Computers & Graphics xxx (2017) xxx–xxx

Contents lists available at ScienceDirect

Computers & Graphics

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

Technical Section

Q1 Relation-based parametrization and exploration of shape collections

Q2 Kurt Leimer*, Lukas Gersthofer, Michael Wimmer, Przemyslaw Musialski

Q3 *TU Wien, Karlsplatz 13, Vienna 1040, Austria*

A R T I C L E I N F O

Article history:
Received 30 September 2016
Revised 29 June 2017
Accepted 3 July 2017
Available online xxx

Keywords:
3D database exploration
Model variability
Shape analysis
Shape collections

A B S T R A C T

With online repositories for 3D models like 3D Warehouse becoming more prevalent and growing ever larger, new possibilities have emerged for both experienced and inexperienced users. These large collections of shapes can provide inspiration for designers or make it possible to synthesize new shapes by combining different parts from already existing shapes, which can be both easy to learn and a fast way of creating new shapes. But exploring large shape collections or searching for particular kinds of shapes can be difficult and time-consuming tasks as well, especially considering that online repositories are often disorganized. In our work, we propose a relation-based way to parametrize shape collections, allowing the user to explore the entire set of shapes based on the variability of spatial arrangements between pairs of parts. The way in which shapes differ from each other is captured automatically, resulting in a small number of exploration parameters. Furthermore, a copy-and-paste system for parts allows the user to change the structure of a shape, making it possible to explore the entire collection from any initial shape.

1. Introduction

High-quality 3D geometric models are needed in many industries, such as entertainment, computer-aided design, urban planning, the fabrication of physical objects, as well as research areas like medical or scientific visualizations and simulations. Manually creating a 3D model not only requires a lot of artistic, but also technical skills, as traditional modeling software often contain hundreds of different tools that all take significant effort to learn how to be used properly.

However, with online repositories for 3D models like 3D Warehouse or Turbosquid becoming more prevalent and growing ever larger, new possibilities have emerged for both experienced and inexperienced users. The availability of large collections of shapes allows users to quickly populate virtual scenes with a multitude of different objects. Expert designers can draw inspiration from the large variety of different shapes and creating novel shapes can be made easier through shape synthesis, where parts of existing models are combined to create a new one, for example by interactive approaches like modeling-by-example [1] or structural blending [2], or even by algorithms that automatically generate novel shapes that in turn can also serve as additional inspiration for the user [3]. On the other hand, these online repositories are often disorganized, with thousands of shapes listed under the same category and no meaningful way to distinguish between differently looking shapes of the same family. This makes it a difficult task to look for specific shapes or explore the collection in a meaningful manner.

Shape retrieval is one proposed solution to this problem [4–7]. From a high-level point of view, given an existing shape or a sketch provided by the user, a number of shapes that are similar to the input are retrieved from the collection based on some distance measure. Such a method can be useful when the user already has a specific shape in mind, but is less suited for exploration where one might not even know what kind of shapes are available. A more general idea of how to do this would be to find some kind of ordering of the shapes so that similar shapes are close and dissimilar ones are far apart.

In this paper we address this problem, denoted as *shape exploration*. In order to provide such an ordering, we propose a relation-based way to parametrize shape collections, allowing the user to explore the entire set of shapes by controlling a small number of parameters. Given a co-segmented set of shapes from the same family, we look at pairs of adjacent parts and how their spatial arrangements vary in regards to other part pairs of the same category.

To make the exploration process more intuitive, we visualize the parameter changes in the shape viewer. The currently selected part is represented by its bounding box which is then transformed relatively to its adjacent part based on the current setting of the

* Corresponding author.
E-mail address: kleimer@cg.tuwien.ac.at (K. Leimer).

http://dx.doi.org/10.1016/j.cag.2017.07.001
0097-8493/

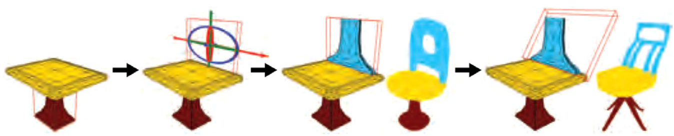

Fig. 1. Exploring the shape collection by changing the structure of an initial shape. Existing parts can be copied and assigned a different label. Such a new part can then be used as a proxy to explore the collection.

exploration parameters. This gives a rough first look at the kind of shapes that can be found near the current point in the exploration space and also provides an understanding of how the shapes change when moving in a specific direction of the exploration space.

Since our method only captures local variability, we also consider the possibility of not using just one such relation to differentiate between shapes, but multiple relations. As an example, one could adjust the spatial arrangements of a chair seat compared to both the legs and backrest to find more specific shapes. Finally, we also look at the problem of shapes containing different numbers of parts. This is something that is often ignored in existing approaches or dealt with by clustering shapes according to their part numbers and only allowing exploration within each such cluster. In our method, we provide a simple tool to allow the copying and pasting of existing parts to add parts not present in the current shape, making it possible to explore the entire shape collection from any initial shape. This can be seen in Fig. 1.

In summary, our contributions are as follows:

- A parametrization of shape collections based on relations between shape parts that automatically captures the way in which the shapes vary.
- A simple exploration process that allows browsing the collection based on local variability (spatial arrangements between a pair of adjacent parts) or searching for more specific shapes by considering multiple such relations at the same time.
- A visualization of the exploration parameters that can aid the user in understanding the effects of altering the parameters, making exploration more intuitive.
- A way to change the structure of the initial shape, allowing the user to find shapes with different structures and explore the entire collection from every initial shape.

Our approach can be separated into two stages. The first stage is the parametrization stage described in Section 3. In this stage, the collection of shapes is analyzed to find relations between shape parts that are useful in exploring the collection. This allows us to compute exploration parameters that can be used to explore the collection in the exploration stage which is described in Section 4. We also consider the possibility of using nonlinear methods described in Section 5.

We perform a number of experiments with our system, results of which are presented in Section 6. Furthermore, in Section 7 we conduct a user study to show that the ordering of shapes produced by our approach is not less intuitive than ordering by the individual spatial features like distance, angle or scale between parts. The discussion of our results, as well as limitations of our approach can be found in Section 8. Finally, we conclude this paper with a short summary in Section 9.

2. Related work

With online model repositories growing larger, a need for automatically organizing these huge collections of shapes has arisen. As such, the problem of how to organize and explore large collections of shapes has been a popular topic of research in recent years. In this section we give an overview of related work on the topic.

Template deformations: Ovsjanikov et al. [8] construct a template made out of boxes from an initial shape and compute a shape descriptor for each shape based on how the template has to be deformed to fit the shape. These descriptors are projected to a low dimensional manifold from which a deformation model is then learned, encoding the way in which the shapes vary the most. The user can then deform the template to explore the collection, with deformations with high variability being suggested by arrows. However, this method only works for shapes that are sufficiently close to the template and is thus less suited for shapes with varying topology. Similarly, Averkiou et al. [9] also abstract each shape by a box template and then encode the spatial arrangements of the boxes as a vector. Based on a distance measure between such configuration vectors, the shapes are embedded into a two-dimensional space that can be explored by the user by clicking on a position inside this space.

Semantic attributes: Another possibility is to use semantic attributes to encode the variations of shapes inside a collection. Chaudhuri et al. [10] conduct a user study to obtain semantic attributes that describe shape parts. To explore the collection of shape parts, the user can manipulate a slider for each attribute to find parts that posses more or less of that certain attribute. Yumer et al. [11] also perform a user study to gather semantic attributes for a shape category, but additionally connect those attributes to shape geometry to allow shape deformation by altering the semantic attributes of a shape. For exploration, the shapes are embedded into a two-dimensional space based on their attributes and a color map is used to visualize high and low values of a chosen attribute in the 2D space.

Images: More recently, the possibility of incorporating 2D images in the exploration of shape collections has also been considered. Hueting et al. [12] train classifiers on both images and shapes to align the image viewpoints with views of the 3D shape and then in turn estimate geometric properties of the shapes shown in the images based on the properties of the 3D shapes. By computation of the differences between their geometric properties, the user can then find 3D shapes that are similar to the shapes displayed in images or find images by setting the orientation of a 3D shape.

Shape descriptors: Kleiman et al. [13] use a nearest neighbor graph of the input shapes, computed using a combination of similarity measures, to generate a dynamic map that can be explored. A subset of shapes of the collection is arranged on a grid of limited size, with similar shapes close to each other. When the user scrolls in any direction or changes the zoom level, the new cells of the grid that enter the screen space are dynamically filled with shapes based on the similarity graph. Huang et al. [14] also combine multiple shape descriptors for their similarity measure and construct a category tree based on quartets of shapes, consisting of two pairs of shapes with high intra-pair similarity and low inter-pair similarity. Based on an initial shape, the other shapes are then arranged in a circular chart with the selected shape in the center and similar shapes in the circles close to the center.

Local features: The methods mentioned above mostly capture the global variability of the shapes, but it is also possible to concentrate on local features. Kim et al. [15] compute fuzzy correspondences between sample points on each shape. These correspondences are then used to define the similarity between shapes. The user can select one or more regions on an initial shape (regions with high variability are highlighted with a color map) to find other shapes with corresponding regions that are either very similar or dissimilar to the selected regions. Rustamov et al. [16] use a shape difference operator to extend this approach by allowing magnification or interpolation of shape differences, meaning that other shapes where a certain difference is more or less pronounced can also be found. Further extension of this method is described by

Please cite this article as: K. Leimer et al., Relation-based parametrization and exploration of shape collections, Computers & Graphics (2017), http://dx.doi.org/10.1016/j.cag.2017.07.001