Computers & Graphics 38 (2014) 248–254

Contents lists available at ScienceDirect

Computers & Graphics

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

CAD/Graphics 2013

Interactive shape co-segmentation via label propagation

Zizhao Wu [a], Ruyang Shou [a], Yunhai Wang [b], Xinguo Liu [a,*]

[a] State Key Lab of CAD & CG, Zhejiang University, China
[b] Shenzhen VisuCA Key Lab/SIAT, China

ARTICLE INFO

Article history:
Received 5 August 2013
Received in revised form
8 November 2013
Accepted 8 November 2013
Available online 20 November 2013

Keywords:
Shape co-segmentation
Interactive segmentation
Label propagation

ABSTRACT

In this paper, we present an interactive approach for shape co-segmentation via label propagation. Our intuitive approach is able to produce error-free results and is very effective at handling out-of-sample data. Specifically, we start by over-segmenting a set of shapes into primitive patches. Then, we allow the users to assign labels to some patches and propagate the label information from these patches to the unlabeled ones. We iterate the last two steps until the error-free consistent segmentations are obtained. Additionally, we provide an inductive extension of our framework, which effectively addresses the out-of-sample data. The experimental results demonstrate the effectiveness of our approach.

1. Introduction

In recent years, there have been increasing interests in shape co-analysis, i.e., simultaneously analyzing a set of shapes. One of the most fundamental problems in this field is co-segmentation. Different from the traditional segmentation tools which treat shapes individually, co-segmentation approaches process shapes from an input set in a batch, and generate segmentations carrying consistent semantics across the shapes. The consistent segmentation has demonstrated great utility in modeling [1,2], shape retrieval [3,4], texturing [5], etc.

Previous attempts for solving this problem can be classified into three categories as supervised, semi-supervised and unsupervised. The supervised ones [5,6] take advantages of manually labeled training sets to generate consistent segmentation results. However, the accuracy of the results relies on the training sets, and not surprisingly, the training process is tedious and time consuming. The unsupervised methods [7,8] generally build their approaches on the patch-level. These methods have superior performance, but the results hinge upon the in-sample data.

Recently, Wang et al. [9] presented a semi-supervised learning method with the aid of constrained clustering, where the user can actively assist in the co-segmentation process by assigning pairwise constraints like must-link and cannot-link. This approach can generate error-free results with a sparse set of constraints. However, as some authors [10] mentioned, pairwise constraints are not expressive to the users. In addition, their approach is a transductive algorithm which does not handle with the out-of-sample data, i.e., given a new datum, it needs performing the algorithm over the whole pipeline, which is ineffective.

In this paper, we address the above issues by introducing an interactive shape co-segmentation method. Our motivation drives from label propagation which propagates labels through the dataset along high density areas defined by unlabeled data. Our method allows the users to participate in the co-segmentation procedure, and is built upon the patch-level, which guarantees the high speed. Specifically, starting from over-segmenting a set of shapes into primitive patches, we allow the users to assign labels to some patches, and then propagate the labels from these patches to the unlabeled ones. We iterate the last two steps until the error-free consistent segmentations are obtained.

In addition, as mentioned previously, when building their approaches on the patch-level, state-of-the-art methods [7–9] are effective in dealing with in-sample dataset in their respective problem domains, but all these methods have not explored the out-of-sample data. We investigate the out-of-sample issue by introducing an inductive extension of our pipeline, where the new datum can be labeled effectively.

Comparing with the state-of-the-art algorithms, our approach is featured as follows:

- *Intuitive*: We provide an intuitive user interface. For the users to directly assign labels is more expressive than the pairwise constraints.
- *Inductive*: We introduce an inductive extension of our algorithm to deal with out-of-sample data.
- *Error-free*: We can achieve error-free results depending on the input dataset and the labels given by the users.
- *Efficient*: Our approach is graph-based, but requires no extra eigen-decomposition, which is different from the unsupervised methods [7,8].

* Corresponding author. Tel.: +86 571 88206681.
E-mail address: xgliu@cad.zju.edu.cn (X. Liu).

0097-8493/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cag.2013.11.009

The remainder of this paper is organized as follows. We review the related work in Section 2. We present the details of the proposed algorithm in Section 3. We show some experimental results on benchmark datasets in Section 4, followed by conclusions and future work in Section 5.

2. Related work

In this section, we provide a brief review of the existing work on shape co-segmentation, interactive segmentation and label propagation.

Shape co-segmentation: Shape co-segmentation refers to simultaneously segmenting a set of shapes into meaningful parts and building their correspondence. The existing co-segmentation methods can be classified into three categories: the unsupervised, the supervised and the semi-supervised.

In the unsupervised setting, the early work reported by Golovinskiy and Funkhouser [11] builds reliable correspondences across segments of shapes using rigid shape alignment. However, their approach cannot handle shapes with large variations. Xu et al. [12] factor out the scale variation in the shape segments by clustering the shapes into different styles, depending on the scales of the shape parts. Still, their approaches are limited to the shapes that can be properly aligned.

To overcome this limitation, Huang et al. [13] introduce an optimization strategy for simultaneously optimizing the saliency of each segmentation as well as consistency between segmentations. However, due to the computational complexity, this approach does not scale well for large datasets. Sidi et al. [7] present a descriptor-based method that employs multiple feature descriptors to measure the similarities of the segments and poses co-segmentation as a clustering problem in a concatenated descriptor space. Because the descriptors are independent of the pose and location of the shapes, this method can handle shapes with rich variations in part composition and geometry. Instead of concatenating the different feature descriptors into one vector, Hu et al. [8] propose a feature fusion method to co-segment a set of shapes via subspace clustering. However, these unsupervised techniques hinge upon the in-sample data.

Kalogerakis et al. [5] present a supervised learning method to simultaneously segment and label shapes. Their approach needs prior knowledge learned from the training dataset, and has demonstrated a labeling high accuracy on a broad class of shapes. van Kaick et al. [6] optimize the previous method by incorporating the prior knowledge to train a classifier. However, the above supervised methods require a substantial number of manually labeled training shapes, and the training set has a large impact on the segmentation performance.

Very recently, Wang et al. [9] propose a semi-supervised method where the user can actively assist in the learning process by interactively providing inputs. The input consists of a sparse set of pairwise constraints, which are marked as must-link and cannot-link constraints. The authors show that a sparse set of constraints can quickly converge toward an error-free result. However, the pairwise constraints are not clearly expressed to the users. In addition, their approach is a transductive algorithm that is ineffective at handling out-of-sample data.

Interactive segmentation: Interactive shape segmentation approaches are simple and intuitively help users express their intentions. Consequently, they have received significant attention [14].

Many interactive techniques have been proposed. Some of them require the user to specify a few points on the desired cutting contour and then employ the geometric snake [15], scissoring [16,17], graph cut [18] or some other method [19] to find the final cutting boundaries. These methods are called boundary-based approaches. In the last few years, a series of region-based approaches [20–22] have been proposed, which take regional information as the input and require a much smaller amount of user effort to complete the labeling process for all of the unlabeled faces of a shape.

In this paper, rather than segmenting an individual shape, we present an interactive region-based technique to simultaneously segment a set of shapes in a consistent manner.

Label propagation: Label propagation was first introduced by Zhu and Ghahramani [23]. This technique propagates the labels through dense unlabeled regions and locates data with properties that are similar to those of the labeled data. Their approach is graph-based, which can be constructed straightforwardly by computing pairwise similarities among all of the data. Due to its simplicity and robustness, it has been used in processes such as patch labeling [24], image segmentation [25], and image annotation [26].

Some authors [27,28] have tried to optimize the original label propagation. Among them, Wang and Zhang [29] propose approximating the graph with a set of overlapped linear neighborhood patches (LNPs) and computing the edge weights in each patch using the neighborhood linear projection. Our work is directly inspired by the LNP. We apply this algorithm to our interactive shape co-segmentation setting.

3. Algorithm

3.1. Overview

Define a set of shapes $S=\{s_1, s_2, \ldots, s_N\}$, where s_i represents the i-th shape and N is the total number of shapes. Our algorithm simultaneously produces segments of the set of shapes S and builds their correspondences across these segments.

The pipeline of our approach is illustrated in Fig. 2. First, the algorithm pre-processes the set of shapes by partitioning the dataset into primitive patches and building a graph that represents the geometric similarities across them. Then, the user interactively labels some patches, which are used as initial seeds that guide the iterative propagation to find labels for the others.

Our algorithm is an iterative approach. Each iteration includes two steps: user interaction and label propagation based on the user input. These steps repeat until satisfactory results are obtained. Additionally, we apply an extension to the pipeline to handle out-of-sample data.

We discuss the preprocessing step in the next section, the label propagation in Section 3.3, and the inductive extension in Section 3.4.

3.2. Preprocessing

In this step, we start by over-segmenting the input shapes, where normalized cuts [30] are employed to decompose each shape s_i into primitive patches. In our settings, the number of patches per shape is set to 30. Let $P=\{p_1, p_2, \ldots, p_M\}$ be the set of patches from all of the shapes; it is clear that $M=30\,N$. Fig. 1 gives an example of our over-segmentation results.

Our approach associates the representation of relations between the patches with graphs. We represent this graph in matrix form, i.e., by constructing an affinity matrix W whose entries $w_{i,j}$ carry the similarities of p_i and p_j. Thus, to measure the similarities among patches, we first choose five robust and discriminative shape descriptors to extract extrinsic geometric information about the patches; shapes can be informatively represented based on these data. These widely recognized descriptors are the Shape Diameter Function (SDF) [31], the Conformal Factor (CF) [32], the Shape Contexts (SCs) [33,34], the Average Geodesic Distance (AGD) [35], and the geodesic distance to the base of the shape (GB) [7]. The descriptors are all defined on the mesh faces, so no additional conversions are required to make them mutually representationally compatible. Then, to describe the patches