

RankCompete: Simultaneous Ranking and Clustering of Web Photos

Liangliang Cao⁺, Andrey Del Pozo⁺, Xin Jin[‡], Jiebo Luo^{*}, Jiawei Han[‡], Thomas S. Huang⁺

⁺ Beckman Institute and Coordinated Science Lab, UIUC. {cao4,delpozo2,huang}@ifp.uiuc.edu

[‡] Department of Computer Science, UIUC. {xinjin3,hanj}@cs.uiuc.edu

^{*} Kodak Research Laboratories. jiebo.luo@kodak.com

ABSTRACT

With the explosive growth of digital cameras and online media, it has become crucial to design efficient methods that help users browse and search large image collections. The recent VisualRank algorithm [4] employs visual similarity to represent the link structure in a graph so that the classical PageRank algorithm can be applied to select the most relevant images. However, measuring visual similarity is difficult when there exist diversified semantics in the image collection, and the results from VisualRank cannot supply good visual summarization with diversity. This paper proposes to rank the images in a structural fashion, which aims to discover the diverse structure embedded in photo collections, and rank the images according to their similarity among local neighborhoods instead of across the entire photo collection. We design a novel algorithm named RankCompete, which generalizes the PageRank algorithm for the task of simultaneous ranking and clustering. The experimental results show that RankCompete outperforms VisualRank and provides an efficient but effective tool for organizing web photos.

Categories and Subject Descriptors

H.3 [Information Storage and Retrieval]:

General Terms

Experiment and Algorithm

Keywords

PageRank, Image Ranking, Image Summarization

1. INTRODUCTION

The popularity of digital cameras, camera-phones and high capacity memory cards has led to an explosion of digital images on the web, especially in online photo sharing communities. The common approach used in web image search is based on textual information (e.g., image file name and surrounding text). However, such an approach cannot handle images where the related textual information is missing or inconsistent with the visual content. The current image search techniques are also ineffective for browsing the photo albums in the online sharing communities (e.g., Flickr, Facebook). When reviewing photos from friends or from the community, users often have to click the images page by page, with many irrelevant images or duplicates.

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

WWW 2010, April 26–30, 2010, Raleigh, North Carolina, USA.
ACM 978-1-60558-799-8/10/04.

Jing and Baluja proposed VisualRank [4], which identifies the authority of images on a similarity graph based on visual similarity. VisualRank is motivated by the recent success of the PageRank algorithm [6]. Unlike the classical PageRank and HITS that build an adjacent matrix based on the hyperlinks between web documents, VisualRank treats images as web documents and their similarities as “visual” links. Such visual links will not suffer from malicious hyperlinks from web spammers. Based on such visual links, a PageRank score is estimated for each image, based on which images are ordered for retrieval systems.

Despite its success in product image search, VisualRank is still not perfect for tasks such as browsing and organizing large collection of images in several aspects. First, VisualRank does not consider the visual diversity of the retrieval results and cannot handle the ambiguity of user queries. If two images look similar, they will share similar hyperlinks with others. As a result, if an image is ranked high by VisualRank, its duplicate or near duplicate images will also be ranked high, giving the user a subset of images with rather limited visual diversity. Moreover, it is difficult for users to explore the searching results provided by VisualRank due to the lack of structure. A user often have to scroll or click through many pages of image results to find the object of interest.

We propose a new algorithm named *RankCompete* that generalizes the PageRank algorithm for the task of simultaneous ranking and clustering. Our contribution is a scheme that performs the tasks of ranking and clustering in a mutually enhancing fashion: (1) the ranking results make more sense when comparing only the images with similar semantics. (2) the clustering results can also be improved using ranking information since relevant documents are more similar to each other than the irrelevant documents.

2. ALGORITHM

Following the work in VisualRank [4], this work models visual similarity using the matched SIFT features [5] from a pair of images, which are 128-dimensional vectors describing the image gradient orientation histograms for local patches. Given two images, VisualRank [4] computes their similarity as the number of local features shared between them. However, computing pairwise similarity is expensive. To accelerate the computation, we quantize SIFT descriptors into salient visual words and compute the image similarity as the number of shared visual words between the two images. Thus the images form a graph where the column-normalized similarity matrix \mathbf{S} corresponds to an adjacency matrix.

To make the presentation simple, we first consider the

Algorithm 1 : The RankCompete Algorithm

- 1: Initialize \mathbf{p}_1 and \mathbf{p}_2 satisfying $\sum_u p_k(u) = 1$ for $k = 1, 2$.
 - 2: Update until convergence
 - 3: do *ranking* step using (1)
 - 4: do *competing* step using (2) and (3).
 - 5: Obtain two clusters D_1 and D_2 with the corresponding normalization factors ρ_1 and ρ_2 .
 - 6: if $\rho_k > Threshold$, $k = \{1, 2\}$
 - 7: do RankCompete on subgraph D_k
 - 8: Output all clusters with corresponding ranking scores.
-

two-class clustering and introduce two groups of two random walks in the similarity graph. We define two measures \mathbf{p}_1 and \mathbf{p}_2 to represent the probability for each random walk to visit each image node u with the constraint $\sum_u p_k(u) = 1, k = \{1, 2\}$. We also assume that one node can only host one walk, so that \mathbf{p}_1 and \mathbf{p}_2 will compete with each other on every node.

The RankCompete algorithm can be viewed as a two-step process. In the *ranking* step, we update \mathbf{p}_k in a way similar to PageRank:

$$\bar{\mathbf{p}}_k = \mathbf{S}^* \mathbf{p}_k \quad (1)$$

Then in the *competing* step, two random walks $\bar{\mathbf{p}}_1$ and $\bar{\mathbf{p}}_2$ will compete on each node by

$$p_k(u) = \begin{cases} \bar{p}_k(u), & \text{if } \bar{p}_k(u) = \max_{k'} \bar{p}_{k'}(u) \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

To guarantee that \mathbf{p}_1 and \mathbf{p}_2 satisfy the constraint of random walk, we use a normalization process of

$$p_k(u) = \bar{p}_k(u) / \rho_k \quad (3)$$

where $\rho_k = \sum_u \bar{p}_k(u)$ is called a normalization factor.

RankCompete initializes \mathbf{p}_1 and \mathbf{p}_2 by two random vectors or from two selected images, and then iteratively update them based on Eqs. (1)(2) (3) until \mathbf{p}_1 and \mathbf{p}_2 no longer change or iterations exceed a threshold (50 in our experiment)¹. Based on $\mathbf{p}_1, \mathbf{p}_2$, we can obtain two clusters D_1 and D_2 . The ranking score of each node u is also obtained simultaneously as $p_1(u)$ or $p_2(u)$. Note that RankCompete can be easily generalized to multiple-class clustering by performing hierarchical clustering in a top-down manner. Algorithm 2 outlines the procedure of general random compete algorithm. More details can be found in [2].

3. EXPERIMENTS

We first use an example to show how our RankCompete algorithm produces a structured view of the search results. We download images from Flickr using the query of “Raleigh” and apply the RankCompete algorithm to these images. The top three clusters in Figure 1 show that our algorithm can effectively summarize the diversified images and help find the most relevant images more effectively.

To further evaluate our algorithm, we employ two public datasets to compare the performance of VisualRank and RankCompete. We first employ the dataset from ImageCLEF’08 [1], which provides labels for 39 topics and each topic is composed with multiple clusters (2 ~ 23). Our use of the dataset is different from the original intention since we are not working on the annotation such as titles, descriptions, and locations. In contrast, we fuse the ground

¹It can be proven that the converged \mathbf{p}_k coincides with the PageRank score in the sub-graph within the cluster [2].



Figure 1: Top retrieved images returned by RankCompete for organizing “Raleigh” photos in Flickr.

Table 1: Comparing the accuracy on CLEF dataset.

Algorithm	Average Precision	Average S-Recall
VisualRank	0.688	0.578
Rankcompete	0.733	0.612

truth images of each topic with 40% randomly selected other images. To evaluate our algorithm, we compare the top 20 images returned by VisualRank and RankCompete algorithms for each topic. The second dataset comes from the WIDE Data set [3], which is the largest labeled visual dataset available at present. For each topic, we collect 1000 images with ground truth concepts or with the same tags. For each topic, we evaluate the top 50 images returned by VisualRank and RankCompete algorithms. To evaluate the performance, we employ two measures: precision (the percentage of relevant images in the retrieved list) and S-recall (the percentage of different clusters found by the ranking algorithms). Table 1 compares the average precision and S-recalls across 39 topics on the ImageCLEF dataset. Our method can improve both the accuracy and diversity of the retrieval results. Since there are no cluster tags associated with each topic in NUS WIDE, we can only compare the precision but not S-recall. Our RankCompete algorithm again outperforms VisualRank (precision 0.894 vs. 0.872).

4. CONCLUSION

We present a new algorithm named RankCompete, which is a generalization of the PageRank algorithm to the scenario of simultaneous ranking and clustering. The results show that RankCompete works well for the task of simultaneous ranking and clustering of web photos, and outperforms VisualRank on two challenging datasets.

5. ACKNOWLEDGEMENTS

This research was sponsored in part by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes. Cao is also supported by UIUC CSE fellowship.

6. REFERENCES

- [1] ImageCLEF 2008. <http://www.imageclef.org/2008/photo>.
- [2] L. Cao, *et al.* RankCompete for simultaneous ranking and clustering. *IFP group Technical Report, UIUC*, 2010.
- [3] T.-S. Chua, *et al.* NUS-WIDE: A real-world web image database from national university of singapore. *CIVR*, 2009.
- [4] Y. Jing and S. Baluja. Pagerank for product image search. *WWW*, 2008.
- [5] D. Lowe. Object recognition from local scale-invariant features. *ICCV*, 1999.
- [6] L. Page *et al.*, The pagerank citation ranking: Bringing order to the web. *Stanford Digital Libraries Working Paper*, 1998.