A Self-immunizing Manifold Ranking for Image Retrieval

Jun Wu[1], Yidong Li[1], Songhe Feng[1], and Hong Shen[1, 2]

[1] School of Computer and Information Technology, Beijing Jiaotong University, Beijing 100044, China
[2] School of Computer Science, University of Adelaide, SA 5005, Australia
{wuj,ydli,shfeng,hshen}@bjtu.edu.cn

Abstract. Manifold ranking (MR), as a powerful semi-supervised learning algorithm, plays an important role to deal with the relevance feedback problem in content-based image retrieval (CBIR). However, conventional MR has two main drawbacks: 1) in many cases, it is prone to exploit "unreliable" unlabeled images when deployed in CBIR due to the semantic gap; 2) the performance of MR is quite sensitive to the scale parameter used for calculating the Laplacian matrix. In this work, a self-immunizing MR approach is presented to address the drawbacks. Concretely, we first propose an elastic *k*NN graph as well as its constructing algorithm to exploit unlabeled images "safely", and then develop a local scaling solution to calculate the Laplacian matrix adaptively. Extensive experiments on 10,000 Corel images show that the proposed algorithm is more effective than the state-of-the-art approaches.

Keywords: content-based image retrieval, relevance feedback, self-immunizing manifold ranking, elastic *k*NN graph, local scaling.

1 Introduction

With the ubiquitous use of digital images in a large number of practical applications, Content-Based Image Retrieval (CBIR) has drawn substantial research attention in many computer communities during the past two decades [2]. A main challenge in CBIR is the so-called semantic gap, i.e. the low-level visual features are not sufficient to characterize the high-level semantics of images. Relevant feedback has been shown as a powerful tool for bridging the semantic gap by exploiting the user's interaction with CBIR system. During the past years, a wide variety of relevant feedback techniques have been proposed, most of which belong to the family of supervised learning [14, 2].

One critical research topic related to relevance feedback is to learn with few labeled training examples, as few users are patient to label a lot of images during the interaction. To this end, semi-supervised learning [1] has been applied to relevance feedback [3, 4, 8, 9, and 13]. A popular semi-supervised learning method used in CBIR is the manifold ranking (MR) that aims to learn a ranking function by making use of the underlying geometrical structure of the given image database. Previous studies have shown that MR is one of the most promising and successful semi-supervised learning techniques for relevance feedback [3, 7, 10 and 11].

J. Pei et al. (Eds.): PAKDD 2013, Part II, LNAI 7819, pp. 426–436, 2013.

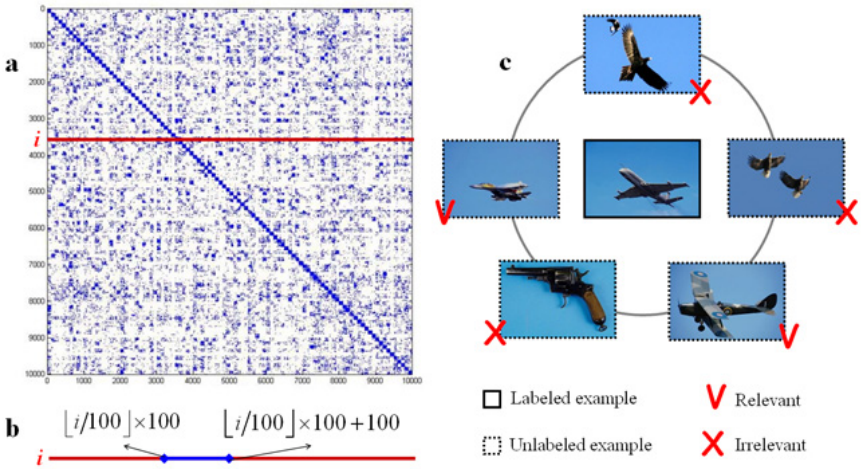

Fig. 1. An **illustration of our motivation**: (a) The adjacency matrix of a *k*NN graph built on a set of 10,000 images, each consecutively-numbered 100 images in which belong to the same semantic category; (b) The "trusted" interval of a labeled image; and (c) An example of a labeled image with "unreliable" unlabeled nearby neighbors.

However, it has been found that the performance of semi-supervised learning may be even worse than the supervised learning when "unreliable" unlabeled data is exploited [5]. Taking MR as an example, it assumes that a labeled example and its (unlabeled) nearby neighbors trend to have similar properties, and thus their ranking scores should be approximate, but this assumption may not be true in CBIR due to the semantic gap. To verify the efficacy of this assumption, we conducted an empirical study on a set of 10,000 images, each consecutively-numbered 100 images in which belong to the same semantic category. Given the image set, a *k*NN graph is constructed and corresponding adjacency matrix is shown by Figure 1a. Ideally, for each image i, we expect its k nearest neighbors appear within a "trusted" interval $\left[\text{floor}(i/100)\times 100,\ \text{floor}(i/100)\times 100+100\right]$ (e.g., if an image id is 1588, its "trusted" interval should be [1501, 1600]), since the images within this interval belong to the same class as illustrated in Figure 1b, where $\text{floor}(\bullet)$ denotes the integer operation. Figure 1a shows that most nonzero elements distribute around the principal diagonal of the adjacency matrix, which means most neighbor points are inside their trusted interval. But there are still many nonzero elements far away from the principal diagonal, i.e. the neighbor points are outside the trusted interval (an example is illustrated by Figure 1c), and, in this case, the performance of MR may degenerate. Moreover, the performance of MR is sensitive to the scale parameter used for calculating the Laplacian matrix. Such a parameter is usually hard to tune with very few labeled examples [11], which is a common issue in graph-based semi-supervised learning.

To address the above problems, this paper presents a **S**elf-**i**mmunizing **ma**nifold **r**anking (Simar) approach for relevance feedback in CBIR, which is able to exploit unlabeled images "safely" and tune the scale parameter adaptively. Concretely, we first

propose a new graph structure named elastic *k*NN graph and corresponding constructing algorithm. In this structure, the creditable relationship between each labeled image and its nearby neighbors can be dynamically adjusted by monitoring the change of retrieval performance. Also, a local scaling solution is employed by Simar to tune the scale parameter for Laplacian matrix calculation, which is beneficial to the data distribution with multi-scales, e.g. image database. Our empirical study shows encouraging results in comparison to some existing semi-supervised learning algorithms widely used in CBIR.

The remainder of this paper is organized as follows. Section 2 elaborates the proposed Simar approach. Section 3 shows experimental evaluations. Finally, section 4 concludes this paper.

2 The Proposed Simar Approach

In this section, we first formulate relevance feedback in CBIR as a semi-supervised graph-based ranking problem, and then present an elastic *k*NN graph structure and a local scaling method to facilitate the setting of parameters.

2.1 Preliminaries

Let $\mathbf{X}=\{\mathbf{x}_1,\cdots,\mathbf{x}_n\}$ denote an image database, where each $\mathbf{x}_i \in \mathbb{R}^d$ represents an image by a *d*-dimensional feature vector. To discover the geometrical structure of the given image database, a graph (such as *k*NN graph) is usually built on $\mathbf{X}$ and we define $\mathbf{W} \in \mathbb{R}^{n\times n}$ as corresponding adjacency matrix with element w_{ij} saving the weight of the edge between point i and j. Normally the weight can be calculated using a Gaussian kernel

$$w_{ij} = \exp\left(-d^2\left(\mathbf{x}_i,\mathbf{x}_j\right)\Big/\sigma^2\right) \tag{1}$$

if $\mathbf{x}_j \in N_k\left(\mathbf{x}_i\right)$ or $\mathbf{x}_i \in N_k\left(\mathbf{x}_j\right)$, otherwise $w_{ij}=0$, where $N_k\left(\mathbf{x}\right)$ denotes the set of the *k* nearest neighbors of $\mathbf{x}$, and $d\left(\mathbf{x}_i,\mathbf{x}_j\right)$ is a distance metric (such as L1 distance) between $\mathbf{x}_i$ and $\mathbf{x}_j$. Finally, we define a label vector as $\mathbf{y}=\left[y_1,\cdots,y_n\right]^T$ to record the user's judgment in relevance feedback loops, in which an element $y_i=1$ if $\mathbf{x}_i$ is the query or labeled as positive, $y_i=-1$ if $\mathbf{x}_i$ is labeled as negative, and $y_i=0$ otherwise.

Given $\mathbf{W}$ and $\mathbf{y}$, the goal of our Simar approach is to learn a ranking function $\mathrm{f}:\mathbf{X}\to\mathbb{R}$ that assigns each image $\mathbf{x}_i$ a ranking score f_i according to its relevance to user's query. Similar to other MR methods, Simar aims to find an optimal f^* by solving the following optimization problem:

$$\mathrm{O}(\mathrm{f}) = \frac{1}{2}\left(\sum_{i,j=1}^{n} w_{ij} \left\| \frac{1}{\sqrt{D_{ii}}} f_i - \frac{1}{\sqrt{D_{jj}}} f_j \right\|^2 + \mu \sum_{i=1}^{n} \left\| f_i - y_i \right\|^2 \right) \tag{2}$$

where $\mu > 0$ is the regularization parameter and $\mathbf{D}$ is a diagonal matrix with $D_{ii} = \sum_{j=1}^{n} w_{ij}$. The first term is a smoothness constraint that makes the nearby images share close ranking scores. The second term is a fitting constraint which means the ranking result should fit to the label assignment. By minimizing $\mathrm{O}(\mathrm{f})$, we get the optimal f by the following closed form

$$\mathrm{f}^* = \left(\mathbf{I}_n - \alpha \mathbf{S}\right)^{-1} \mathbf{y} \tag{3}$$

where $\alpha = 1/(1+\mu)$, $\mathbf{I}_n$ is an identity matrix with $n \times n$, and $\mathbf{S} = \mathbf{D}^{-1/2}\mathbf{W}\mathbf{D}^{-1/2}$ is the symmetrical normalization of $\mathbf{W}$. In large scale problems, we prefer to use the iteration scheme:

$$\mathrm{f}(t+1) = \alpha \mathbf{S}\,\mathrm{f}(t) + (1-\alpha)\mathbf{y}\,. \tag{4}$$

During each round of iteration, each data point receives information from its neighbors (first term). And retains its initial information (second term). The iteration process is repeated until convergence.

As illustrated by Eq. (3) and (4), one of the key issues is to design an appropriate $\mathbf{S}$, and more precisely to design $\mathbf{W}$, which depends on two key parameters: the number of the nearest neighbors k used for constructing kNN graph and the scale parameter σ used by Gaussian kernel. We will discuss how to tune the parameters in the following subsections.

2.2 Constructing an Elastic *k*NN Graph

Constructing an appropriate graph is one of the keys to develop a high-performance MR scheme. As mentioned, the kNN graph is a popularly used structure, but it is prone to exploit "unreliable" unlabeled images, as illustrated by Figure 1. To "safely" exploit unlabeled images, we expect that the constructed graph could dynamically update the k value in a query session, in order to maintain a relatively confidential connecting relationship between each labeled image and its nearby unlabeled neighbors. To this purpose, for each image $\mathbf{x}_i$, we suggest using a large k in our approach when its most neighbors are inside its "trusted" interval because corresponding (unlabeled) nearby neighbors are "reliable" in this case. Conversely, a small k is preferable in order to reduce the likelihood of exploiting the "unreliable" neighbors. At worst, no unlabeled images are considered and our Simar approach will degenerate to a supervised ranking method. In this way, we can guarantee that our semi-supervised ranking method will never worse than a supervised one.

Given the labeled image set, a challenge is to probe whether most of their (unlabeled) nearby neighbors are inside the corresponding "trusted" intervals, since the

images are not indexed by semantic in real-world applications. Considering this, Simar adopts an indirect strategy, that is, in each round of feedback the "reliability" of the unlabeled images used by current ranker is evaluated by monitoring the changes in its retrieval performance. Concretely, the retrieval performance is measured by using the precision rate defined as $Precision = \frac{\text{Number of postive retrievals}}{\text{Number of total retrievals}} \times 100\%$ via the user's feedback on image retrievals. If the current precision $Precision_{cur}$ is greater than the previous precision $Precision_{pre}$, then the "reliability" of the unlabeled images exploited currently is enhanced, and the value of k should be enlarged. On the other hand, if $Precision_{cur} < Precision_{pre}$, then it means that the "reliability" of the unlabeled images exploited currently is receded, and the value of k should be decreased. With these considerations, we adaptively tune the parameter k according to

$$k_{cur} = \text{floor}\left(k_{pre}\left(1 + Precision_{cur} - Precision_{pre}\right)\right). \quad (5)$$

to "safely" exploit the unlabeled images.

Note that the Precision mentioned here is calculated with the number of relevant images that appear in a fixed number of retrievals. Suggested by Luxberg [6], the initial value of k, used at the first round of feedback, is set to $\text{floor}(\log n)$ for the asymptotic connectivity purpose.

2.3 Local Scaling

As mentioned before, the performance of MR is sensitive to the scale parameter σ. Some previous works [3, 7 and 10] suggested running their MR algorithms repeatedly for a number of σ values and selecting the one leading to the highest average precision. However, the performance of this approach is heavily depended on the testing data and the range of values to be tested still has to be set manually. What is worse, there may not be a single value of σ that works well for all data points when the input data with different local statistics, which is the common case in the image database. Therefore, we try to address this shortcoming from a local scaling view, i.e. calculating a local scale parameter for each image, instead of selecting a single scale parameter for all images.

Inspired by the self-tuning spectrum clustering technique [12], the scale parameter can be regarded as some measure when two data points are considered similar. This provides an intuitive way for selecting possible σ. Let σ_i and σ_j denote the local scaling parameters of image $\mathbf{x}_i$ and $\mathbf{x}_j$ respectively. The distance from $\mathbf{x}_i$ to $\mathbf{x}_j$ as 'seen' by $\mathbf{x}_i$ can be defined as $d(\mathbf{x}_i, \mathbf{x}_j)/\sigma_i$ while the converse is $d(\mathbf{x}_j, \mathbf{x}_i)/\sigma_j$. Hence, the square distance d^2 between two images can be generalized as:

$$d(\mathbf{x}_i, \mathbf{x}_j)\, d(\mathbf{x}_j, \mathbf{x}_i)/\sigma_i\sigma_j = d^2(\mathbf{x}_i, \mathbf{x}_j)/\sigma_i\sigma_j, \quad (6)$$

and the weight of the edge between a pair of images, i.e. Eq. (1), can be rewritten as:

$$w_{ij} = \exp\left(-d^2(\mathbf{x}_i, \mathbf{x}_j)/(\sigma_i\sigma_j)\right). \quad (7)$$

Intuitively, a small σ_i is preferable when $\mathbf{x}_i$ is residing in a tight local region, while a large σ_i is preferable when $\mathbf{x}_i$ is residing in a sparse local region. To this purpose, the selection of the local scale σ_i can be done by studying the local statistics of the neighborhood of $\mathbf{x}_i$. Considering the efficiency, we use the distance from $\mathbf{x}_i$ to its *k*-th nearest neighbor $\mathbf{x}_{ik}$ to represent the local statistic, i.e. $\sigma_i = d(\mathbf{x}_i, \mathbf{x}_{ik})$ where $k = \text{floor}(\log n)$ that gave good result in our experiment.

2.4 Implementation Issues

For the real-time response purpose, previous work suggested using a sparse representation for the affine matrix $\mathbf{W}$ and calculating it off-line [3]. However, different from conventional MR, Simar requires updating matrix $\mathbf{W}$ on-line because elastic *k*NN graph is considered. Our idea is to calculate an initial affine matrix with a large *k* value off-line, and then add or remove elements into/from the matrix according to the changes of *k* values on-line. In the way, we can update matrix $\mathbf{W}$ with low computational cost. The key steps are summarized as follows.

Step 1 (off-line): Starting with a large k_0 (=100) value, for each image, we search its k_0 nearest neighbors from database and store their identities in a matrix $\mathbf{G} \in \mathbb{R}^{n \times k_0}$, where each element g_{ij} denotes the identity of the *j*-th nearest neighbor of image $\mathbf{x}_i$. Based on $\mathbf{G}$, the initial affinity matrix $\mathbf{W}_0$ is calculated by Eq. 7.

Step 2 (off-line): In the first round of feedback, given $k_1 = \text{floor}(\log n)$, the affinity matrix $\mathbf{W}_1$ (initialized by $\mathbf{0}^{n \times n}$) is generated based on $\mathbf{W}_0$ and $\mathbf{G}$ by copying elements from $\mathbf{W}_0$ to $\mathbf{W}_1$. For example, given image $\mathbf{x}_i$, the identity of its j-th nearest neighbor is g_{ij}, and corresponding affinity between $\mathbf{x}_i$ and its *j*-th nearest neighbor can be gained by: $\mathbf{W}_1(i, g_{ij}) \leftarrow \mathbf{W}_0(i, g_{ij})$.

Step 3 (on-line): After the second round of feedback, the affinity matrix $\mathbf{W}_{cur}$ used currently is updated by monitoring the changes of retrieval performance. In details, k_{cur} is first calculated by Eq. 5; then, based on $\mathbf{W}_0$ and $\mathbf{G}$, we add elements into $\mathbf{W}_{cur}$ when $k_{cur} > k_{pre}$, while remove elements from $\mathbf{W}_{cur}$ when $k_{cur} < k_{pre}$. The detailed updating rules can be described as:

```
if k_cur > k_pre then /* adding elements into W_cur */
  for i = 1 to n
    for j = k_pre + 1 to k_cur
      W_cur(i, g_ij) ← W_0(i, g_ij);
    end for
  end for
```

```
else  /* removing elements from W_cur */
  for i = 1 to n
    for j = k_pre to k_cur + 1 with step=-1
      W_cur(i, g_ij) ← 0 ;
    end for
  end for
end if
```

Another issue is with respect to the out-of-sample search. If the query image is not in the database, we first connect the query with its k_0 nearest neighbors from database images, meanwhile, add a new row to $\mathbf{G}$, with each element store the identify of the corresponding neighbor. Then, we calculate the edge weights by Eq. 7 and add one row and one column to $\mathbf{W}_0$, with each element equal to the corresponding edge weight. All the other operations will be performed similarly using the enlarged matrix $\mathbf{W}_0$ and $\mathbf{G}$.

3 Experimental Study

In this section, we show several experimental results and comparisons to evaluate the effectiveness of Simar scheme on a real world image database. All algorithms in experiments are implemented in MATLAB 2008 and run on a PC with Intel Core (TM) Duo 2.93 GHZ processor and 2GB RAM.

3.1 Experimental Setup

Experiments are performed on a set of 10,000 images picked from the Corel database. These images belong to 100 semantic classes, each of which has 100 images.

Three different features are used to represent the images, including a 64-dimensional color histogram, an 18-dimensional wavelet-based texture and a 5-dimensional edge direction histogram. At last, each image is represented as an 87-dimensional feature vector.

We use PR-graph and P@TopN to evaluate the effectiveness of image retrieval methods. PR-graph depicts the relationship between precision and recall of a specific retrieval method. In general, a PR-graph can also be summarized into one statistic value, i.e. MAP (mean average precision). However, PR-graph can hardly reflect the changes of retrieval performance caused by feedbacks directly. P@TopN emphasizes the retrieval performance at a particular scope N, which describes the relationship between precision and round of feedback at top N retrieval results. Thus it can compensate for the deficiency of PR-graph.

3.2 Comparison Methods

To examine the efficacy of the proposed Simar approach, several existing semi-supervised learning solutions for relevance feedback in CBIR are compared in our

empirical study. (1) Conventional **MR** [3] serves as a baseline method that applies regular MR algorithm to learn a ranking function. The setting of parameters is consistent with [3], i.e. α=0.99 , σ=0.05 and k=200 . (2) **Co-training** [15] first trains two independent rankers using different distance metrics, and then each ranker labels for the other ranker its two most confident images from unlabeled data for the purpose of enriching the training set. (3) **SemiBoost** [8] iteratively learns an ensemble of SVMs using a similar procedure of boosting algorithm. In particular, both labeled and

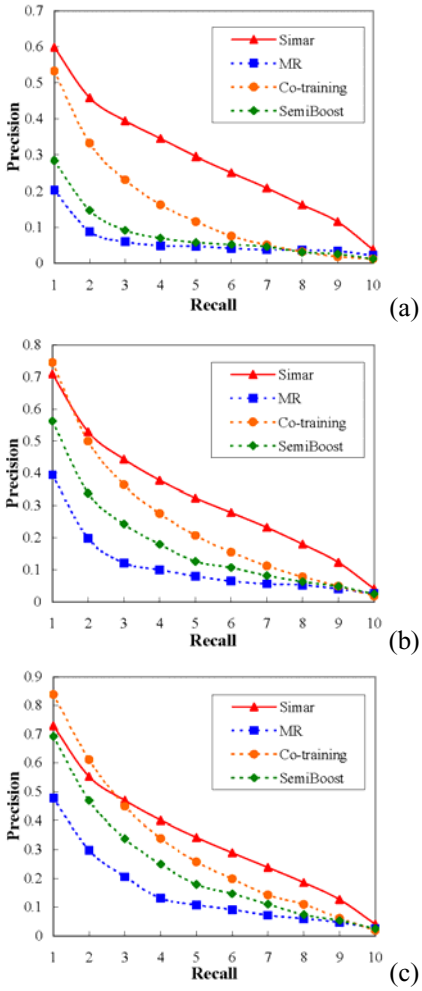

Fig. 2. PR-graphs of the proposed method compared with some existing methods at the (a) 1st, (b) 2nd and (c) 3rd round of feedback.

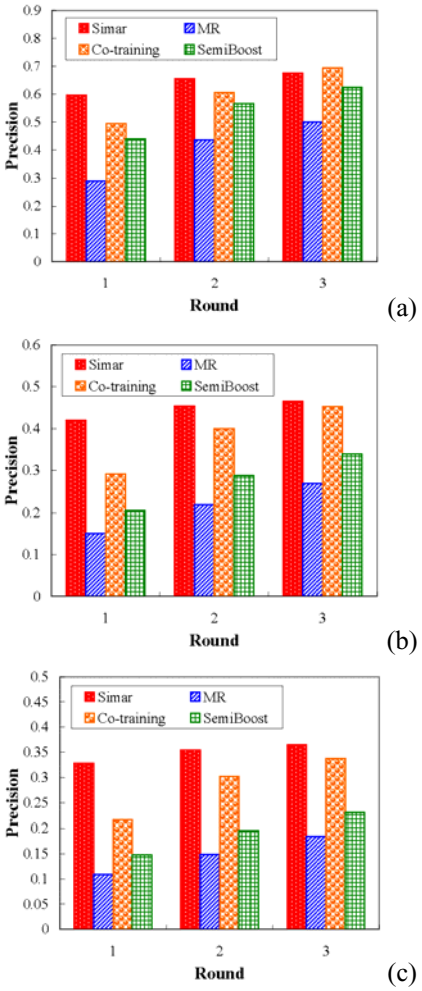

Fig. 3. Precisions of the proposed method compared with some existing methods at the Top (a) 20, (b) 60 and (c) 100 retrievals.

unlabeled images are exploited in the boosting procedure. The SVM is implemented using LIBSVM toolbox. Furthermore, in order to study whether the elastic *k*NN graph is useful, a degenerated variant of Simar, termed SimarDeg, is evaluated in the comparison. (4) **SimarDeg** is almost the same as Simar except that the former use the fixed *k*NN graph ($k = \text{floor}(\log n)$), instead of the elastic *k*NN graph, to calculate the Laplacian matrix.

3.3 Performance Evaluation

To evaluate the average performance, we conducted every experiment on a set of 200 random queries sampled from our image dataset. At the beginning of retrieval, the database images are ranked according to their Euclidean distances to the query image and top ten images are labeled as the initially labeled training data. Then, various methods are then applied to rerank the database images. For each compared method, after obtaining a query, several rounds of feedback were performed, and in each round the user labeled ten images as the feedback.

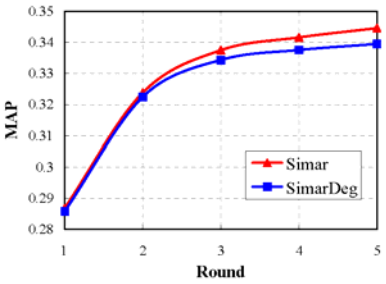

Fig. 4. MAP of the proposed method compared with its degenerated variant.

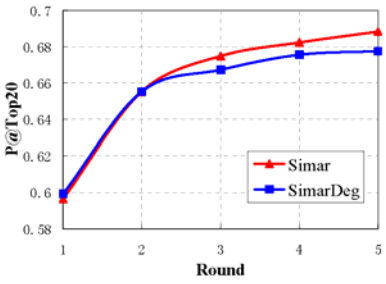

Fig. 5. P@Top 20 of the proposed method compared with its degenerated variant.

Table 1. MAPs of the four compared methods

	Simar	**MR**	**Co-training**	**SemiBoost**
Round 1	**0.287**	0.062	0.156	0.081
Round 2	**0.324**	0.113	0.25	0.177
Round 3	**0.338**	0.152	0.303	0.234

At first, the performance of Simar, MR, Co-training and SemiBoost are compared. The PR-graph at the 1st, 2nd, and 3rd round of feedback are shown in Figure 2, and the corresponding MAP statistic is tabulated in Table 1, where the best performance has been boldfaced. The precision curve at top 20, top 60, and top 100 retrieval results are presented in Figure 3. Several observations can be drawn from the experimental results. First, by comparing the two MR approaches, the performance of Simar is much better

than conventional MR. Note that the main difference between them is that Simair calculates the Laplacian matrix using an adaptive scale parameter while conventional MR does this using a fixed scale parameter, which verifies the usefulness of our local scaling solution. Furthermore, in most cases, Simar outperforms Co-training and SemiBoost, especially at the first round of feedback, which is meaningful to the real world applications because it is not practical to require the user to provide many rounds of feedback and therefore the retrieval performance at the 1st round of feedback is the most important. Finally, it is impressive that at all rounds of feedback, the MAP of Simar is always the best. That means the MR approach is more effective than other semi-supervised ranking methods when the parameters are tuned appropriately.

In order to study whether the elastic *k*NN graph employed in our approach is beneficial or not, Simar is compared with its degenerated variant SimarDeg. Figure 4 and Figure 5 print the MAP and the P@Top20 of the two algorithms at 1st to 5th round of feedback, respectively. As can been seen, the performance of Simar and SimarDeg are close to each other at the first two rounds of feedback, and then Simar growingly outperforms SimarDeg with the increase of the rounds of feedback. It is conjectured that the number of labeled images is small at the first two rounds of feedback, the probability of exploiting the "unreliable" unlabeled images (the nearby neighbors of the labeled images) would be low, and thus the impact of the elastic *k*NN graph is trivial. By gradually adding the user's feedbacks, the elastic *k*NN graph is increasingly helpful to Simar.

4 Conclusions

In this paper, we presented a novel MR approach for relevance feedback in CBIR, which addressed the two main drawbacks of regular MR algorithm. In particular, we employed an elastic *k*NN graph in MR to reduce the risk of exploiting "unreliable" unlabeled data, and developed a local scaling solution to facilitate the setting of the scale parameter used for calculating Laplacian matrix. We conducted extensive experiments to evaluate the performance of our techniques for relevance feedback in CBIR, from which the promising results showed the advantages of the proposed approach in comparison to several existing methods. In the future work, we will take more visual features into consideration and evaluate our method on other databases.

Acknowledgements. This work was supported in part by the Fundamental Research Funds for the Central Universities (#2012JBM038 and #2012JBM035), the Natural Science Foundation of China (#61170232, #60973067 and #61175053), and the State Key Laboratory Research Funds #RS2012K011.

References

1. Chapelle, O., Scholkope, B., Zien, A.: Semisupervised Learning. MIT Press, Cambridge (2006)
2. Datta, R., Joshi, D., Li, J., Wang, J.Z.: Image Retrieval: Ideas, Influences, and Trends of The New Age. ACM Comput. Surv. 40(2), 5:1-5:60 (2008)

3. He, J., Li, M., Zhang, H., Tong, H., Zhang, C.: Manifold-Ranking Based Image Retrieval. In: Proc. ACM Int. Conf. Multimedia, MM (2004)
4. Hoi, S.C.H., Jin, R., Zhu, J., Lyu, M.R.: Semi-Supervised SVM Batch Mode Active Learning for Image Retrieval. In: Proc. IEEE Conf. Computer Vision and Pattern Recognition, CVPR (2008)
5. Li, Y.F., Zhou, Z.H.: Towards Making Unlabeled Data Never Hurt. In: Proc. Int. Conf. Machine Learning, ICML (2011)
6. Luxberg, U.: A Tutorial on Spectral Clustering. Statistics and Computing 17(4), 395–416 (2007)
7. Wang, B., Pan, F., Hu, K.M., Paul, J.C.: Manifold-Ranking Based Retrieval using k-Regular Nearest Neighbor Graph. Pattern Recognition 45(4), 1569–1577 (2012)
8. Wu, J., Lin, Z., Lu, M.: Asymmetric Semi-Supervised Boosting for SVM Active Learning in CBIR. In: Proc. ACM Int. Conf. Image and Video Retrieval, CIVR (2010)
9. Wu, J., Lu, M., Wang, C.: Collaborative Learning between Visual Content and Hidden Semantic for Image Retrieval. In: Proc. IEEE Int. Conf. Data Mining, ICDM (2010)
10. Xu, B., Bu, J., Chen, C., Cai, D., He, X., Liu, W., Luo, J.: Efficient Manifold Ranking for Image Retrieval. In: Proc. ACM Int. Conf. Research and Development in Information Retrieval, SIGIR (2011)
11. Yang, Y., Nie, F., Xu, D., Luo, J., Zhuang, Y., Pan, Y.: A Multimedia Retrieval Framework Based on Semi-Supervised Ranking and Relevance Feedback. IEEE Trans. Pattern Analysis and Machine Intelligence 34(4), 723–742 (2012)
12. Zelnik-Manor, L., Perona, P.: Self-Tuning Spectral Clustering. Adv. Neural. Inf. Process. Syst. (NIPS) 2 (2004)
13. Zhang, L., Wang, L., Lin, W.: Semisupervised Biased Maximum Margin Analysis for Interactive Image Retrieval. IEEE Trans. Image Processing 21(4), 2294–2308 (2012)
14. Zhou, X.S., Huang, T.S.: Relevance Feedback in Image Retrieval: A Comprehensive Review. Multimedia Syst. 8(6), 536–544 (2003)
15. Zhou, Z.H., Chen, K.J., Dai, H.B.: Enhancing Relevance Feedback in Image Retrieval using Unlabeled Data. ACM Transactions on Information Systems 24(2), 219–244 (2006)