Using Relevance Feedback in Multimedia Databases

Chotirat Ann Ratanamahatana Eamonn Keogh

Department of Computer Science and Engineering, University of California, Riverside, CA 92521

{*ratana, eamonn*)@cs.ucr.edu

ABSTRACT

Much of the world's data is in the form of time series, and many other types of data, such as video, image, and handwriting, can easily be transformed into time series. This fact has fueled enormous interest in time series retrieval in the database and data mining community. We argue, however, that much of this work's narrow focus on efficiency and scalability has come at the cost of usability and effectiveness. In this work, we introduce a general framework that learns a distance measure with arbitrary constraints on the warping path of the Dynamic Time Warping calculation. We demonstrate utility of our approach on both classification and query retrieval tasks for time series and other types of multimedia data including images, videos, and handwriting archives.

1. INTRODUCTION

Much of the world's data is in the form of time series, and many other types of data, such as video, image, and handwriting, can also be trivially transformed into time series. This fact has fuelled enormous interest in time series retrieval in the database and data mining community. We argue, however, that much of this work's narrow focus on efficiency and scalability has come at the cost of usability and effectiveness. For example, the lion's share of previous work has utilized the Euclidean distance metric, presumably because it is very amenable to indexing [2][5][7]. However, there is increasing evidence that the Euclidean metric's sensitivity to small differences in the time axis makes it unsuitable for most real world problems [1][4][6][11][22][26]. This fact appears to have gone almost unnoticed because, unlike their counterparts in information retrieval, many researchers in the database and data mining community evaluate algorithms without considering precision/recall or accuracy [13].

In this work, we introduce a new distance measure and empirically show its utility with thorough experiments measuring the precision/recall and accuracy. While we will demonstrate that our measure is the best in literature, it has a potential weakness; It requires some training or human intervention to achieve its finer results. However, to achieve this end, we will show that the classic information retrieval technique of relevance feedback can be used.

The rest of the paper is organized as follows. The remainder of this section will familiarize readers with the time series and its tight connection with other types of multimedia data. Section 2 gives a review of Dynamic Time Warping (DTW), and related work. In Section 3, we introduce our approach which is based on learning domain specific (and possibly class specific) constraints on the popular DTW distance measure using a representation we call the *R-K Band*. Section 4 demonstrates how this framework is used for relevance feedback then reports the empirical evaluation on three real-world datasets. Lastly, Section 5 gives conclusions and direction for future work.

1.1 The Ubiquity of Time Series Data

In this section, we wish to expand the readers' appreciation for the ubiquity of time series data. Rather than simply list the traditional application domains, i.e. stock market data, electrocardiograms, weather data, etc., we will consider some less obvious applications that can benefit from efficient and effective retrieval.

Video Retrieval: Video retrieval is one of the most important issues in multimedia database management systems. Generally, research on content-based video retrieval represents the content of the video as a set of frames, leaving out the temporal features of frames in the shot. However, for some domains, including motion capture editing, gait analysis, and video surveillance, it may be fruitful to extract time series from the video, and index *just* the time series (with pointers back to the original video). Figure 1 shows an example of a video sequence that is transformed into a time series. There are several reasons why using the time series representation may be better than working with the original data. One obvious point is the massive reduction in dimensionality, which enhances the ease of storage, transmission, analysis, and indexing. In addition, it is much easier to make the time series representation invariant to distortions in the data, such as time scaling and time warping.

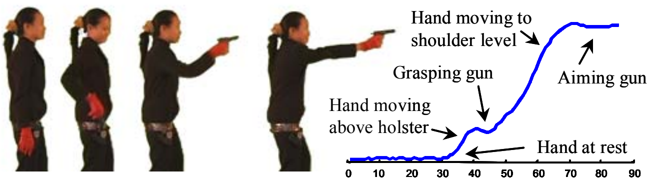

Figure 1 Stills from a video sequence; the right hand is tracked, and converted into a time series

Image Retrieval: Image Retrieval has become increasingly crucial in our information-based community. Large and distributed collections of scientific, artistic, technical, and commercial images have become more prevalent, thus requiring more sophisticated and precise methods for users to perform similarity or semantic based queries. For some specialized domains, it can be useful to convert the images into "pseudo time series". For example, consider Figure 2 below. Here, we have converted an image of a leaf into a time series by measuring the local angle of a trace of its perimeter. The utility of such a transform is similar to that for video retrieval.

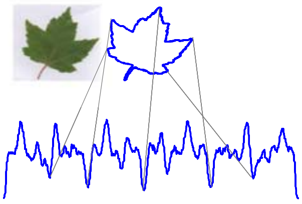

Figure 2. Many image indexing/classification tasks can be solved more effectively and efficiently after converting the image into a "time series"

Handwriting Retrieval: While the recognition of *online* handwriting [10] may be largely regarded as a solved problem, the problem of transcribing and indexing existing historical archives remains a challenge. The usefulness of such ability is obvious. For even such a major historical figure as Isaac Newton, there exists a body of unpublished, handwritten work exceeding one million words. For other historical figures, there are even larger collections of handwritten text. Such collections are potential goldmines for researchers/biographers.

Figure 3.A shows an example of text written by George Washington, which is all but illegible to modern readers with little experience with cursive writing.

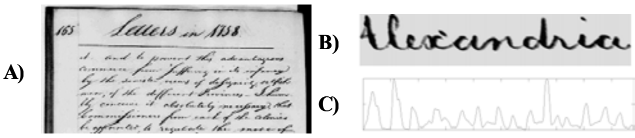

Figure 3.A) An example of handwritten text by George Washington. B) A zoom-in on the word "Alexandria", after being processed to remove slant. C) Many techniques exist to convert 2-D handwriting into a time series; in this case, the projection profile is used (Fig. created by R. Manmatha)

Many off-line handwritten document image-processing algorithms have recently been proposed in the interest of word recognition and indexing [12]. While handwriting is not a time series, there exist several techniques to convert handwriting to (one or more) time series; many of these transformations were pioneered by Manmatha and students [17].

1.2 Existing Work on Time Series Retrieval

The explosion of interest in time series indexing in the last decade has been extraordinary, with well over a thousand papers devoted to the subject [13]. However, the vast majority of the work has focused on the Euclidean distance; recent work has demonstrated that this similarity model generally does not work well for many real-world problems since even very similar time series often demonstrate some variability in the time axis. The problem of distortion in the time axis can be addressed by Dynamic Time Warping (DTW), a distance measure that has long been known to the speech processing community [15][16][20][23]. This method allows for non-linear alignments between the two time series to accommodate sequences that are similar but out of phase, as shown in Figure 4.C.

Our approach takes this recent work on DTW as its starting point. In particular, DTW is currently viewed as a "one-size-fits-all" algorithm, which is applied to diverse domains in a black box fashion. We note, however, that we may be able to fine-tune the algorithm, for a particular domain and even a particular query, by selectively limiting the amount of warping we allow along various parts of the query. For example, in Figure 4, we can see that the first 1/5 and the last 1/6 of the time series do not require warping. This happens to be true for all instances in this particular domain. As we will demonstrate, by selectively limiting the amount of warping allowed, we can actually improve the accuracy of DTW, and as an important side effect, we can drastically improve the indexing performance. Before formally introducing our technique, we must review the basic DTW algorithm in some detail.

2. BACKGROUND

Suppose we have two time series, a query sequence $Q = q_1, q_2, \ldots, q_i, \ldots, q_n$ of length n and a candidate sequence $C = c_1, c_2, \ldots, c_j, \ldots, c_m$ of length m. The DTW algorithm finds the optimal time alignment between these two given time series. To align two sequences using DTW, an n-by-m matrix is constructed where the (i^{th}, j^{th}) element of the matrix corresponds to the cumulative squared distance, $d(q_i, c_j) = (q_i - c_j)^2$, the alignment between points q_i and c_j. To find the best match between these two sequences, the path through the matrix that minimizes the total cumulative distance between them is discovered. This is illustrated in Figure 4. A warping path, W, is a contiguous set of matrix elements that characterizes a mapping between Q and C. The k^{th} element of W is defined as $w_k = (i,j)_k$. By definition, the optimal path W_o is the path that minimizes the warping cost:

$$DTW(Q,C) = \min\left\{ \sqrt{\sum_{k=1}^{K} w_k} \right. \quad (1)$$

This path can be found using dynamic programming to evaluate the following recurrence which defines the cumulative distance $\gamma(i,j)$ as the distance $d(i,j)$ found in the current cell and the minimum of the cumulative distances of the adjacent elements:

$$\gamma(i,j) = d(q_i, c_j) + \min\{ \gamma(i-1, j-1), \gamma(i-1, j), \gamma(i, j-1) \} \quad (2)$$

Although the dynamic programming algorithm reduces the (potentially exponential) number of paths which we must consider to a "mere" $m * n$, this may still be prohibitively large for many problems. The following well known constraints further reduce the number of warping paths that must be considered:- *Boundary Conditions, Continuity Condition, Monotonic condition, and Adjustment Window Condition.*

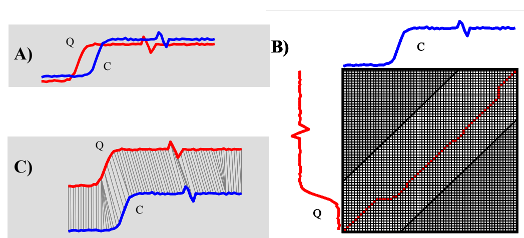

Figure 4. A) Two similar sequences Q and C, but out of phase. B) To align the sequences, we construct a warping matrix, and search for the optimal warping path, shown with solid squares. Note that the "corners" of the matrix (shown in dark gray) are excluded from the search path as part of an Adjustment Window condition. C) The resulting alignment

By applying the above conditions, we can restrict the moves that can be made from any point in the path and so restrict the number of paths that need to be considered. Figure 4.B illustrates a particular example of the last condition with the Sakoe-Chiba Band. Since a good path is unlikely to wander very far from the diagonal, the distance that the path is allowed to wander is within the window of size *r,* above and to the right of the diagonal. As we will see in Section 3, it is this type of constraint that we will exploit to improve DTW.

2.1 Related Work

There has been relatively little work on relevance feedback for both time series and multimedia retrieval. However, relevance feedback in text-mining community has been the subject of much research since the 1970's [3][18][19] and still is an active area of research. It is only in recent years that the researchers started to expand relevance feedback into time series [14], image [25], and multimedia retrieval domains.

Before addressing the relevance feedback system with DTW, we first must introduce our representation, the *R-K Band*, which will be used for the DTW distance measure in the classification task and relevance feedback.

3. RATANAMAHATANA-KEOGH BAND

The 'Adjustment Window Condition' discussed in Section 2 has been almost universally applied to DTW, primarily to prevent unreasonable warping and to speed up its computation. However, surprisingly little research has looked at discovering the best *shape* and *size* of the window. Most practitioners simply use one of the well-known bands, e.g. Sakoe-Chiba Band [20] or Itakura Parallelogram [9], proposed in the context of speech recognition several decades ago. In addition, there is a widespread but unwarranted belief that having wider bands improves accuracy, and having narrower bands decreases accuracy. The use of smaller-size band is seen as a compromise made to make the algorithm tractable. This belief has been proved to be false by our extensive experiments on wide variety of datasets; surprisingly, the accuracies often peak at smaller-size window, and degrade or become stable for wider window sizes. The motivation for our work has sparked from this discovery; we find that in general, the effect of the window size on accuracy is very substantial, and is strongly domain dependent. And if the *width* of the band can greatly affect accuracy, then the *shape* of the band could also have similarly large effects. Our ideal solution would be to find an optimal band (both shape and size) for a given problem that will potentially increase the accuracy. In the next section, we first introduce our representation, the *R-K Band,* which allows user to specify arbitrary shaped constraints.

3.1 A General Model of Global Constraints

We can represent any warping window as a vector R:

$$R_i = d \qquad 0 \le d \le m,\ \ 1 \le i \le m \qquad (3)$$

where R_i is the height above the diagonal in the y direction, as well as the width to the right of the diagonal in the x direction. Note that $|R| = m$, and the above definition forces R to be symmetric, i.e. the constraint above the diagonal is the mirror image of the one below the diagonal.

To represent a Sakoe-Chiba Band of overall width of 11 (width 5 strictly above and to the right of the diagonal) with the definition:

$$R_i = \begin{cases} 5 & 1 \le i \le m-5 \\ m-i & m-5 < i \le m \end{cases} \qquad (4)$$

or an Itakura Parallelogram with the definition:

$$R_i = \begin{cases} \left\lfloor \frac{2}{3} i \right\rfloor & 1 \le i \le \left\lfloor \frac{3}{8} m \right\rfloor \\ \left\lfloor \frac{3}{8} m \right\rfloor - \left\lfloor \frac{2}{5} i \right\rfloor & \left\lfloor \frac{3}{8} m \right\rfloor < i \le m \end{cases} \qquad (5)$$

The classic Euclidean distance can also be defined in terms of $R_i = 0;\ 1 \le i \le m$; only the diagonal path is allowed. More generally, we can define any arbitrary constraint with a suitable vector R. Figure 5 illustrates some examples of *R-K Bands*.

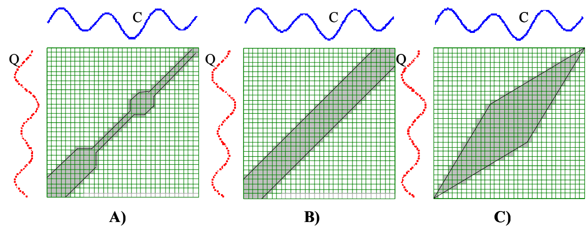

Figure 5. We can use R to create arbitrary global constraints. A) Note that the width of the band may increase or decrease. We can also use R to specify all existing global constraints, e.g. Sakoe-Chiba Band B) and Itakura Parallelogram C)

An interesting and useful property of our representation is that it also includes the ubiquitous Euclidean distance and classic DTW as special cases. We also can exploit the *R-K Bands* for both classification and indexing (query retrieval) problems, depending on the task at hand. In particular,

- for classification, we can use a different *R-K Band* for each class; we denote the band learned for c^{th} class as the $R\text{-}K_c$ *Band.*
- for indexing, we can use *one R-K Band* that maximizes the trade off between efficiency and precision/recall.

Having introduced an *R-K Band*, we can easily represent any arbitrary warping windows. However, we are left with the question of how to *discover* the optimal *R-K Band* for the task at hand. In some cases, it maybe is possible to manually construct the bands, based on domain knowledge. For example, a cardiologist may know from experience that the Romano-Ward syndrome may manifest itself with high variability in the length of one part of the heartbeat (the QT-wave), but little variability in the other section of a heartbeat (the UP-wave)[24]. We could explicitly attempt to encode this insight into an *R-K Band* for retrieving instances of the disease, allowing R_i to be large where variability is expected.

Unfortunately, our preliminary attempts to manually construct *R-K Bands* met with limited success, even for simple toy problems. Furthermore, since the number of possible *R-K Bands* is exponential, exhaustive search over all possibilities is clearly not an option. In the following sections, however, we will show how we can *learn* the high-quality bands automatically from the data.

3.2 Learning Multiple $R\text{-}K_c$ *Bands* for Classification

While it is generally not possible to handcraft accurate *R-K Bands*, it *is* possible to pose the problem as a search problem, and utilize classic search techniques from the artificial intelligence community. Using generic heuristic search techniques, we can perform both forward and backward searches. The forward search starts with the initial Sakoe-Chiba band (uniform) of width 0 (Euclidean), and the backward search starts from the uniform band of the maximum width *m*, above and to the right of the diagonal.

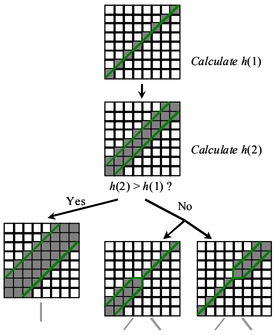

Figure 6. An illustration of our forward search algorithm

Due to space limitations, we will only give a simple intuition behind this approach as illustrated in Figure 6. For a forward search, we start off with Euclidean Bands, one for each class. Then it takes turns, one class at a time, trying to expand or increment the whole section of the band before re-evaluating its overall accuracy for that particular band. If an improvement is made, we keep on expanding that section; otherwise, we undo the expansion, split that section in half, and then recursively expand each portion individually before another re-evaluation. Backward search is very similar; except that we start off with a very wider band then try to tighten the band instead of expanding. Here, we do not include a bi-directional search, a straightforward combination of the forward and backward search; this is omitted in this work for brevity.

The searches are complete when one of the following is true:- No improvement in accuracy can be made; the width of the band reaches *m* for the forward search and 0 (Euclidean) for the backward search or; each section of the band (after recursively cut the portion in half) reaches some threshold. We set a threshold to be some function of *m*.

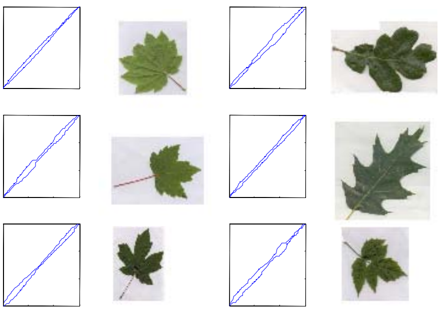

Figure 7. The *R-K*$_c$ *Bands* learned from 6 different species

We can illustrate the utility of our *R-K*$_c$ *Bands* for classification by the following simple experiment. We tested various similarity measures (Euclidean, DTW with 10% warping, DTW with best uniform warping, and DTW with *R-K*$_c$ *Bands*) on the Leaf dataset (dataset details in 4.3.2) and measure their classification error rates. Euclidean is very fast but inaccurate, giving 34.16% error rate. DTW with 10% uniform warping gives a big improvement with only 4.52%. However, the best uniform warping size for this dataset is at 8.6% window size, giving 4.3%. With *R-K*$_c$ *Bands*, we produce 6 different bands, one for each class shown in Figure 7. Classification using these bands gives us almost a perfect result, a mere 0.9% error rate. These promising results suggest that *R-K*$_c$ *Bands* are very effective in improving accuracy in classification.

3.3 Learning One *R-K Band* for Indexing

In addition to creating *R-K*$_c$ *Bands* for classification, we can learn one single *R-K Band* for indexing or query retrieval. The one-band learning algorithm is very similar to the multiple-band learning in the previous section, except that we only maintain one single band that represents the whole problem and that we measure the precision/recall instead of the accuracy.

We re-illustrate this approach by another simple experiment, measuring precision and recall for indexing. We take 10 examples of *Cylinders* from the Cylinder-Bell-Funnel dataset [6][11] and place them in a database containing another 10,000 *random-walk* sequences that are similar in shape but do not belong in the class. Another 30 examples of *Cylinders* with 470 *random-walk* sequences are used in the *R-K Band* training process. To evaluate our method, another 10 different *Cylinder* examples are used to make 10 iterations of k-nearest neighbor queries to the dataset, using various distance measures (Euclidean, DTW with 10% warping, and DTW with *R-K Band*).

We measure the precision from 1-object (10%) to 10-object (100%) recall levels. The results are shown in Figure 8. It is apparent that utilizing an *R-K Band* in this problem improves both precision and recall by a wide margin, compared to Euclidean and DTW with 10% warping. However, an *R-K Band* needs to be learned from a training data, which may not be practical or available in many circumstances. To resolve this problem, we can build a training data through relevance feedback system, with a little help from the user in identifying the positive and negative examples to the system. We will explain how this works in Section 4, but first we will attempt to develop the readers' intuition as to *why* the *R-K Bands* can produce superior performance

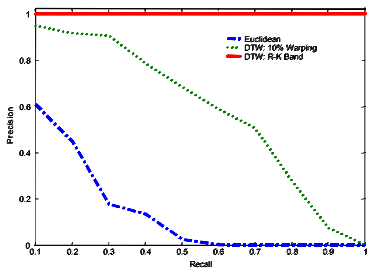

Figure 8. The Precision-Recall curves from 10% to 100% recall for various distance measures: Euclidean, DTW with 10% window size, and our proposed method – *R-K Band* that gives *perfect* precision for all recall levels

3.4 Intuition behind *R-K Bands* Learning

After seeing some examples of our *R-K* Band's utility, we would like to further convince readers by giving an intuition why *R-K Band* improve accuracy. Consider the following problem of face classification based on the head profile. We took a number of photos (20-35) of each individual with different expression on the face, e.g. talking, smiling, frowning, etc. We then use the similar method (see section 1.1) to extract each of the head profile into time series as shown in Figure 9.

We will show by experiment how *R-K Bands* may play an important role in this problem. First, we consider a 2-class problem: a dataset that contains only the collection of profiles from 2 different individuals that look rather different (1 male, 1 female). The *R-K_c Bands* learned from our framework discover the bands both of size zero, the Euclidean distance measure, with 2% error rate. The result suggests that these two individuals are very distinguishable, i.e. the set of time series within each class is much different from another, just by looking at their Euclidean distances. Hence, no warping is necessary; in fact, too much warping could potentially hurt the accuracy because one person could be forced to match with another.

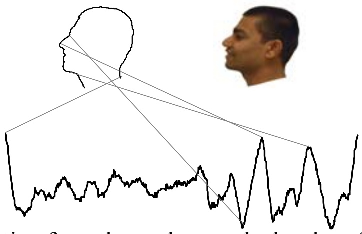

Figure 9. Starting from the neck area, the head profile is converted into a "pseudo time series"

We then extend our experiment by adding 2 more male individuals into our problem (112 instances total). The corresponding *R-K Bands* are learned which give very low error rate of 1.8% (vs. 6.25% for Euclidean).

4. RELEVANCE FEEDBACK

In text-mining community, relevance feedback is well known to be effective method to improve the query performance [3][18][19][21]. However, there has been relatively little research in non-text domains, such as images or multimedia data. In section 1.1, we have introduced time series as an alternative in representing certain types of multimedia data, including special cases of images and video. We will explain in this section how we utilize and incorporate the technique into the relevance feedback system using our proposed framework, *R-K Band*.

4.1 Query Refinement

Relevance feedback methods attempt to improve performance for a particular informational need by refining the query, based on the user's reaction to the initial retrieved documents or objects. In text retrieval in particular, the user's ranking of the document allows reweighing the query terms.

Working with time series retrieval is rather similar to the text retrieval; a user can draw or provide an example of a query and retrieve the set of best matches' retrieval of images/videos/time series. Once the user ranks each of the results, a query refinement is performed such that a *better-quality* query is produced for the next retrieval round. For real time series retrieval (i.e. electrocardiograms or stock market data), the querying interface can show the time series directly. For transformed data of images or video, the underlying time series representation is hidden from the user, and the user sees only thumbnails of the actual images or video snippets. In our system, the user is asked to rank each result in a 4-point scale as shown based on relevance to their informational needs. These rankings are converted into appropriate weights which are used in the query refinement process (averaging the weighted positive results with the current query).

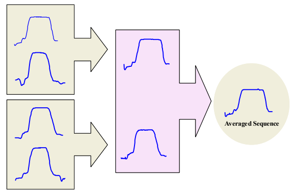

Figure 10. An example of averaging 4 sequences with DTW. Pairs of sequences are hierarchically combined by DTW with their weights until the final averaged sequence is obtained

However, averaging a collection of time series that are not perfectly time-aligned is non-trivial and DTW is needed [8]. Each pair of time series are averaged according to their weights and warping alignment. The results from each pairs are hierarchically combined

Figure 10 illustrates this averaging process using equal weights for all sequences; in practice, the weights may all be different. In the next section, we will show how the relevance feedback system can benefit from our proposed *R-K Band* framework.

4.2 R-K Band in Relevance Feedback

We will empirically demonstrate that our proposed *R-K* Band combined with the query refinement can improve precision and recall of retrieval. Table 1 shows our relevance feedback algorithm.

Table 1: *R-K Band* learning with Relevance Feedback

```
Algorithm RelFeedback(initial_query)
1. Repeat until all rankings are positive.
2.    Show the 10 best matches to the current query to
      the user.
3.    Let the user rank how relevant each result is.
4.    According to the ranking, accumulatively build
      the training set; positive result → class 1,
      negative result → class 2.
5.    Learn a single envelope (R-K Band) that
      represents the given training data.
6.    Generate a new query, by averaging (with DTW)
      the positive results with the current query
      according to their weights (rankings).
7. end;
```

In the first iteration, given a query, the system uses the initial *R-K Band* (the special case of Euclidean distance) to retrieve the 10 nearest neighbors, and then shows them to the user (line 1). When the user finishes their ranking, the positive and negative responses are noted and collected as a training data (lines 3-4). The algorithm uses this training data to learn an *R-K Band* that best represents the positive objects in the training set while being able to correctly differentiate the positive from the negative instances (line 5). The training data will be accumulated during each round, developing a larger training set, thus producing progressively finer results. The process is complete when only positive feedbacks are given to the system or the user abandons the task.

In our experiments, we consider 3 multimedia datasets to be tested using the relevance feedback technique

4.3 Datasets

To evaluate our framework, we measure the precision and recall for each round of the relevance feedback retrieval. Since we only return the 10 best matches to the user and we would like to measure the precision at all recall levels, we purposely leave only 10 relevant objects of interest in all the databases.

4.3.1 Gun Problem

This dataset comes from the video surveillance domain (see Figure 1). The dataset has two classes, 100 examples each:

- **Gun-Draw:** The actors have their hands by their sides. They draw a replicate gun from a hip-mounted holster, point it at a target for approximately one second, and then return the gun to the holster, and their hands to their sides.
- **Point:** The actors have their hands by their sides. They point with their index fingers to a target for approximately one second, and then return their hands to their sides.

For both classes, the centroid of the right hand is tracked both in X- and Y-axes; however, in this experiment, we only consider the X-axis for simplicity. The dataset contains 200 instances, 100 for each class. Each instance has the same length of 150 data points. For the relevance feedback purpose, we only leave 10 Gun-Draw examples in the Point database, and randomly pick another example for an initial query.

4.3.2 Leaf Dataset

This dataset contains a collection of 6 different species of leaf images, including 2 genera of plant, i.e. oak and maple. Maple has 4 different species, and Oak has 2, with 442 instances in total (original images are available at http://web.engr.oregonstate.edu/~tgd/leaves/dataset/herbarium). Each instance is linearly interpolated to have the same length of 150 data points. In our experiment, we choose Circinatum maple as the specie of interest, i.e. only 10 images of such specie are left in the database, and we randomly selected another separate image as our initial query.

4.3.3 Handwritten Word Spotting Dataset

This is a subset of the WordSpotting Project dataset. In the full dataset, there are 2,381 words with four features that represent each word image's profiles or the background/ink transitions. For simplicity, we pick the "background/ink transitions" (feature 4) and use the word "*would*" which occurs in the dataset 11 times for our query. Hence, one is removed to be used as our initial query.

4.4 Experimental Results

We measure the performance of our relevance feedback system with the precision-recall plot from each round of iteration. Figure 11 below shows the precision-recall curves of the three datasets for the first five iterations of relevance feedback. Our experiments illustrates that each iteration gives significant improvement in both precision and recall.

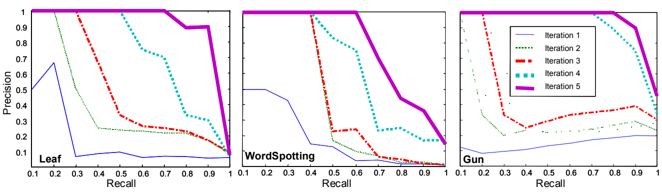

Figure 11. The precision-recall plots for the Gun, Leaf and Word spotting datasets with 5 iterations of relevance feedback

5. DISCUSSION AND CONCLUSIONS

In this work, we have introduced a framework for both classification and time series retrieval. The *R-K Band* allows for any arbitrary shape of the warping window in DTW calculation. With our extensive evaluation, we have shown that our framework incorporated into relevance feedback can reduce the error rate in classification, and improve the precision at all recall levels in video and image retrieval.

6. REFERENCES

[1] Aach, J. and Church, G. (2001). Aligning gene expression time series with time warping algorithms. Bioinformatics. Volume 17, pp. 495-508.

[2] Agrawal, R., Lin, K. I., Sawhney, H. S., & Shim, K. (1995). Fast similarity search in the presence of noise, scaling, and translation in times-series databases. VLDB, pp. 490-501.

[3] Attar, R. & Fraenkel, A.S. (1977). Local Feedback in Full-Text Retrieval Systems. JACM '77), Vol. 24, No. 3, July 1977, pp. 397-417.

[4] Bar-Joseph. Z., Gerber, G., Gifford, D., Jaakkola T & Simon. I. (2002). A new approach to analyzing gene expression time series data. In the 6th Annual Intl Conf. on Research in Comp Molecular Bio. pp. 39-48.

[5] Chan, K.P., Fu, A & Yu, C. (2003). Haar wavelets for efficient similarity search of time-series: with and without time warping. IEEE TKDE, Vol.15, No.3. pp. 686-705.

[6] Diez, J. J. R. & Gonzalez, C. A. (2000). Applying boosting to similarity literals for time series Classification. Multiple Classifier Systems, 1st Inter' Workshop. pp. 210-219.

[7] Faloutsos, C., Ranganathan, M., & Manolopoulos, Y. (1994). Fast subsequence matching in time-series databases. In Proc. ACM SIGMOD Conf., Minneapolis. pp. 419-429.

[8] Gupta, L., Molfese, D., Tammana, R., and Simos, P. (1996). Nonlinear Alignment and Averaging for Estimating the Evoked Potential. IEE Transactions on Biomedical Engineering, Vol. 43, No.4, April 1996.

[9] Itakura, F. (1975). Minimum prediction residual principle applied to speech recognition. IEEE Trans. ASSP-23, pp. 52-72

[10] Jain, A.K. & Namboodiri, A.M. (2003). Indexing and Retrieval of On-line Handwritten Documents. IEEE 7th ICDAR, Vol. II, pp. 655-659.

[11] Kadous, M. W. (1999). Learning comprehensible descriptions of multivariate time series. In Proc. of the 16th ICML pp. 454-463.

[12] Kavallieratou, E., Dromazou, N., Fakotakis, N., & Kokkinakis, G. (2003). An Integrated System for Handwritten Document Image Processing. Int'l J. of Pattern Recognition and AI, 17(4)pp. 617-636.

[13] Keogh, E. and Kasetty, S. (2002). On the Need for Time Series Data Mining Benchmarks: A Survey and Empirical Demonstration. In the 8th ACM SIGKDD. pp. 102-111.

[14] Keogh, E. & Pazzani,M (1999). Relevance feedback retrieval of time series data. In Proceedings of the 22nd ACM-SIGIR. pp 183-190.

[15] Kruskall, J. B. & Liberman, M. (1983). The symmetric time warping algorithm: From continuous to discrete. In Time Warps, String Edits and Macromolecules. Addison-Wesley.

[16] Myers, C., Rabiner, L & Roseneberg, A. (1980). Performance tradeoffs in dynamic time warping algorithms for isolated word recognition. IEEE Trans. ASSP-28, pp. 623-635.

[17] Rath, T. & Manmatha, R. (2003). Word image matching using dynamic time warping. CVPR, Vol. II, pp. 521-527.

[18] Rickman, J.T. (1972). Design Considerations for a Boolean Search System with Automatic Relevance Feedback Processing. SIGIR2: August 1972. Proceedings of the ACM conference-Vol. 1, pp. 478-481.

[19] Rocchio, J.J. (1971). Relevance Feedback in Information Retrieval. The SMART Retrieval System: Experiments in Automatic Document Processing. pp. 313-323.

[20] Sakoe, H. & Chiba, S. (1978). Dynamic programming algorithm optimization for spoken word recognition. IEEE ASSP-26. pp. 43-49.

[21] Salton, G., & Buckley, C. (1990). Improving retrieval performance by relevance feedback. JASIS41(4), pp.288-297.

[22] Schmill, M., Oates, T. & Cohen, P. (1999). Learned models for continuous planning. In *7th International Workshop on AI and Statistics*.

[23] Tappert, C. & Das, S. (1978). Memory and time improvements in a dynamic programming algorithm for matching speech patterns. *IEEE Trans*. ASSP-26, pp. 583-586.

[24] Viskin, S. (2000). Cardiac Pacing in the Long QT Syndrome: review of available data and practical recommendations. Journal Cardiovasc Electrophysiol 11:593-600.

[25] Wu, P. & Manjunath, B.S. (2001). Adaptive Nearest Neighbor Search for Relevance Feedback in Large Image Databases. ACMMM, pp89-97.

[26] Yi, B, K. Jagadish, H & Faloutsos (1998). Efficient retrieval of similar time sequences under time warping. ICDE, pp. 23-27.