Computer Aided Geometric Design 43 (2016) 159–171

Contents lists available at ScienceDirect

Computer Aided Geometric Design

www.elsevier.com/locate/cagd

Volumetric shape contexts for mesh co-segmentation

Xuanmeng Xie, Jieqing Feng *

State Key Lab of CAD&CG, Zhejiang University, China

ARTICLE INFO

Article history:
Available online 12 February 2016

Keywords:
3D shape descriptor
Volumetric shape context
Mesh segmentation

ABSTRACT

In the field of mesh segmentation, co-segmentation techniques achieve state-of-the-art performance; however, the segmentation results rely on the shape descriptors used in the segmentation process. In this paper, we propose a novel type of descriptor called the "volumetric shape context" (VSC). For each triangle in the mesh, the VSC describes the distribution of the shape's volume relative to the center of the triangle. This descriptor is descriptive, robust, and invariant under rigid transformations, uniform scaling, mirror imaging and model degeneration. We compare the VSC with state-of-the-art descriptors in a supervised mesh segmentation framework, and the results show that the VSC is most frequently selected as the first descriptor and that combining the VSC with other descriptors improves the segmentation results, thereby demonstrating the descriptiveness of the VSC.

1. Introduction

Many mesh segmentation techniques have been proposed, and these techniques can be classified into two types: part segmentation and surface-patch segmentation (Shamir, 2008). In this paper, we discuss part segmentation, in which a mesh is segmented into semantic parts. For example, a human body model can be segmented into several such parts: the head, torso, arms and legs. In recent years, researchers have realized that segmenting a set of shapes as a whole (co-segmentation) can yield better performance than segmenting individual shapes or pairs of shapes. For most co-segmentation approaches, their performance relies heavily on the 3D shape descriptors that are used. In these approaches, the triangles in the mesh are clustered into several semantic parts according to the similarity of their descriptors.

There are many descriptors that are widely used for mesh segmentation, but most are designed for solving problems in other fields, such as object recognition, and few studies have focused on the design of specialized descriptors for mesh co-segmentation. Furthermore, no descriptor is effective for all categories of models, and combinations of different descriptors can yield better performance than only one. Thus, co-segmentation approaches will benefit from the proposal and adoption of new descriptors, especially descriptors that provide significantly different information from that offered by existing descriptors.

In this paper, we propose a 3D shape descriptor for mesh segmentation named the "volumetric shape context" (VSC), which describes the volume distribution of the model around the center of each mesh triangle. It is robust to noise because of its context-based nature, and it is invariant under rigid transformations, uniform scaling, mirror imaging and model degeneration. The proposed descriptor is inspired by context analysis, which has been proven to be robust and effective in measuring the similarity between 3D shapes (Osada et al., 2002) or semantic parts (Laga et al., 2013).

* Corresponding author.
E-mail address: jqfeng@cad.zju.edu.cn (J. Feng).

http://dx.doi.org/10.1016/j.cagd.2016.02.006
0167-8396/

The VSC was evaluated in the framework of Kalogerakis et al. (2010), which is a supervised co-segmentation method that achieves state-of-the-art segmentation and labeling results. In this framework, several shape descriptors are extracted to form a descriptor vector, and JointBoost classifiers (Torralba et al., 2007) automatically select the descriptors that are most useful for a particular segmentation task. Hence, this framework is suitable for comparing the descriptiveness of 3D shape descriptors. The experimental results demonstrate that the VSC is most frequently selected as the first descriptor by JointBoost. With regard to the effect of removing a single descriptor, the labeling accuracy decreases most significantly upon the removal of the VSC, thereby indicating that the VSC represents the largest contribution to the high labeling accuracy achieved in the experiments.

The main contribution of this paper is our proposal of a new descriptor for mesh co-segmentation, which is descriptive and significantly improves the labeling accuracy. A secondary contribution is our proposal of an automatic context-based method and an interactive method to estimate the auxiliary orientations of the models. Both methods produce consistent auxiliary orientation estimation results, thereby verifying the descriptiveness of the VSC.

2. Related work

In this section, we review several 3D descriptors that are widely used in mesh segmentation methods and several mesh co-segmentation methods that take advantage of 3D descriptors.

There are many descriptors that are widely used for mesh segmentation, such as the curvature (Gal and Cohen-Or, 2006), principal component analysis (PCA) singular values, the shape diameter function (Shapira et al., 2010), the volumetric shape image (Liu et al., 2009), the average geodesic distance (Hilaga et al., 2001; Zhang et al., 2005), the shape context (Belongie et al., 2002), the spin image (Johnson and Hebert, 1999), the light field descriptor (Ding-Yun et al., 2003), the conformal factors (Ben-Chen and Gotsman, 2008), and the geodesic distance to the base of the shape (Sidi et al., 2011). Of these, the shape context and the spin image are originally 2D descriptors, and they have been adapted for extension to the 3D mesh segmentation task. These descriptors have proven to be effective for achieving good segmentation results, but none of them describes the volume distribution information of the models.

The shape context is widely used in mesh segmentation approaches and exhibits high descriptiveness. It treats a 3D model as a shell model and evaluates the distribution of the surface area; thus, the volume distribution information is not considered. Moreover, it does not construct a local 3D reference frame on each triangle; instead, it divides the neighboring space in only two dimensions, yielding insufficiently detailed distribution statistics.

The 3D shape context (Frome et al., 2004) extends the concept of the shape context to 3D point clouds and offers a solution for constructing local 3D reference frames. In this solution, the spherical support region is regarded as a globe, and the direction of the North Pole is the normal estimated at the base point. The globe is equally divided in the azimuthal and elevational dimensions and is logarithmically divided in the radial dimension. The local reference frames that are constructed in this manner are not unique, and one degree of freedom remains (the globe can rotate around its own central axis); as a result, the descriptors based on this solution are not unique and, consequently, are not suitable for mesh co-segmentation.

Many different mesh segmentation approaches have been proposed. Several surveys can be found in Agathos et al. (2007), Attene et al. (2006), Shamir (2008), Theologou et al. (2015). In recent years, researchers have realized that high-level analyses of 3D models can yield better performance than can the use of low-level geometric cues alone and thus have proposed co-segmentation approaches. Co-segmentation approaches can be classified into three categories: supervised (Kalogerakis et al., 2010; Van Kaick et al., 2011), unsupervised (Hu et al., 2012; Huang et al., 2011; Luo et al., 2013; Meng et al., 2013; Shapira et al., 2008; Sidi et al., 2011; Wang et al., 2012; Wu et al., 2014, 2013; Xu et al., 2010) and semi-supervised (Lv et al., 2012). Supervised approaches can achieve the highest labeling accuracy but require training data consisting of a set of manually segmented and labeled models, and the training phase is time-consuming. Unsupervised approaches offer lower performance than do supervised methods, which is not surprising because they use less information, but they also do not require training data and are much faster. The semi-supervised approach is a trade-off between the supervised and unsupervised approaches in that it learns from both labeled and unlabeled models. Descriptors are used in the algorithms of almost all co-segmentation approaches in these 3 categories, and their performance relies heavily on the adopted descriptors. In fact, among all of the approaches mentioned above, there was only one early study (Shapira et al., 2008) in which descriptors were not adopted, and the approaches in which descriptors are applied exhibit better performance. Supervised and semi-supervised methods use descriptors to form descriptor vectors, whereas unsupervised methods use descriptors predominantly for patch clustering and over-segmentation. These approaches will all benefit from the proposal and adoption of novel descriptors.

The proposed VSC descriptor was evaluated in the framework of Kalogerakis et al. (2010), which is a representative supervised mesh segmentation method. In this framework, the labeling of parts of the mesh is regarded as a problem of optimizing a conditional random field (CRF) (Lafferty et al., 2001). Several descriptors are extracted to form the descriptor vectors for all mesh triangles, and these descriptor vectors are used to train the basic terms of the CRF using JointBoost classifiers (Torralba et al., 2007). Different types of descriptors can be used in this framework, and adding more descriptors typically improves the performance; therefore, this framework is suitable for evaluating the descriptiveness of different descriptors.