Information Sciences 265 (2014) 91–104

Contents lists available at ScienceDirect

Information Sciences

journal homepage: www.elsevier.com/locate/ins

A scalable re-ranking method for content-based image retrieval

Daniel Carlos Guimarães Pedronette [a,*], Jurandy Almeida [b,c], Ricardo da S. Torres [c]

[a] Department of Statistics, Applied Mathematics and Computing, Universidade Estadual Paulista (UNESP), Av. 24-A, 1515, Rio Claro, SP 13506-900, Brazil
[b] Institute of Science and Technology, Federal University of São Paulo (UNIFESP), Rua Talim, 330, São José dos Campos, SP 12231-280, Brazil
[c] RECOD Lab, Institute of Computing (IC), University of Campinas (UNICAMP), Av. Albert Einstein, 1251, Campinas, SP 13083-852, Brazil

ARTICLE INFO

Article history:
Received 5 November 2012
Received in revised form 29 May 2013
Accepted 30 December 2013
Available online 7 January 2014

Keywords:
Content-based image retrieval
Re-ranking methods
Indexing structures

ABSTRACT

Content-based Image Retrieval (CBIR) systems consider only a pairwise analysis, *i.e.*, they measure the similarity between pairs of images, ignoring the rich information encoded in the relations among several images. However, the user perception usually considers the query specification and responses in a given *context*. In this scenario, re-ranking methods have been proposed to exploit the *contextual information* and, hence, improve the effectiveness of CBIR systems. Besides the *effectiveness*, the usefulness of those systems in real-world applications also depends on the *efficiency* and *scalability* of the retrieval process, imposing a great challenge to the re-ranking approaches, once they usually require the computation of distances among all the images of a given collection. In this paper, we present a novel approach for the re-ranking problem. It relies on the similarity of top-*k* lists produced by efficient indexing structures, instead of using distance information from the entire collection. Extensive experiments were conducted on a large image collection, using several indexing structures. Results from a rigorous experimental protocol show that the proposed method can obtain significant effectiveness gains (up to 12.19% better) and, at the same time, improve considerably the efficiency (up to 73.11% faster). In addition, our technique scales up very well, which makes it suitable for large collections.

1. Introduction

Advances in multimedia technologies for creating and sharing digital contents have triggered an exponential increase of image collections. In order to deal with these collections, it is necessary to develop methods for efficiently indexing and retrieving these data. Traditional search approaches based on image metadata and keywords can be unfeasible for large collections, since manual annotation is prohibitively expensive. In this scenario, Content-Based Image Retrieval (CBIR) systems [10,35] have emerged as an alternative to overcome those limitations by taking into account the content of the images for supporting retrieval tasks.

A common task for CBIR systems is to retrieve the most similar images to a query pattern (*e.g.*, query image) defined by users. In general, the output provided is a ranked list, where the images are disposed in decreasing order of similarity, according to a visual property, such as shape, color, and texture. In this scenario, accurately ranking the collection images is of great relevance. Existing systems often consider only pairwise analysis, measuring the similarity between pairs of images and ignoring the relevant information encoded in the relations among several images. The user perception, on the other hand, considers the query specification and responses in a given *context*.

* Corresponding author at: Department of Statistics, Applied Mathematics and Computing, Universidade Estadual Paulista (UNESP), Av. 24-A, 1515, Rio Claro, SP 13506-900, Brazil. Tel.: +55 (19) 3526 9082.

E-mail addresses: daniel@rc.unesp.br, dcarlos@ic.unicamp.br (D.C. Guimarães Pedronette), jurandy.almeida@unifesp.br, jurandy.almeida@ic.unicamp.br (J. Almeida), rtorres@ic.unicamp.br (R. da S. Torres).

0020-0255/$ - see front matter
http://dx.doi.org/10.1016/j.ins.2013.12.030

Motivated by these limitations, many supervised learning approaches have been proposed. Relevance Feedback [13,37,34,41,53] and Active Re-Ranking [42] methods, for example, were incorporated into CBIR systems with the aim of exploiting interactions for learning user needs. Basically, the image retrieval process with relevance feedback is comprised of four steps: (i) showing a small number of retrieved images to the user; (ii) user indication of relevant and non-relevant images; (iii) learning the user needs by taking into account his/her feedbacks; and (iv) and selecting a new set of images to be shown. This procedure is repeated until a satisfactory result is reached. Although very effective, these approaches require a lot of human efforts for obtaining enough training data, which can be infeasible for some real-world systems.

Aiming at overcoming these problems, efforts were put on unsupervised approaches. Recently, various approaches [19,29,49,50] have been proposed to improve the *effectiveness* of retrieval tasks by taking into account the relationships among all dataset objects. In other words, research efforts have been focused on post-processing the similarity (or distance) scores, by using the contextual information available in relationships among images of a given collection. The goal of those methods is somehow mimic the human behavior on judging the similarity among objects by considering specific contexts. The key advantage of re-ranking approaches consists in the fact that they require no user intervention, training or labeled data, operating on an absolutely unsupervised way.

However, the usefulness of re-ranking approaches for CBIR systems depends not only on the *effectiveness*, but also on the *efficiency* and *scalability*. While the effectiveness is related to the quality of retrieved images, the efficiency refers to the time spent to obtain the results. Scalability considers the system capability of handling growing image collections. Although the effectiveness has been the focus of various recently works [27,46,49], dealing with those three requirements at the same time is essential in real-world applications. Aiming at computing the relationship among images, re-ranking algorithms often consider all the distances among images of a given dataset, which represent a large computational effort (typically, between $O(N^2)$ and $O(N^3)$), hindering its use in searching services that deal with real-world image collections.

On the other hand, significant research efforts have been spent trying to improve the performance in processing similarity queries. Most of existing indexes employed to accelerate data retrieval are constructed by partitioning a set of objects using distance-based criteria. Those approaches avoid the computation of distances among all the images of a given collection.

In this paper, we aim at combining the potential of effectiveness gains obtained by re-ranking approaches with the power of indexing structures in processing similarity queries efficiently. Here, we present a novel approach for the re-ranking problem that relies on ranked lists produced by efficient indexing structures. The ranked lists used by the proposed method contain only a subset of the most similar images, avoiding the computation, storage, and processing of distance information from the entire collection. The main contribution of the proposed index-based re-ranking approach consists in its capacity of combining effectiveness and efficiency features, making it suitable for large collections. The proposed re-ranking method requires very low computational efforts, presenting an asymptotic complexity of only $O(N)$. On the other hand, the effectiveness gains are comparable to state-of-the-art approaches.

We carried out extensive experiments on a large image collection, considering several indexing structures. The reported results demonstrate that the proposed method obtains significant effectiveness gains (up to 12.19% better) and, at the same time, improves considerably the efficiency (up to 73.11% faster). Moreover, our technique scales up very well, which makes it suitable for large collections. We also evaluated the proposed method in comparison with several other state-of-the-art approaches considering a common shape dataset. Experimental results demonstrate that the proposed method yields effectiveness results comparable to post-processing algorithms recently proposed in the literature.

The remainder of this paper is organized as follows. Section 2 describes related work. Section 3 introduces the re-ranking problem based on ranked lists. Section 4 presents our image re-ranking approach. Section 5 discusses indexing structures used to produce ranked lists. Section 6 reports the results of our experiments. Finally, we offer our conclusions and directions for future work in Section 7.

2. Related work

This section presents related work. Section 2.1 overviews image re-ranking approaches, while existing indexes structures are discussed in Section 2.2.

2.1. Image re-ranking

In recent years, several CBIR approaches [19,27,29,46,49,51] have been proposed aiming at improving the effectiveness of retrieval tasks by replacing pairwise similarities by more global affinities that also consider the relation among the database objects [51].

Although using a very diverse taxonomy (re-ranking [24,29], graph transduction [49], diffusion process [50], affinity learning [51], contextual similarity/dissimilarity measures [30,46]), these *post-processing* methods have in common the goal of improving the effectiveness of retrieval tasks by exploiting considering relationships among dataset objects on an unsupervised way, requiring no training data.

Graph-based methods are used by several approaches [19,46,49]. In [49], a *graph-based transductive learning algorithm* is proposed for shape retrieval tasks. It learns a better metric through a graph transduction by propagating the model through existing shapes, in a similar manner to the computation of geodesics in a dataset manifold. Another approach based on