). In many current content-based image retrieval systems, relevant images are determined based on their similarity to a single image. However, it has been shown that in many applications such a single-point query model lacks the flexibility to capture queries' intent (Bjoerge and Chang, 2004). Recent query models (Porkaew et al., 1999; Kim and Chung, 2003; Chakrabarti, 2004; Ortega-Binderberger and Mehrotra, 2004) allow (the high-level concepts of) queries to be expressed using multiple reference images, enabling the specification of desired constraints on the sets to be retrieved. The extended model provides enormous flexibility in query formulation, able to overcome serious limitations of traditional similarity measures (Ortega-Binderberger and Mehrotra, 2004). These techniques can be seen as a result of the following findings.

Extensive research has helped identify a set of characteristic image features that allow robust comparison of images. For instance, color features are now represented using uniform color systems (e.g., HSV, HVC) in which the dissimilarity of two colors is simply the distance between them (Hiyahara and Yoshida, 1988); texture features are captured in DFT coefficients for robust texture evaluations; shape is represented by edge-based structures (Zhou and Huang, 2001) or skeletal graphs (Hiransakolwong et al., 2004) that make possible translation, rotation- and scaling-invariant matching. These sets of features enable images of similar characteristics to be represented continuously in low-level feature space. Perhaps, the most prominent illustration of the continuity property of image representation is the a priori relevancy of the centroid: for a given multipoint query, its centroid can be shown to be a virtual, ideal query, one with all the specified characteristics. This is the underlying basis of picking the centroid to represent a set of query points in many multipoint-to-single-point approaches (Porkaew et al., 1999; Hiransakolwong et al., 2004). The rationale has been extended to weighted centroids, which are computed based on user-defined relevance rankings of query images (Ortega-Binderberger and Mehrotra, 2004). It will be shown later that weighted centroids always satisfy the feature constraints set by query points, independent of orthogonal transformations of features.

On the other hand, there appears to be no direct correspondence between similarity of high-level concepts and the proximity of their representing points in low-level feature space, which is widely known as semantic gaps. In fact, examples can be found that images belonging to the same semantic class could be farther apart than irrelevant ones as quantified by well-known measures, such as normed distances. This nonhomogeneity property implies that these distances not be used as a (dis)similarity measure in applications such as image retrieval.

Example: Compute L_p distances between points representing Images in Figure 1: P, Q1, Q2, and Q, the centroid of Q1 and Q2. (Their features are described in detail in Section V.)

Correspondence to: Dr. Khanh Vu; e-mail: khanh-vu@acm.org

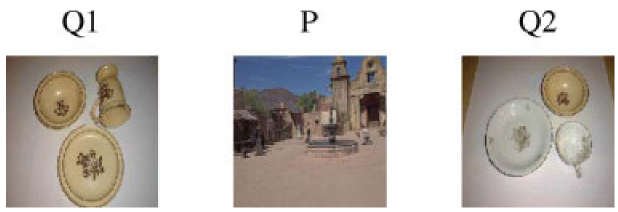

Figure 1. Some example images. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Answer: See Figure 2. Observe that L_p (Q1,Q2) are greater than both respective L_p (Q1,P) and L_p (P,Q2). The same is true for L_p (Q,Q1) or L_p (Q,Q2) versus L_p (P,Q).

In the above example, it would not be possible to separate false hit P from semantic class (Q_1,Q_2) based on the normed distances centered at Q1, Q2, or Q. This is because these measures disregard the nonhomogeneous nature of the feature space and reach into all dimensions in search of ''similar'' images. In multipoint queries, semantic classes of interest are defined by users specifying anchor points, which can be seen as users explicitly establishing the constraints for relevant points. Figure 3 illustrates the relevant points of query (Q1,Q2,Q3), i.e. the shaded area, that always satisfy the constraints of the query. They are weighted centroids of the query, while P and other gray points are not. When relevant sets are defined with a search radius, it has the effect of expanding the query set [e.g., (Q1,Q2,Q3,Q4,Q5)] and thereby the centroid set. Note that the L_2 measure from the centroid (i.e., the big circle) would encompass many false hits.

The nonhomogeneity and continuity characteristics of feature space have been observed for several semantic classes, such as linearly enhanced images (Vu et al., 2005) and ''similar'' time series (Goldin and Kanellakis, 1995). It was shown in (Vu et al., 2005) that such semantic classes are strictly confined in 2-D planes passing through the space's diagonal. Class members could be as far apart as spanning the largest length of the dataspace. For general image classes, we postulate that they are confined in arbitrary, bounded, low-dimensional subspaces within the feature space, and accurately isolating them is essential to improving retrieval effectiveness. This postulate might not be valid when images are not represented adequately or images of the same class are very different visually (e.g., a single red rose and a bunch of purple violets in the flower category). In the latter case, retrieval of the whole category must be executed with several subqueries, one for each cluster (Hua et al., 2006). We will address this issue in our future work.

In this article, we present a CBIR approach that observes the above characteristics of low-level feature space. To achieve this, we must first make sure that images are adequately represented. Fortunately, this issue can be addressed by describing images with well-known characteristic features suggested in recent works (see Section V). With the relevancy of weighted centroids assumed, we devise an algorithm to retrieve relevant sets that closely conform

L_p	(Q1,P)	(P,Q2)	(Q1,Q2)	(P,Q)	(Q,Q1) or (Q,Q2)
L_1	2.235	5.353	6.235	2.711	3.117
L_2	0.478	1.198	1.374	0.600	0.687
L_∞	0.209	0.457	0.460	0.231	0.231

Figure 2. Normed distances between the images.

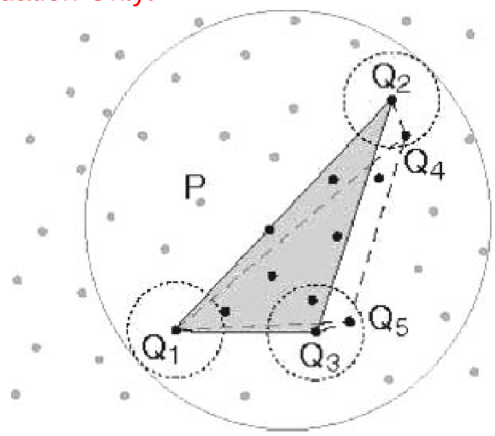

Figure 3. Illustration of the continuity and non-homogeneity properties of low-level feature space.

the centroid set of the query, that is, to find weighted centroids of the query expanded based on user-defined search thresholds. We discuss in detail how to represent and determine this set precisely. As indexing is crucial for efficient searches over large databases, we propose an approximation method to reduce the overall time complexity of finding the exact sets. The results of our experiments and system implementation indicate that the centroid sets match image semantic classes extremely well, and our procedure is very efficient. The remaining article is organized as follows. Section II reviews related works and summarizes their common shortcomings. We describe our approach in detail in Section III and an indexing scheme in Section IV. Section V discusses the results of our performance study. Concluding remarks are given in Section VI.

II. RELATED WORK

Various methods have been proposed to support multipoint queries, many being designed to facilitate image retrieval with relevance feedback. During a retrieval session, users identify relevant/irrelevant images from a set presented by the system (e.g., the result of the preceding round) and form a query for the next round. A chief goal of these systems is to minimize false hits, those that are beyond the intent of the query. To accomplish this, the focus is typically on two critical areas: extracting salient image features and selecting appropriate similarity metrics.

Salient image features and their significance to a particular query can be inferred from user feedback (Maybank, 2007; Tao and LI, 2008). It has been observed that, in general, relevant samples are confined in a small neighborhood, while irrelevant ones scatter across multiple regions in the feature space. In (Maybank, 2007), negative samples are clustered into several groups, and a marginal convex machine (MCM) sub classifier is applied to distinguish negative groups and the positive one. The sub classifiers are combined to form a biased MCM. In (Tao and Li, 2008), an orthogonal complement component analysis is applied to determine the inherent structure of the positive samples, and all data points are then mapped into the resulting new space. In Tao et al. (2006a), the direct and biased discriminant analyses are integrated to derive the salient projections of the data points. Tao et al. (2006b) argues that SVM, performing poorly with a small training set, can still be used as a weak classifier and the classifier committee learning can boost the performance using a process in which the bagging and random subspace are utilized to aggregate the results of multiple weak classifiers.

With images properly represented in the feature space, retrieval systems apply appropriate metrics to measure similarity. We review several techniques below, which differ in the way they use similarity measures to estimate retrieval sets.

Many multipoint-to-single-point techniques have been proposed including Image Grouper (Nakazato et al., 2003) (using weighted L_2 distance from the centroid to images under consideration) and (Koompairojn et al., 2004) (using weighted L_1 distance as the similarity measure). MARS (Ortega-Binderberger and Mehrotra, 2004) selects weighted centroids taking into consideration different relevance rankings of query points. Each dimension's weight is inversely proportional to the standard deviation of the relevant images' feature values in that dimension. The rationale is that a small variation among the values is more likely to express restrictions on the feature, and thereby should carry a higher weight. On the other hand, a large variation indicates this dimension is not significant in the query, thus should assume a low weight.

In Query Expansion (Porkaew et al., 1999; Chakrabarti, 2004); Qcluster (Kim and Chung, 2003) and Query Decomposition (Hua et al., 2006) multiple query points are used to define the ideal space most likely to contain relevant results. Query Expansion clusters query points and chooses their centroids as representatives. The distance between a point and a query is defined as the weighted sum of individual L_2 distances from the representatives. As the weights are proportional to the number of relevant objects in the clusters, Query Expansion treats local clusters differently, as opposed to the equal treatment in other multipoint-to-single-point techniques.

In some queries, clusters are too far apart for a unified, all-encompassing contour to be effective; separate contours can yield more selective retrieval. This observation motivated Qcluster (Kim and Chung, 2003) to employ an adaptive classification and cluster-merging method to determine the optimal contour shapes for complex queries. Qcluster supports disjunctive queries, where similarity to any of the query points is considered as good. To handle disjunctive queries both in vector space and in arbitrary metric space, a technique was proposed in FALCON (Wu et al., 2000). It uses an aggregate distance function to estimate the (dis)similarity of an image to a set of desirable images. Query Decomposition (Hua et al., 2006) automatically decomposes a complex query into localized subqueries, which proves to be highly effective for images with similar semantics but in very different appearance (e.g., the front view and the side view of a car).

A common trait shared by these techniques is the use of (weighted) normed distances in identifying retrieval sets or forming image clusters. As observed in Section I, these measures construct retrieval volumes spanning all dimensions, thereby vastly overestimating the relevant sets (how these techniques cover relevant sets will be depicted in Section IV). A large number of false hits complicates user selection of relevant images, increases false dismissals, and as a result, causes slow convergence in multi-round image retrieval.

In our approach, a system of equations is used to describe constraints established by query points. Since, as will be shown shortly, only weighted centroids satisfy the system, retrieval sets can be precisely isolated. Our experiments show that the new measure is capable of eliminating almost all false hits when images are represented properly (with a moderate number of low-level features). To support multipoint search over very large databases, we also consider an indexing scheme to speed up retrieval, which proves to be a serious challenge for recent multipoint CBIR techniques.

This work can be seen as a generalization of our previous work (Vu et al., 2005; 2006b). Our previous measure (Vu et al., 2005) is applied to retrieve a special semantic class, linearly dependent images, which are confined strictly in 2-D planes passing through the diagonal. Hence, it cannot be used for general multipoint queries, which could occupy arbitrary subspace, i.e., arbitrary shape, size, and orientation. In Vu et al., 2006b, we did not evaluate the centroid sets or consider an indexing support to handle large databases. In this work, we set out the following contributions:

- We present a representation scheme to capture the constraints of multipoint queries. This representation, by means of a linear system of equations, precisely describes the shape of the retrieval set, invariant to any orthogonal transformation of the features.
- We work out an efficient procedure to solve the system. Depending on the system's determinism, our method is able to determine whether there is no solution (point is irrelevant), a unique solution, or infinite solutions (point is relevant).
- To support fast multipoint search for very large databases, an approximation scheme is proposed, which can be implemented on top of popular index structures such as R*-tree (Beckman et al., 1990) to significantly reduce query processing time.

III. MULTIPOINT QUERY RETRIEVAL

In this article, relevant points are defined as those that satisfy the feature constraints set by query points. Specifically, if feature i is identified as acceptable at two values, then any point with feature i within that range is so considered. In general, if n-dimensional points $L(l_1,l_2,\ldots,l_n)$ and $H(h_1,h_2,\ldots,h_n)$ are relevant, so are points $P(p_1,p_2,\ldots,p_n)$ where $p_i \in [l_i,h_i]$ or $p_i \in [h_i,l_i]$, $1 \le i \le n$. This condition needs to hold true for arbitrary orthogonal transformation of the features. The requirement has the implication that the relevant set is always the same, independent of the representation being used. More importantly, as shown in our implementation, such relevant sets match the intended semantic classes of multipoint queries extremely well.

In this article, we show that for a given set of query points, their weighted centroids are always within the bounds established by the query points, regardless of their representations. (Much of the results are applicable to arbitrary finite metrics. However, since any finite metric can be embedded in L_2 norm (Goldin and Kanellakis, 1995) we will skip the discussion of it.) Let S denote the set of n-dimensional points, each representing an image with n features; $\boldsymbol{Q} = \{Q_1,Q_2,\ldots,Q_k\} \subseteq S$, a set of k distinct query points; w_i non-negative weight associated with Q_i where $\int_{i=1}^{k} w_i = 1$; $[l_i,h_i]$ the relevant range established by $\boldsymbol{Q}$ in the dimension i, where $l_i = \min\{q_{j,i}|Q_j \in \boldsymbol{Q}, 1 \le j \le k\}$ and $h_i = \max\{q_{j,i}|Q_j \in \boldsymbol{Q}, 1 \le j \le k\}$.

A. The Relevant Set

Definition 1. (Relevance) Point $P(p_1,p_2,\ldots,p_n)$ is said to be relevant to $\boldsymbol{Q}$ if for any orthogonal representation of S, then $p_i \in [l_i,h_i]$.

Definition 2. (Weighted Centroid) $P(p_1,p_2,\ldots,p_n)$ is a weighted centroid of $\boldsymbol{Q}$ if, for$1 \le i \le n$,

$$p_i = w_1 \cdot q_{1,i} + w_2 \cdot q_{2,i} + \cdots + w_k \cdot q_{k,i} \tag{1}$$

When $w_i = 1/k$, P is the centroid of $\boldsymbol{Q}$.

Lemma 1. P is relevant if and only if it is a weighted centroid of $\boldsymbol{Q}$.

Proof. We first show that if P is a weighted centroid of $\boldsymbol{Q}$, then $p_i \in [l_i, h_i]$ for arbitrary orthogonal representation of the features. We prove $p_i \le h_i$; the lower bound can be similarly obtained. We have:

$$\begin{aligned} p_i &= w_1 \cdot q_{1,i} + w_2 \cdot q_{2,i} + \cdots + w_k \cdot q_{k,i} \\ &\le w_1 \cdot h_i + w_2 \cdot h_i + \cdots + w_k \cdot h_i \\ &\le h_i \sum_{j=1}^{k} w_j \\ &\le h_i. \end{aligned}$$

On the other hand, if P is relevant, then it is linearly dependent with constraints on $\boldsymbol{Q}$'s points (Eq. 2). In other words, it can be formulated as a weighted centroid of $\boldsymbol{Q}$.

Note the absence of any assumptions on the features p_i and $q_{j,i}$, and thereby l_i and h_i, implies that Lemma 1 is true for arbitrary feature presentation of S. ■

We now define the retrieval problem in multipoint queries:

Definition 3. (Multipoint Query Retrieval) The no-false-dismissal retrieval for a given multipoint query $\boldsymbol{Q}$ is to retrieve all weighted centroids of $\boldsymbol{Q}$.

In the remaining of this article, we discuss how to guarantee the retrieval of this centroid set, multipoint retrieval with a search radius (i.e., query expansion), and an indexing scheme to speed up query processing.

B. Representation of the Relevant Set

Fully enumerated, Eq. 1 produces n equations with k unknowns that completely describe the centroid set of the query. In other words, P is a weighted centroid of $\boldsymbol{Q}$ if it satisfies the linear system of equations (in matrix form) with constraints:

$$\begin{cases} \boldsymbol{Q}\vec{w} &= \vec{p} \quad (a) \\ w_i &\ge 0 \quad (b) \\ \sum_{i=1}^{k} w_i &= 1 \quad (c) \end{cases} \tag{2}$$

where

$$\vec{p} = \begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_n \end{bmatrix}, \quad \boldsymbol{Q} = \begin{bmatrix} q_{11} & q_{12} & \cdots & q_{1k} \\ q_{21} & q_{22} & \cdots & q_{2k} \\ & \ddots & \ddots & \vdots \\ q_{n1} & q_{n2} & \cdots & q_{nk} \end{bmatrix} \text{ and } \quad \vec{w} = \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_k \end{bmatrix}.$$

Observe that some query points themselves could be weighted centroids of the remaining. The minimal set of independent query points will be called the principal set.

A general procedure to retrieve the relevant set is presented in Algorithm 1. *searchSet*($\boldsymbol{Q}$,S,r) first expands the query set based on user-defined search radius r (see Section III.E.), determines the principal set of the query points (see Section III.D.), performs a quick-and-dirty filter of the dataspace (see Section IV), solves the system 2(a) for w_i for each of the candidate points, and checks the weight constraints 2(b) and (c) to determine its relevancy (see Section III.C).

Algorithm 1 *searchSet*($\mathcal{Q}, \mathcal{S}, r$)

```
Input:
  set of query images: Q
  set of database images: S
  user-defined search radius: r
Output:
  set of relevant images: C
begin
  Set Q ← expandQuery(Q, r).
  Set Q ← principalSet(Q).
  Set C ← ∅.
  Set S ← filterDB(S, r).
  for all P ∈ S do
    if P ∈ Q then
      Set C ← C ∪ {P}.
    else
      Solve p⃗ = Qw⃗ for w⃗.
      if (min(w⃗) ≥ 0) & (∑_{i=0}^{k} w_i = 1) then
        Set C ← C ∪ {P}.
      else
        Discard P.
      end if
    end if
  end for
  return C.
end
```

If $\boldsymbol{Q}$ is a square and singular matrix, there exists a unique solution to the system, nd therefore only one check is required. In general, however, $n \neq k$ and thus the system can be over-determined or under-determined. This is likely because while n is fixed (in our experiments, $n = 37$), k varies greatly depending on the outcome of *expandQuery*($\boldsymbol{Q}$,r). These cases are explained below.

- $k > n$. Initial k is typically small (users normally select around three to five query images). However, query expansion could produce a much higher k, up to tens or hundreds. The system then has more unknowns than equations; it is under-determined. An underdetermined system either has no solution or infinitely many solutions.
- $k < n$. This can happen if query expansion is inadequate (e.g., the search radius is too small or search space is sparse). The system is then overdetermined; it has more equations than variables. An overdetermined system can have either no solution, exactly one solution, or infinitely many solutions.
- $k = n$. Generally, there is a unique solution to the system.

As a result, we have to expect many possible outcomes solving System 2 and may attempt to circumvent solving an over- or under-determined system by reshaping $\boldsymbol{Q}$ into a square matrix, e.g., by forcing it into a square and filling in the expanded matrix, ignoring some dimensions (reducing n), dropping some query points (reducing k), or expanding the query set (increasing k). The pitfall is, however, we may be altering the semantics of the initial query, and unwittingly enclosing irrelevant points. Therefore, it is desirable to solve the original system. In the following, we discuss how to solve such a system. In Section IV, we propose an indexing scheme for the entire search.

C. Solving the Linear System

Recall that System 2 has a set of n linear equations in k unknowns, and our task is to test, given $\vec{p}$, whether there exists $\vec{w}$ which satisfies 2(a) and constraints (b) and (c). In general, $\vec{p} \neq \vec{0}$ and the non-homogeneous system $\boldsymbol{Q}\vec{w} = \vec{p}$ has no solution, one unique solution, or an infinite number of solutions. Let the intrinsic dimensionality of a set of points be its rank. For example, the rank of a triangle is 2, and that of a nonplanar 4-point pyramid is 3. We investigate the following possible cases.

- Case 1: An additional intrinsic dimension is required to represent $\boldsymbol{Q}$ *augmented with* $\vec{p}$, that is: $rank([\boldsymbol{Q}\vec{p}]) > rank(\boldsymbol{Q})$. Then $\vec{p}$ is outside the space represented by $\boldsymbol{Q}$, and thus $\boldsymbol{Q}\vec{w} = \vec{p}$ has no solution. Consequently, if the above condition is true for the problem input, no further tests on the other constraints are needed.
- Case 2: Matrix $\boldsymbol{Q}$ and the augmented matrix $[\boldsymbol{Q}\vec{p}]$ have the same rank. Then $\boldsymbol{Q}\vec{w} = \vec{p}$ has at least one solution. There are two subcases:

  i. If $rank(\boldsymbol{Q})=k$, then $\boldsymbol{Q}\vec{w} = \vec{p}$ has one unique solution $\vec{w}$, and the homogeneous system $\boldsymbol{Q}\vec{w} = \vec{0}$ does not have non-zero solutions. We can directly compute its unique solution by various methods, e.g. Gaussian Elimination. With the obtained solution $\vec{w}$, we can test whether or not both $w_i \geq 0$ and $\sum w_i = 1$ are satisfied.

  ii. If $rank(\boldsymbol{Q})<k$, then the homogeneous system $\boldsymbol{Q}\vec{w} = \vec{0}$ has solutions and the corresponding non-homogeneous system has an infinite number of solutions. We will study this case in detail below.

When $rank(\boldsymbol{Q})<k$, the null space of matrix $\boldsymbol{Q}$ has $s=(k\text{-}rank(\boldsymbol{Q}))$ free orthogonal variables. This set of variables span the entire null space with basis vectors $N_1, N_2, \ldots$, and N_s, where

$$\vec{N}_1 = \begin{bmatrix} N_{11} \\ N_{21} \\ \vdots \\ N_{n1} \end{bmatrix}, \quad \vec{N}_2 = \begin{bmatrix} N_{12} \\ N_{22} \\ \vdots \\ N_{n2} \end{bmatrix}, \quad \text{and} \quad \vec{N}_s = \begin{bmatrix} N_{1s} \\ N_{2s} \\ \vdots \\ N_{ns} \end{bmatrix}.$$

This implies that for any $t_1, \ldots, t_s \epsilon (-\infty, \infty), t_1 \overrightarrow{N_1} + t_2 \overrightarrow{N_2} + \cdots + t_s \overrightarrow{N_s}$ is a solution to the homogeneous system $\boldsymbol{Q}\vec{w} = \vec{o}$. Let w' be the particular solution to $\boldsymbol{Q}\vec{w} = \vec{p}$. Then

$$\vec{w} + t_1\vec{N}_1 + t_2\vec{N}_2 + \cdots + t_s\vec{N}_s$$

is the general form of the solutions to $\mathbf{Q}\vec{w} = \vec{p}$, where

$$\begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_n \end{bmatrix} = \begin{bmatrix} w'_1 \\ w'_2 \\ \vdots \\ w'_n \end{bmatrix} + t1 \begin{bmatrix} N_{11} \\ N_{21} \\ \vdots \\ N_{n1} \end{bmatrix} + \cdots + t_s \begin{bmatrix} N_{1s} \\ N_{2s} \\ \vdots \\ N_{ns} \end{bmatrix}.$$

We still have to test whether there exist $t_1, t_2, \ldots$, and t_s such that the resulting $\vec{w}$ satisfies the constraints of our original problem. Plugging in $t_1, t_2, \ldots$, and t_s to the constraints 2(b) and (c), we have,

$$\left[\sum\nolimits_{i=1}^{n} N_{i1} \cdots \sum\nolimits_{i=1}^{n} N_{is}\right] \begin{bmatrix} t_1 \\ t_2 \\ \vdots \\ t_s \end{bmatrix} = 1 - \sum_{i=1}^{n} w'_i$$

and

$$\begin{bmatrix} -N_{11} & -N_{12} & \cdots & -N_{1s} \\ -N_{21} & -N_{22} & \cdots & -N_{2s} \\ \ddots & \ddots & \ddots & \\ -N_{n1} & -N_{n2} & & -N_{ns} \end{bmatrix} \begin{bmatrix} t_1 \\ t_2 \\ \vdots \\ t_s \end{bmatrix} \leq \begin{bmatrix} w'_1 \\ w'_2 \\ \vdots \\ w'_n \end{bmatrix}.$$

Finding the solution to the above equations means checking whether there exists a feasible region for the given $t_1, t_2, \ldots$, and t_s. That is, we need to test whether the feasible region is empty or not by setting the objective function to be zero, a task that can be done using the linear programming.

Observe that the transformed problem above has much fewer unknowns (i.e., typically $s \ll k$), and only for a few cases, thus the cost of computing those is small on average. In our experiments, the above procedure is quite efficient compared to using the linear programming to solve the original problem directly.

D. The Principal Set

Determining the principal set of $\boldsymbol{Q}$ amounts to finding the minimal system of Eq. 2(a). The principal points form the convex hull of $\boldsymbol{Q}$, which can be found using the Gift-Wrapping algorithm (Skiena, 1997) Qhull (Barber et al., 1996) or others. The gift-wrapping is the basic algorithm for constructing higher-dimensional convex hulls, which can take $O(k^{n/2+1})$. Algorithm 2 computes the principal set based on the gift-wrapping algorithm for $n = 2$ (available at http://www.cse.unsw.edu.au/lambert/java/3d/giftwrap.html). Convex hulls of higher dimensions can be found with its generalized version.

It is easy to see that computing the principal set using a convex-full algorithm is costly. An alternative is to use the procedure in the preceding section to determine query points that are not part of the convex hull. Specifically, let's start with any two farthest points in $\boldsymbol{Q}$, which make up the initial $\boldsymbol{Q}^{\circ}$. Consider $\vec{p}$, the farthest point of the remaining points of $\boldsymbol{Q}$ from $\boldsymbol{Q}^{\circ}$. Point $\vec{p}$ is not a convex-hull point if and only if the following system is satisfied:

$$\begin{cases} \boldsymbol{Q}^{\circ}\vec{w} & = & \vec{p} \\ w_i & \geq & 0 \\ \sum_{i=1}^{k} w_i & = & 1 \end{cases} \tag{3}$$

Algorithm 2 *principalSet*($\mathcal{Q}$)

```
Input:
  set of query images: Q
Output:
  principal set of query images: Q°
begin
  Find an extreme point Q_A ∈ Q (one with a minimun
    it to Q°.
  repeat
    Find Q_B ∈ Q where all points lie to the left of Q_A
    if Q_B ∉ Q° then
      Add Q_B to Q° and rename it to Q_A.
    else
      Exit loop.
    end if
  until Done
  return Q°.
end
```

Otherwise, add $\overrightarrow{p}$ to $\boldsymbol{Q}^\circ$ and repeat the calculation until $\boldsymbol{Q}$ is empty. Observe that our procedure in Section III.D is capable of evaluating System 2 even if matrix $\boldsymbol{Q}$ is not minimal (i.e., skipping line 2 of Algorithm 1). However, since the principal set is typically much smaller than $\boldsymbol{Q}$, it is faster to compute the system with $\boldsymbol{Q}^\circ$ for a large number of candidate points.

E. Query Expansion

Query sets can be enlarged by including points within radius r_i from Q_i or its nearest neighbors. In general, r_i can be calculated taking into consideration the rankings, the directions and how aggressive the expansion should be (Rocchio, 1971). Given r_i, our *expandQuery*($\boldsymbol{Q}$,r) expands $\boldsymbol{Q}$ in all directions but farther in the directions with larger r_i. The expanded query now spans n dimensions but its shape still conforms the shape of the original query. The expanded set is treated as the input query to the retrieval algorithm.

IV. INDEXING AND APPROXIMATE SEARCH

For very large databases, it is crucial to index the data to support efficient search, especially when the complexity of computing relevant sets is high, such as in multipoint retrieval. A general approach to fast search is to perform it in two steps: the exact search is preceded by a quick-and-dirty filtering. According to (Faloutsos, 1994; Keogh et al., 2001), a filter must guarantee no false dismissals, is able to prune out a majority of irrelevant points efficiently, and can be implemented on top of a high-performance index structure such as the R*-tree. In the following, we briefly review recent techniques and discuss the challenges of supporting fast search. We will then show how our approach does not suffer these drawbacks.

A. Existing Schemes for Multipoint Queries

As observed in Section I, existing techniques can be regarded as a distance-based retrieval approach; relevant points are determined based on some normed distance measured from the query points or from weighted centroids of the query points. This centroid-based measure reflects the dominant thinking that all query points should contribute to the defining of relevant points. In other words, in these techniques the concentration of query points should guide the search for relevant sets. Below, we will determine the minimum bounding shapes according to these proposals that guarantee no false dismissal of the relevant sets. These will be used in our experimental study.

First, we need to determine the minimum threshold so that the retrieval shape generated by each of the techniques properly encloses $\boldsymbol{Q}$. Let C denote the centroid of $\boldsymbol{Q}$, H the retrieval set, $\min_i = \min\{q_i | 1 \le i \le k, Q \in \boldsymbol{Q}\}$ and $\max_i = \max\{q_i | 1 \le i \le k, Q \in \boldsymbol{Q}\}$ be the bounds in dimension i.

A.1. Minimum Bounding Hyperrectangle. This approximation, proposed in Koompairojn et al. (2004) is a hyper rectangle with center at C and side i inversely proportional to $(\max_i - \min_i)$ up to scale t. Let $L_1(C,Q)$ be the weighted city-block distance of Q from C:

$$\mathcal{L}_1(C,Q) = \sum_{i=1}^{n} \frac{t}{\max_i - \min_i} |c_i - q_i|.$$

The minimum threshold r must be large enough to encompass $Q_f \in \boldsymbol{Q}$, the farthest query point from C. Thus, $r = L_1(C,Q_f)$. The retrieval set H is thus:

$$H = \{P | P \in S, L_1(C,P) \le r\}.$$

A.2. Minimum Bounding Hypersphere. Proposed in (Nakazato et al., 2003; Chakrabarti, 2004), the retrieval set is points P contained in the hypersphere with center at C and radius $r = L_2(C,Q_f) = \sum_{i=1}^{n} t(c_i - q_{f,i})^2$, where $Q_f \in \boldsymbol{Q}$ is the farthest query point from C. The retrieval set H is therefore:

$$H = \{P | P \in S, \mathrm{L}_2(C,P) \le r\}.$$

A.3. Minimum Bounding Iso Hypersurface. Proposed in MARS (Porkaew et al., 1999). The aggregate distance of point P to query set $\boldsymbol{Q}$ is defined as:

$$L_{agg}(P,Q) = \sum_{Q \in \boldsymbol{Q}} L_2(P,Q).$$

To ensure that all relevant points are retrieved, a minimum aggregate radius must be established. Let $r = \max\{L_{\mathrm{agg}}(Q, \boldsymbol{Q}) | Q \in \boldsymbol{Q}\}$. Then,

$$H = \{P | P \in S, L_{\mathrm{agg}}(P, \mathbf{Q}) \le r\}.$$

We state the following lemma (proof is omitted).

Lemma 2. The set H so determined encloses all relevant points of $\boldsymbol{Q}$. In implementation, r is computed as follows:

1. Construct the adjacency matrix $k \times kM$, where $m_{i,j} = L_2(Q_i,Q_j)$.
2. Compute the sum of column i, which is $L_{\mathrm{agg}}(Q_i,\boldsymbol{Q})$.
3. Find r, the largest value among the column values.
4. Use r as the minimum threshold to filter out irrelevant points.

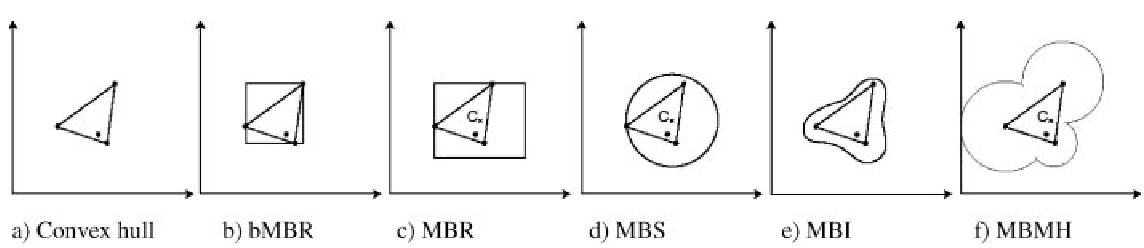

Figure 4. Convex hull of a co-planar four-point query set and approximation shapes.

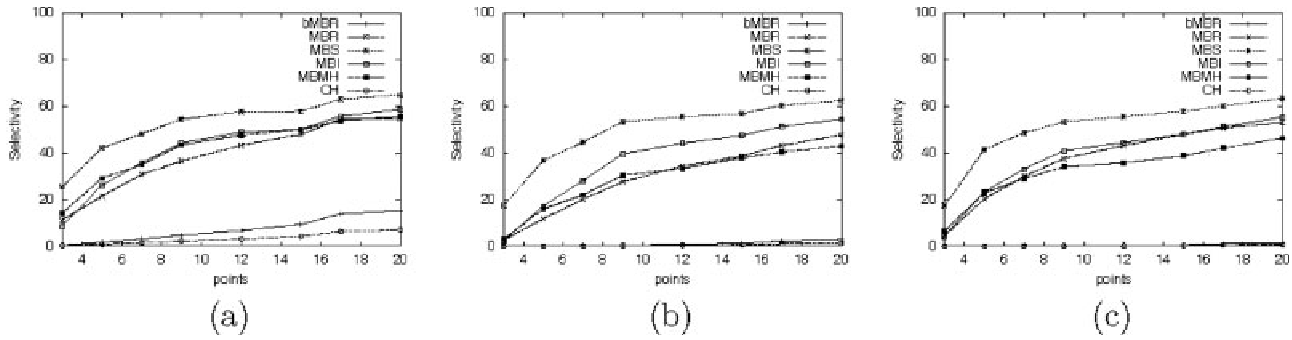

Figure 5. Selectively versus the number of query points.

A.4. Minimum Bounding Merged Hyperspheres. Minimum bounding merged hypersphere (MBMH) is a version of distance-based concave shapes (Kim and Chung, 2003) where all query points have the same relevance rank. Consequently, the disjunctive distance of point P from $\boldsymbol{Q}$ is:

$$\mathcal{L}_{\text{disj}}(p, \mathcal{Q}) = \frac{k}{\sum_{Q \in \mathcal{Q}} \frac{1}{\mathcal{L}_2(P,Q)}}.$$

To ensure that the merged shape covers C, a minimum disjunctive radius must be established. Let $r = \max\{L_{\text{disj}}(C,\boldsymbol{Q}) | Q \in \boldsymbol{Q}\}$. Then,

$$H = \{P | P \in S, L_{\text{disj}}(P, \mathbf{Q}) \leq r\}.$$

B. Fast Search in Existing Techniques

Figures 4a and 4c–4f portrait the convex hull of a co-planar 4-point query set and the shapes approximating the hull by these techniques. Several observations can be made:

- These are n-dimensional volumes.
- The selection of centroid points as queries (see Figs. 4c and 4d) and uneven distribution of query points (see Fig. 4e) have the effect of stretching the approximation shapes.
- The size of the convex hull has a significant impact on MBMH (see Fig. 4f).

As will be seen in the experimental results, the displacement of shapes approximating relevant sets causes reduced effectiveness in these methods. Regarding efficient search, we note the following drawbacks of these techniques:

1. Only the minimum bounding hyperrectangle (MBR) queries can be executed as very efficient window queries (see Fig. 4c). However, MBR is dependent on the representation of the features, thus it retrieves different results for different feature spaces.
2. Although the other shapes are invariant with respect to the features' orthogonal representation (see Figs. 4d–4f, it is difficult to support their fast search. Clearly, it is not an easy task to approximate minimum bounding iso hypersurface (MBI) and MBMH with window queries, since it requires computing the maximum reach of their volumes in each dimension. While minimum bounding hypersphere (MBS) can be approximated with a minimum bounding hypercube Paradopoulos and Manolopoulos, 1997 such an approach, however, degrades very quickly for increasing dimensionality Weber et al., 1998).
3. Transformations do not make it easier to support efficient search in those techniques. The problem compounds for increasing dimensionality, because dimensionality reduction methods such as DFT (Agrawal et al., 1993) SVD (Kanth et al., 1998), or MS (Vu et al., 2006) are rendered useless.

In fact, these techniques did not consider an indexing scheme (e.g. MBMH) or propose one for k-NN retrieval only (e.g., MBI). The obvious option left is to scan the entire database sequentially to guarantee no false dismissals.

C. Boundary-Based Minimum Bounding Hyperrectangle

The discussion in the preceding section has suggested that retrieval shapes in multipoint queries should not be dependent on the

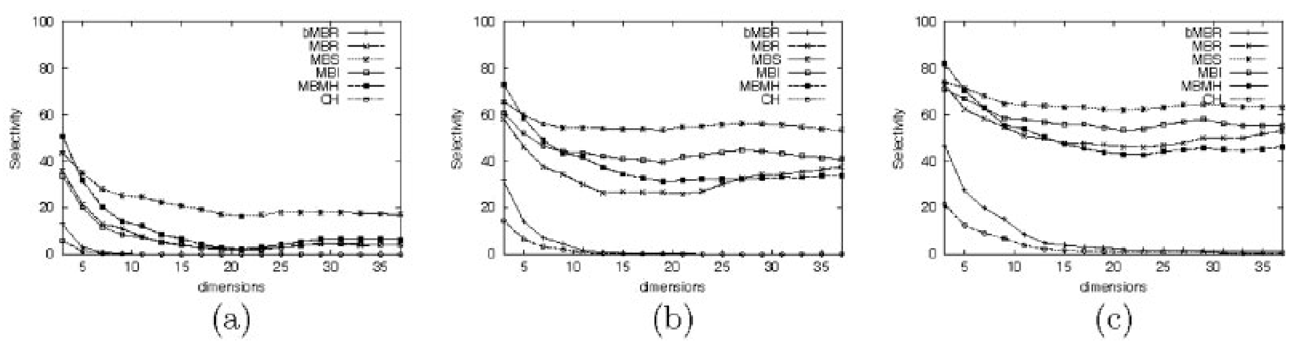

Figure 6. Selectively versus the number of dimensions.

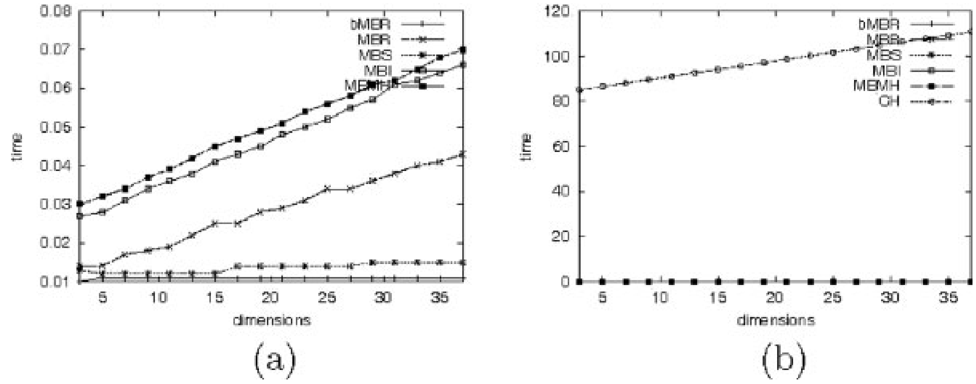

Figure 7. Retrieval times in seconds of the compared techniques.

concentration of query points: points inside the hull do not affect the constraints of the query, and therefore their distribution should not skew the bounding shape. We propose to approximate the relevant sets with a hyper rectangle determined by boundary points. This effectively restricts the spanning dimensions of the approximation shape as well as eliminating the effects of any skew distribution of query points. Specifically, the retrieval set H by bMBR includes all points in the hyperrectangle whose side in dimension i is from $\min_i$ to $\max_i$:

$$H = \{P | P(S, \min_i \leq p_i \leq \max_i, 1 \leq i \leq n\}.$$

Figure 4b shows bMBR wrapping around the hull, and it could occupy as few dimensions as the hull's, which typically is small. In the illustration, if the hull is contained in the plane depicted, bMBR is confined in the same plane while the distance-based volumes extend into all n dimensions. In our experiments, bMBR proved to be much more selective than existing techniques.

Clearly, bMBR can be computed very fast (i.e., for moderate sized databases, sequential scanning is suffcient) or can be executed on top of popular index structures, such as R*-tree or Pyramid (Berchtold et al., 1998). When the data's dimensionality is high, we employ dimensionality-reduction techniques to index the data on only dominant dimensions. However, like MBR, shapes defined by bMBR depend on the features' representation, and therefore, should not be employed as the final measure in multipoint retrieval. In our implementation, we use bMBR to prune out the search space before evaluating the linear system. Obviously, the no-false-dismissal property is guaranteed. The indexing/search procedure is straightforward, and is not shown here.

V. PERFORMANCE STUDY

We have conducted extensive experiments to evaluate the performance of our approach. We selected the techniques discussed in Section IV for comparison. They represent popular approaches for multipoint queries in image retrieval. Our main objectives in this section are to evaluate the effectiveness and efficiency of bMBR and the linear system against the recent techniques, and how well our linear system representation captures the semantics of query images in realistic environments.

The dataset consists of about 20,000 images from various categories (Nakazato et al., 2003; Hiransakolwong et al., 2004) For each image, we extracted 37 visual features as suggested in recent methods, including color features (nine components) (Stricker and Orengo, 1995), texture (10 components) (Smith and Chang, 1994) and edge-based structure (18 components) (Zhou and Huang, 2001; Hiransakolwong et al., 2004). Thus, if nine dimensions are used, only the color group is utilized; 19 denotes the use of both color and texture; and all 37 when all groups are included. Limited features were used to determine how adequately images are represented.

A. Performance in Simulated Environments

There were 100 queries for each experiment from which the average results were calculated. Initial queries composed of 3–5 points randomly picked from the same semantic class. Larger queries were generated from initial ones by expand Query with some fixed radius

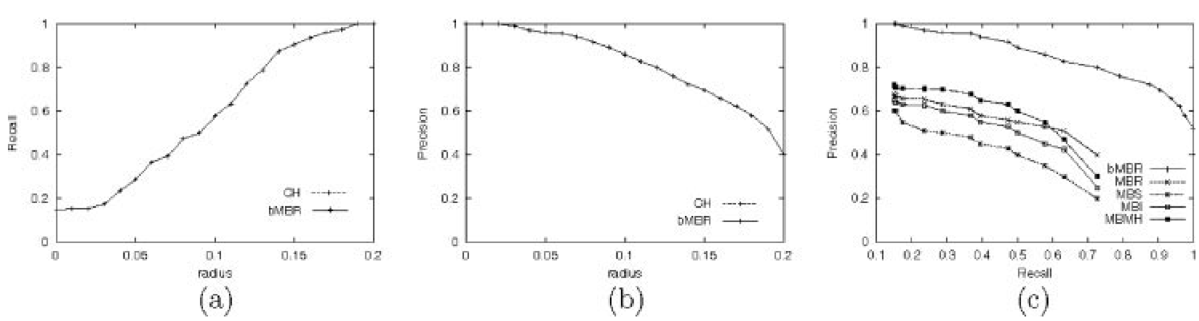

Figure 8. Recall and precision.

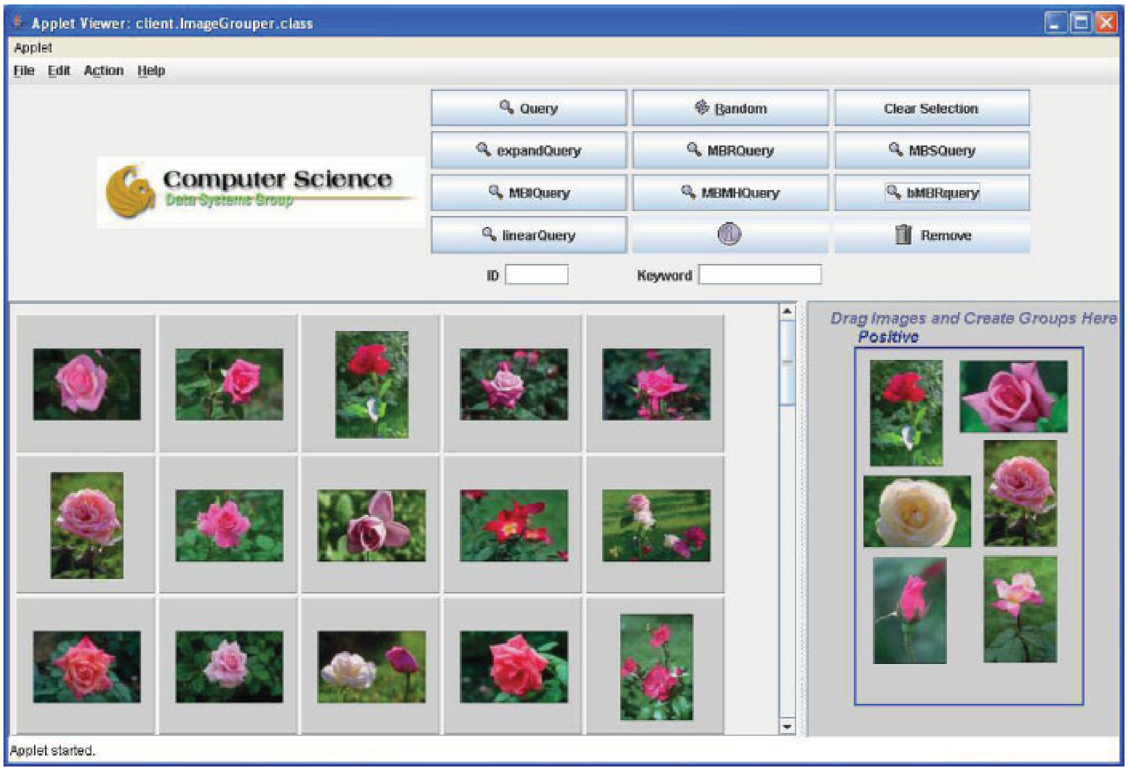

Figure 9. Multipoint image retrieval system. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

from the query points. Since no false dismissals are guaranteed by all techniques, we will use the selectivity (i.e. the ratio of the retrieved to the database's size) of each scheme as the main metric for effectiveness. Their processing times were also recorded.

Figures 5a–5c plot the selectivity of the methods studied as the number of query points varies, with Figures 5a–5c corresponding to the use of 9, 19, and 37 dimensions. They indicate that more query points result in larger retrieval sets. Out of the existing techniques, MBMH appears to be the most effective overall, outperforming MBR, MBS, and MBI. These perform reasonably well for small multipoint queries, around 3 or 4. For larger queries, their approximation volumes expand significantly into all dimensions, resulting in large retrieval sets. Comparing their performances plotted in Figures 5a–5c, their effectiveness almost remains constant even with the use of more features. This is because the more dimensions that are added, the larger their retrieval volumes, offsetting the effects of additional features used.

Clearly, CH and bMBR outdo MBMH, MBR, MBS, and MBI by a wide margin. bMBR closely estimates CH and the more features being used the more effective they are. The results indicate

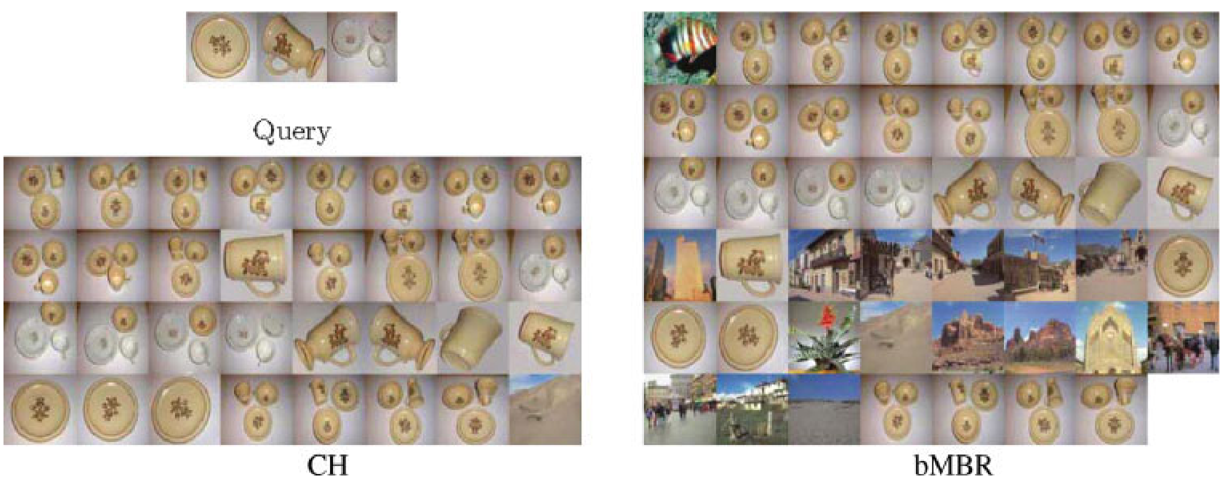

Figure 10. Result by the proposed techniques. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

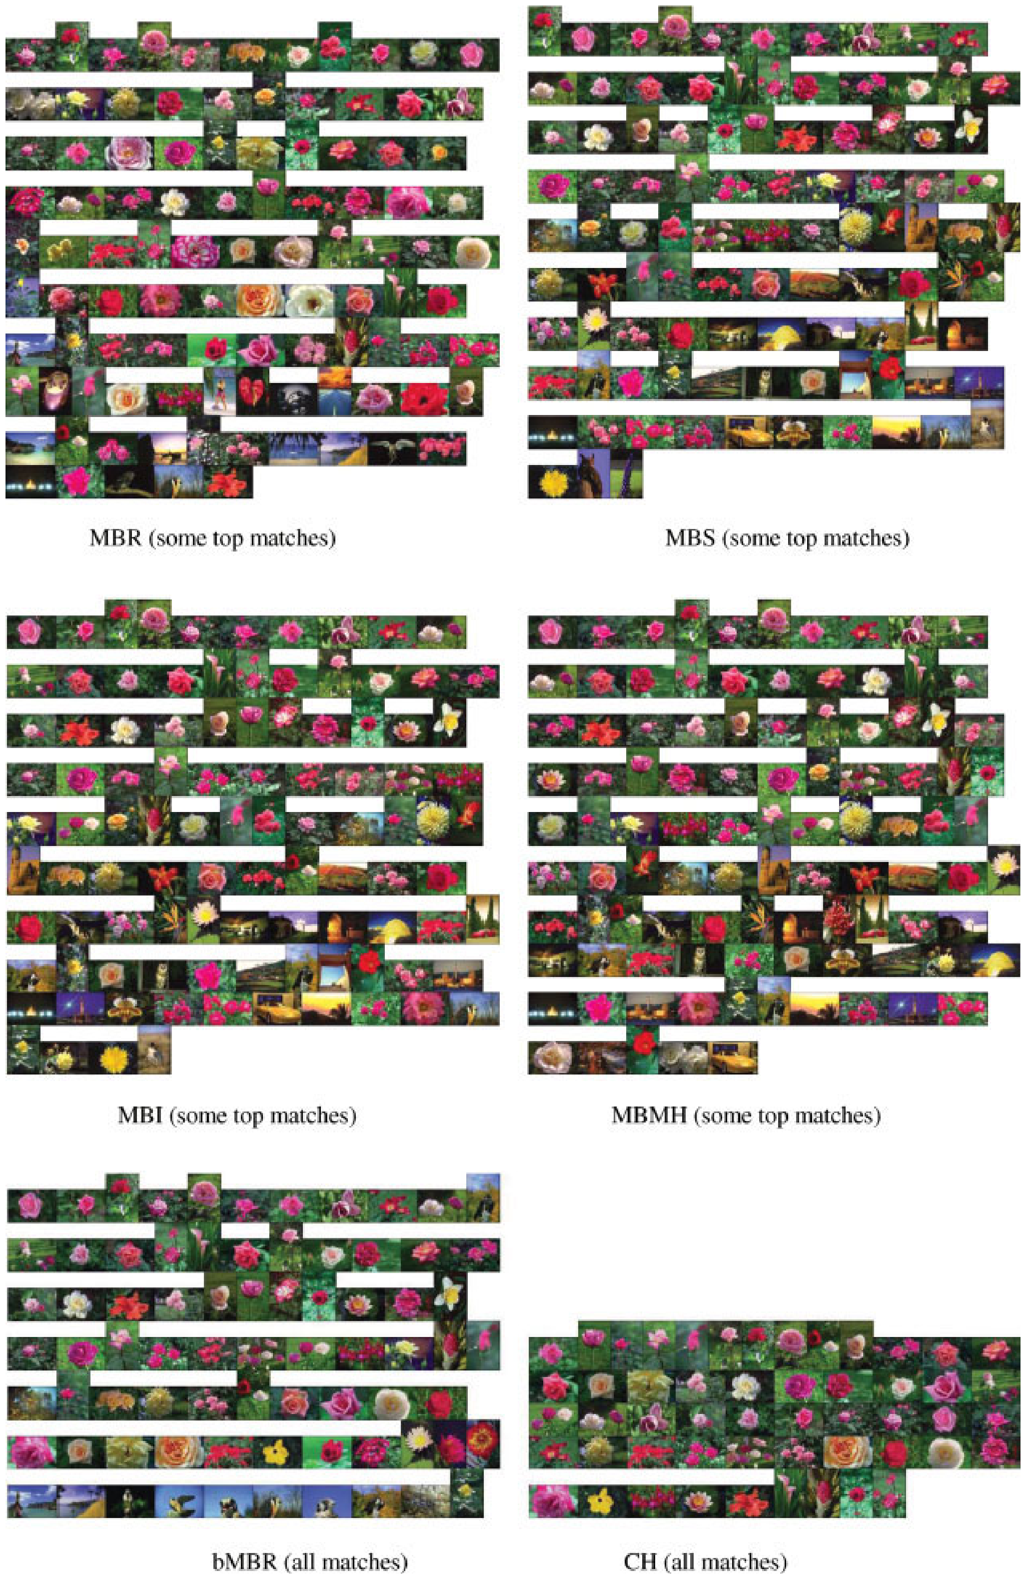

Figure 11. Results by all techniques. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

that exact sets (i.e., CH) occupy only small subspaces of the feature space and that bMBR is capable of restricting the spanning dimensions of its approximation shapes. More dimensions help isolate irrelevant images from relevant ones. This also confirms that image features suggested by recent works suffciently describe essential characteristics of images.

Figures 6a–6c plot the selectivity of the studied methods as the number of dimensions varies. They look closely at the performances when the queries are 3, 9, and 20 points. Note their high retrieval rates at low dimensions, e.g., Figure 6c. This was caused by the large dataset got compacted in a small data space. It also explains that colors alone (i.e., at 9D or fewer) do not represent images suffciently, and both bMBR and CH retrieve large sets as many irrelevant images get mapped onto the tight color subspace. When all features are used, their effectiveness improves significantly.

Since some of the compared techniques do not support indexing (see Section IV), we scanned the dataset sequentially for relevant ones. Figures 7a and 7b show the processing times of the schemes versus dimensions. Observe that the simple measure of bMBR consumes the least time to produce the results compared to the weighted, aggregate, or disjunctive measures used by the existing techniques, see Figure 7a. Figure 7b displays the scanning time of CH solving the linear system for every data points. Taking on average 105 s, CH is not suitable for moderate-sized image retrieval systems without an efficient indexing structure.

B. Performance in Realistic Environments

With images sufficiently represented, it is expected that the proposed measure closely match semantic classes in actual implementations. In these experiments, the ground truth of a multipoint query was the actual relevant images in the database according to the user. Given an initial query (e.g., with three query images), recall and precision are a function of its search radius. Figures 8a and 8b show Recall and Precision of bMBR and CH for various search radii (Plots of CH are not seen because its Recall is identical to bMBR and its Precision is almost perfect). Figure 8c plots Recall vs. Precision of the proposed approach against the compared techniques, showing that the proposed measure is much more effective than those of the existing multipoint ones. The results also indicate that the features used represent images remarkably well that CH closely approximates the semantic groups.

Figure 9 shows our prototype, based on ImageGrouper (Nakazato et al., 2003) executing a 6-point query (in the right pane) with an expansion radius of 0.15. Figure 11 shows the results of this query by all schemes. Since the returns by the existing techniques are very large (over a thousand images by MBI and MBR), we display only some top matches (60) as ranked by the system, while all the results by bMBR and CH are shown (the relevant images by bMBR and CH are not ranked, since they are set-based.) As seen, there are no false matches in the set returned by CH.

Another sample query is shown in Figure 10. The intent of the query is to retrieve China sets whose characteristics are captured in color, texture, and shape. With an expansion radius of 0.175, the semantic class is fully covered, which consists of 38 images. Hence, bMBR's recall and precision for this search radius are respectively 0.895 and 0.723, while CH contains only one false match. Picture the shape of the relevant set of this query. The initial query is a 2D triangle while the expanded query is a 37D ''slice'' of thickness 0.35 and a maximum diameter of 1.374.

VI. CONCLUDING REMARKS

We proposed a novel image retrieval approach for multipoint queries that was inspired by recent findings. We showed that weighted centroids are always relevant with respect to the constraints set by queries, and can be described by a linear system of equations with constraints. A general search algorithm was presented. We evaluated our boundary-based approximation method and recently proposed techniques. Results showed that our proposed methods works effectively and performs well in actual implementation.

REFERENCES

R. Agrawal, C. Faloutsos, and A. Swami. Efficient similarity search in sequence databases, Proc FODO Conf (1993).

Available at http://www.cse.unsw.edu.au/lambert/java/3d/ giftwrap.html.

C. Barber, D. Dobkin, and H. Huhdanpaa. The quickhull algorithm for convex hulls, ACM Trans Math Softw 22 (1996), 469–483.

N. Beckman, H. Kriegel, R. Schneider, and B. Seeger. The r*-tree: An efficient and robust access method for points and rectangles, ACM SIGMOD (1990), 322–331.

S. Berchtold, C. Bohm, and H. Kriegel. The pyramid-technique: Toward breaking the curse of dimensionality, Proc ACM SIGMOD Int Conf Manag Data (1998), 142–153.

T.E. Bjoerge and E.Y. Chang. Why one example is not enough for an image query, Proc IEEE Int Conf Multimedia, Taipei, 2004.

C. Carson, M. Thomas, S. Belongie, J. Hellerstein, and J. Malik. Blobworld: A system for region-based image indexing and retrieval, Third Int Conf Vis Inf Syst (1999), 509–516.

K. Chakrabarti, M. Ortega-Binderberger, S. Mehrotra, and K. Porkaew. Evaluating refined queries in top-k retrieval systems, IEEE Trans Knowled Data Eng 16 (2004), 256–270.

T. Chang and C.J. Kao. Texture analysis and classification with tree-structured wavelet transform, IEEE Trans Image Process 2 (1993), 429–441.

C. Faloutsos, M. Ranganathan, and Y. Manolopoulos. Fast subsequence matching in time-series databases, Proc ACM SIGMOD (1994), 419–429.

D.Q. Goldin and P.C. Kanellakis. On similarity queries for time-series data: Constraint specification and implementation, Proc CP'95 (First Int Conf Principles and Pract of Constraint Programming) (1995), 137–153.

N. Hiransakolwong, K. Vu, K.A. Hua, and S. Lang. Shape recognition based on the medial axis approach, IEEE Int Conf Multimedia Expo (ICME) (2004), 24–28.

M. Hiyahara and Y. Yoshida. Mathematical transform of (r, g, b) color data to munsell (h, v, c) color data, SPIE Vis Commun Image Proc (1988), 650–657.

K.A. Hua, N. Yu, and D. Liu. Query decomposition: A multiple neighborhood approach to relevance feedback processing in content-based image retrieval, Proc IEEE ICDE Conf (2006).

L. Huitao and E. Alexandros. On face detection in the compressed domain, Proc ACM Multimedia (2000), 285–294.

C.E. Jacobs, A. Finkelstein, and D.H. Salesin. Fast multiresolution image querying, Proc SIGGRAPH95 Annu Conf Ser (1995), 277–286.

K.V.R. Kanth, D. Agrawal, A.E. Abbadi, and A. Singh. Dimensionality reduction for similarity searching in dynamic databases, ACM SIGMOD Conf Proc (1998), 166–176.

E.J. Keogh, K. Chakrabarti, S. Mehrotra, and M.J. Pazzani. Locally adaptive dimensionality reduction for indexing large time series databases, SIGMOD Conf (2001), 151–162.

D.-H. Kim and C.-W. Chung. Qcluster: Relevance feedback using adaptive clustering for content-based image retrieval, Proc ACM SIGMOD Conf (2003), 599–610.

S. Koompairojn, K.A. Hua, N. Hiransakolwong, and K. Vu. Range distance in query-by-example(s): An application, IASTED Int Conf Adv Comput Sci Technol (ACST2004) (2004), 271–276.

B. Manjunath and W.Y. Ma. Texture feature for browsing and retrieval of image data, IEEE Trans Pattern Anal Machine Intell 18 (1996), 837–842.

D.T.X.L.S.J. Maybank. Negative samples analysis in relevance feedback, IEEE Trans Knowled Data Eng (TKDE) 19 (2007), 568–580.

M. Nakazato, L. Manola, and T.S. Huang. ImageGrouper: A group-oriented user interface for content-based image retrieval and digital image arrangement. J Vis Lang Comput 14 (2003), 363–386.

A. Natsev, R. Rastogi, and K. Shim. Walrus: A similarity retrieval algorithm for image databases, Proc 1999 ACM SIGMOD Int Conf Manag Data (1999), 395–406.

W. Niblack, R. Barber, W. Equitz, M. Flicker, E. Glasman, D. Petkovic, P. Yanker, and C. Faloutsos. The qbic project: Query images by content using color, texture and shape, SPIE V1908 (1993).

M. Ortega-Binderberger and S. Mehrotra. Relevance feedback techniques in the mars image retrieval system, Multimedia Sys 9 (2004), 535–547.

A. Paradopoulos and Y. Manolopoulos. Performance of nearest neighbor queries in r-trees. Int Conf Database Theory (1997), 394–408.

K. Porkaew, K. Chakrabarti, and S. Mehrotra. Query refinement for multimedia similarity retrieval in mars, Proc 7th ACM Multimedia Conf (1999), 235–238.

J.J. Rocchio. The SMART retrieval system-experiments in automatic document processing, Prentice-Hall, Englewood Cliffs, N.J., 1971.

S.S. Skiena. Convex hull in the algorithm design manual, Springer-Verlag, New York, 1997.

J.R. Smith and S.-F. Chang. Transform features for texture classification and discrimination in large image databases, Proc Int Conf Image Process (1994), 407–411.

J.R. Smith and S.F. Chang. Visualseek: A fully automated content-based image query system, Proc 1996 ACM Int Multimedia Conf (1996), 87–98.

M.A. Stricker and M. Orengo. Similarity of color images, Proc Storage Retrieval Image Video Databases (SPIE) (1995), 381–392.

D. Tao and X.T.X. Li. Which components are important for interactive image searching?, IEEE Trans Circuits Syst Video Technol (TCSVT) 18 (2008), 3–11.

D. Tao, X. Tang, X. Li, and Y. Rui. Direct kernel biased discriminant analysis: A new content-based image retrieval relevance feedback algorithm, IEEE Trans Multimedia (TMM) 8 (2006a), 716–727.

D. Tao, X. Tang, X. Li, and X. Wu. Asymmetric bagging and random subspace for support vector machine-based relevance feedback in image retrieval, IEEE Trans Pattern Anal Machine Intell (TPAMI) 28 (2006b), 1088–1099.

K. Vu, K.A. Hua, H. Cheng, and S.-D. Lang. A non-linear dimensionality-reduction technique for fast similarity search in large databases, ACM SIGMOD Conf Proc (2006a), 527–538.

K. Vu, K.A. Hua, N. Hiransakolwong, and S. Nilpanich. Recognition of enhanced images, Eleventh Int Multimedia Model Conf (MMM2005) (2005), 239–246.

K. Vu, K.A. Hua, and N. Yu. A linear system of equations with constraints for multipoint queries in image retrieval, Twelfth Int Multimedia Model Conf (MMM2006) (2006b), 527–538.

R. Weber, H. Schek, and S. Blott. A quantitative analysis and performance study for similarity-search methods in high-dimensional spaces, Proc 24th VLDB Conf (1998), 194–205.

L. Wu, C. Faloutsos, K. Sycara, and T.R. Payne. Falcon: Feedback adaptive loop for content-based retrieval, Proc 26th VLDB Conf (2000), 297–306.

X.S. Zhou and T.S. Huang. Edge-based structural features for content-based image retrieval, Pattern Recognit Lett 22 (2001), 457–468.