Computers & Graphics 74 (2018) 150–160

Contents lists available at ScienceDirect

Computers & Graphics

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

Special Section on SMI 2018

3D synthesis of man-made objects based on fine-grained parts

Diego Gonzalez, Oliver van Kaick*

School of Computer Science, Carleton University, Ottawa, ON, Canada

ARTICLE INFO

Article history:
Received 27 April 2018
Revised 15 May 2018
Accepted 15 May 2018
Available online 22 May 2018

Keywords:
Shape synthesis
Shape segmentation
Shape analysis

ABSTRACT

We present a novel approach for 3D shape synthesis from a collection of existing models. The main idea of our approach is to synthesize shapes by recombining fine-grained parts extracted from the existing models based purely on the objects' geometry. Thus, unlike most previous works, a key advantage of our method is that it does not require a semantic segmentation, nor part correspondences between the shapes of the input set. Our method uses a template shape to guide the synthesis. After extracting a set of fine-grained segments from the input dataset, we compute the similarity among the segments in the collection and segments of the template using shape descriptors. Next, we use the similarity estimates to select, from the set of fine-grained segments, compatible replacements for each part of the template. By sampling different segments for each part of the template, and by using different templates, our method can synthesize many distinct shapes that have a variety of local fine details. Additionally, we maintain the plausibility of the objects by preserving the general structure of the template. We show with several experiments performed on different datasets that our algorithm can be used for synthesizing a wide variety of man-made objects.

1. Introduction

There is an ongoing need for digital content in the form of 3D models in fields such as entertainment and product design. However, the manual creation of 3D models is often a difficult and time-consuming task. Hence, facilitating the creation of 3D models is a fundamental problem in computer graphics. A popular approach for the automatic creation of novel 3D objects is to reuse parts from existing shapes. Several works on shape synthesis propose to extract and recombine parts from a dataset of existing objects [1–6]. Most of these techniques make two important assumptions about the input models and their parts. First, the input objects should be segmented into coarse *semantic* parts, that is, parts with a meaningful semantic meaning, and possibly a specific functionality. Second, a *correspondence* or consistent *labeling* should exist among the different parts of the input shapes.

One shortcoming of these synthesis methods is that the generated shapes may display a limited range of variations in their fine geometrical details, since semantic parts are often large, coarse and exchanged as a whole [7,8]. Furthermore, although much progress has been made to obtain semantic segmentations of 3D shapes [9,10], computing highly accurate and consistent segmentations and part correspondences for large collections of shapes is still a challenging problem [11,12].

In our work, we propose to synthesize shapes by exchanging *fine-grained* parts among a set of input shapes. As shown in Fig. 1, the use of fine-grained parts allows us to generate shapes with large variation in fine geometric details, which is not achievable with methods that exchange semantic parts as a whole. Fig. 2 shows a comparison between a typical semantic segmentation of a table, and the fine-grained segments that we use for synthesis. The recombination of fine-grained segments is guided by the geometric similarity of the segments, enabling us to exchange segments that have the same overall geometry but which can possess variations in their fine details. The recombination guided by geometry allows us to even exchange segments between models from different families.

Moreover, we do not require a semantic segmentation nor a part correspondence as input, since the fine-grained parts can be obtained by analyzing the local geometry of the shapes, as done by traditional geometry-based segmentation methods [9]. Our main requirement on the segmentation is that the shape segments should possess roughly the same size. However, one problem arising from the elimination of part semantics is that we require an alternative mechanism to guide and constrain the synthesis of a new shape. Thus, we propose to use for guidance an input template, which can be a simple configuration of geometric proxies or an example shape. The template constrains the topology

* Corresponding author.
E-mail addresses: diego.gonzalez@carleton.ca (D. Gonzalez), oliver.vankaick@carleton.ca (O. van Kaick).

https://doi.org/10.1016/j.cag.2018.05.016

Fig. 1. Scene created with shapes synthesized by our method. Note how the three shapes exhibit a variety of fine details and a coherent structure.

Fig. 2. Comparison of two segmentations for the table model shown in the inset. Left: typical semantic segmentation. Right: segmentation into fine-grained segments used in our work.

of the synthesized shape, which helps to ensure a certain level of plausibility in the generated shapes, and provides a domain for formulating the synthesis of a shape as a graph assignment problem, as we discuss in Sections 3 and 5.

In summary, we present a pipeline for synthesizing 3D shapes using fine-grained segments extracted from an input set of shapes. Our method does not require a segmentation of the input models into semantic parts, nor a part correspondence, but mainly a guiding template and a segmentation of a collection where the segments are consistent in size, which is easier to obtain with automatic or semi-automatic methods. Specifically, for the results shown in this paper, we employ a semi-automatic segmentation method. Moreover, we demonstrate the effectiveness of our approach by presenting and analyzing a variety of results synthesized with our method. In addition, we explore variations of our pipeline that enable us to control different aspects of the shape generation.

2. Related work

In this section, we discuss the previous work most related to our method, i.e., shape synthesis and segmentation approaches.

Shape synthesis. Earlier approaches for synthesis employed statistical shape models to generate shapes, according to the variability learned from a collection. For example, Blanz and Vetter [13] developed a deformable model of 3D faces, while Allen et al. [14] introduced a statistical model of human bodies. These methods are quite general as they are applicable to any type of shape (biological, man-made), but require a compatible mesh with correspondences across the entire collection.

The seminal work by Funkhouser et al. [1] proposed a system that allows the user to browse a library of 3D models and compose new shapes by assembling together parts of the existing 3D objects. This *part reuse framework* is widely applicable to objects that can be decomposed into parts in a meaningful manner, such as man-made objects. Thus, much of the subsequent work on shape synthesis has been based on this idea.

One line of work has proposed interfaces that facilitate the extraction and reuse of shape parts. For example, Kraevoy et al. [15] use compatible segmentations of objects to enable part exchange, while Sharf et al. [16] and Takayama et al. [17] introduce interfaces that facilitate the selection of part or regions to be exchanged between shapes. Chaudhuri et al. [18] incorporate semantics into this framework with a learning approach, to suggest suitable part replacements while building a shape from existing parts. Recently, Jaiswal et al. [19] and Sung et al. [20] also developed suggestion mechanisms based on learning approaches which however do not require labeled parts.

Another line of work proposes methods that *automatically* synthesize shapes while requiring little to no user input. A few works are based on the idea of *blending* existing shapes together as a whole, such as the methods of Jain et al. [3] and Alhashim et al. [7]. Blending is advantageous in that it can generate a variety of objects from two input shapes, however, these methods require a semantic segmentation of the input models.

Moreover, much of the recent work on automatic shape synthesis focuses on learning statistics of part co-occurrence based on a semantic segmentation of the objects in a collection, and using this information to automatically generate objects. Kalogerakis et al. [4] and Huang et al. [2] introduce approaches that select parts to compose a shape based on the co-occurrence of semantic labels. Zheng et al. [21] exchange specific arrangements of parts to synthesize objects that satisfy certain functionalities, while Huang et al. [22] extend this approach to non-symmetric arrangements. Su et al. [8] exchange more complex substructures among shapes. Moreover, Xu et al. [6] use an evolution-based approach to synthesize shapes.

A few methods also focus on the specific problem of generating valid configurations of shape parts, i.e., defining the relative positioning and orientation of the parts that compose the shapes. Fish et al. [23] and Yumer and Kara [24] learn models from a collection of shapes that capture the probability of part configurations. Averkiou et al. [25] represent the shapes of a dataset as box-like templates which can be used for exploration of the set but also for synthesizing new objects via part deformation.

The main requirement of all of these blending and automatic synthesis methods is a semantic segmentation of the objects in the analyzed collections. Often, a semantic labeling of the segments is also required. Although there has been significant work in recent years for obtaining such segmentations automatically [10], obtaining an accurate segmentation of a large collection is still a challenging problem with much room for improvement to the segmentation accuracy [12]. In contrast, the objective of our method is to synthesize shapes by exchanging parts obtained with a less demanding type of segmentation, which can be computed mainly from geometric constraints without involving semantics. Moreover, the use of fine-grained segments in our work enables to exchange fine details between objects and increase the variability in the synthesized shapes, which is not possible when using semantic segments, as these are often larger and more coarse in relation to the entire shape.

Finally, a few methods also aim to extract parts that can be suitable for exchange, in a sense, treating the segmentation and part exchange problems together. Bokeloh et al. [26], and later Kalojanov et al. [27], discover a shape grammar that can be used to generate the input shape and, subsequently, novel objects. Liu et al. [5] combine the discovery of shape grammars with the analysis of substructures for shape synthesis. Although the parts extracted by these methods are not semantic in nature, the discovered parts tend to be large as the goal of these methods is to discover the smallest set of elementary parts. Thus, the exchange of finer details is also difficult to achieve.

Shape segmentation. There has been a significant amount of work in shape segmentation, where the goal is to decompose an input object into meaningful parts. Much of earlier work employed geometric criteria combined with clustering or region growing heuristics to partition a single input object into parts, as surveyed by Shamir [9] and more recently by Theologou et al. [10].