Automatic Alignment of Shape Collections

Melinos Averkiou and Niloy Mitra

University College London

Abstract

We present a method for automatically aligning a collection of similar shapes in arbitrary initial poses. By analyzing the shape collection we extract a deformation model to capture the variability in the collection. We use this information to deform an extracted template shape and use it to align pairs of shapes by direct PCA alignment. We evaluate our method on synthetically created model collections in arbitrary initial poses and demonstrate accurate results with near ground truth alignment. Our algorithm significantly outperforms existing direct PCA alignment methods, without significant computational overhead.

1. Introduction

Shape collections are now ubiquitous with the proliferation of cheap 3D acquisition devices and 3D modeling tools. A number of interesting applications, ranging from shape search engines to assembly-based 3D modeling [FKS*04] depend on such collections. Such applications, however, assume the input models to be in mutual alignment, which is rarely the case, especially in public shape repositories (e.g., Google Warehouse, etc.). Manually aligning such shape collections is tedious, error prone, and often impractical for large collections. Current pairwise alignment approaches are not designed to handle the additional information implicitly contained in shape collections.

We present an automatic method to align a collection of similar shapes, for example a class of aeroplanes that can be extracted from a shape database using a keyword search. Our method takes advantage of the structure of the input collection and the low computational complexity of principal component analysis (PCA) alignment, without incurring the inaccuracies that usually accompany it. In a key observation, we use a deformation model and a template shape, extracted using previous work [OLGM11], as a replacement of the actual input shapes in the PCA alignment process. Specifically, instead of applying PCA to align a pair of input shapes $\mathcal{S}_i$ and $\mathcal{S}_j$, we apply PCA on $D_j(\mathbb{S}_i)$ and $\mathbb{S}_j$, where $D_j(\mathbb{S}_i)$ is the template shape extracted from $\mathcal{S}_i$, deformed to match $\mathcal{S}_j$, and $\mathbb{S}_j$ is the proxy shape extracted from $\mathcal{S}_j$. We tested the accuracy and speed of our method against direct PCA alignment methods, using synthetic data. Preliminary results indicate that our method outperforms such simple alternatives.

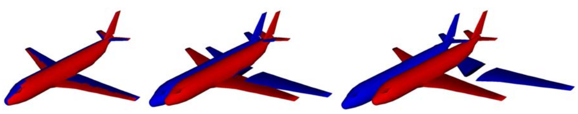

Figure 1: *Alignment results for three pairs of shapes from the synthetic dataset we used. The result of moving the wings outside the fuselage is unrealistic, but challenging to align.*

2. Background

Global alignment of pairs of shapes is a well studied problem, with solutions mainly based on normalization, such as PCA alignment [ETA02], or exhaustive search over the space of possible rotations. Optimization approaches have also been developed in order to increase the accuracy of PCA alignment [CVB09] or accelerate the search over the rotations space [Kaz07]. Such techniques, however, are designed to operate on pairs of shapes, which means they cannot take advantage of the structure of a shape collection. We observe that a collection of similar shapes often lies near a low-dimensional manifold in some shape descriptor space and this can be used to extract a deformation model that allows us to factor-out the variability of the shape collection, which is the main cause for the poor accuracy of direct PCA alignment methods.

3. Method

The input to our algorithm is a collection $\mathcal{S}$ of similar shapes. One shape $\mathcal{S}_0$ is chosen so that all other shapes are aligned to it. We now describe the main steps (see Figure 2).

DOI: 10.2312/conf/EG2012/posters/019-020

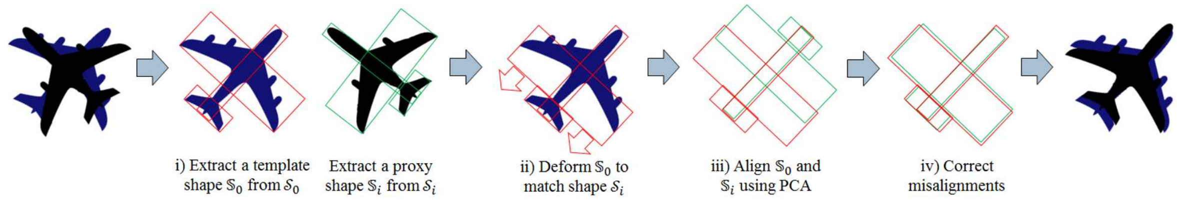

Figure 2: *Our algorithm. We start from randomly rotated similar shapes and align them using deformation information.*

i) Template shape extraction. We extract a deformation model that describes the variability of the shape collection [OLGM11]. We then extract a template shape $\mathbb{S}_0$ from shape $\mathcal{S}_0$ and a proxy shape $\mathbb{S}_i$ for any other shape $\mathcal{S}_i, i \in 1\dots N$ where N is the size of the shape collection, using connected component analysis. Each template shape is essentially a collection of bounding boxes.

ii) Template shape deformation. We use the extracted deformation model to deform the template shape $\mathbb{S}_0$ so that it matches shape $\mathcal{S}_i$. In the example of Figure 2, the template's component corresponding to the wings of the aeroplane moves along the fuselage.

iii) PCA alignment. We use PCA to recover the eigenvectors of the covariance matrix for template shape $\mathbb{S}_0$ and proxy shape $\mathbb{S}_i$. We align these eigenvectors, thus aligning $\mathbb{S}_0$ and $\mathbb{S}_i$.

iv) Misalignment correction. Due to eigenvector ambiguity, it is possible that $\mathbb{S}_0$ and $\mathbb{S}_i$ might be aligned wrongly in the previous step. We use an exhaustive approach, going through the 24 possible alignments and selecting the one that gives the smallest distance between $\mathbb{S}_0$ and $\mathbb{S}_i$. The resulting transformation aligns $\mathcal{S}_i$ to $\mathcal{S}_0$.

4. Results

We created several synthetic datasets to test our approach. Here we report results on 100 synthetic aeroplanes created by moving the wings of an aeroplane $\mathcal{S}_0$ along the fuselage and to its side (for examples of aeroplanes see Figure 1). Each model was then randomly rotated and translated. The rotations were stored as ground truth. The original aeroplane $\mathcal{S}_0$ was chosen so that all other aeroplanes were aligned to it. We tested the accuracy and speed of the following approaches: (i) Our method, (ii) direct PCA alignment, and (iii) best PCA alignment, i.e., the same as direct PCA alignment, but while aligning the eigenvectors, go through the 24 possible alignments and choose the best one.

Figure 3 illustrates the Frobernius distance to the ground truth for each pair of shapes $\mathcal{S}_0$ and $\mathcal{S}_i, i \in 1\dots 99$. As illustrated in the plot, our method clearly outperforms the other two methods, providing alignments very near the ground truth for all shape pairs. Direct PCA is rather unstable because of eigenvector ambiguities, while best PCA is better, but its accuracy steadily decreases until it breaks completely

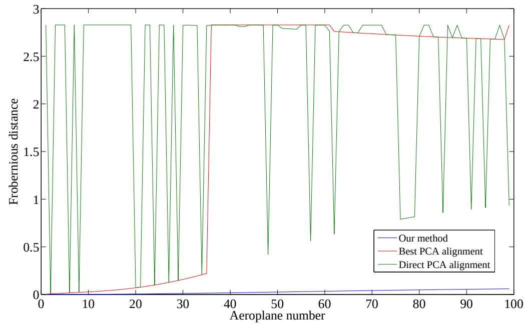

Figure 3: *Accuracy comparison of the three methods in terms of Frobernius distance to the ground truth. Note that the distance between the wings of aeroplane 0 and aeroplane 99 is 495 units on both x and y axes.*

somewhere around shape $\mathcal{S}_{40}$. Note that our algorithm is oblivious to initial orientations of the models in the collection. In terms of speed, our method is nearly as fast as direct PCA alignment and two orders of magnitude faster than best PCA alignment.

5. Conclusions and Future Work

We have presented an alignment method for collections of similar shapes that gives near ground-truth alignments, without incurring a high computational cost. Next, we would like to test our approach with real shapes from repositories to validate our initial results.

References

[CVB09] CHAOUCH M., VERROUST-BLONDET A.: Alignment of 3d models. *Graphical Models 71* (2009), 63–76. 1

[ETA02] ELAD M., TAL A., AR S.: Content based retrieval of vrml objects: an iterative and interactive approach. In *Proceedings of the 6th Eurographics workshop on Multimedia 2001* (2002), pp. 107–118. 1

[FKS*04] FUNKHOUSER T., KAZHDAN M., SHILANE P., MIN P., KIEFER W., TAL A., RUSINKIEWICZ S., DOBKIN D.: Modeling by example. *ACM Trans. Graph. 23* (2004), 652–663. 1

[Kaz07] KAZHDAN M.: An approximate and efficient method for optimal rotation alignment of 3d models. *IEEE TPAMI 29* (2007), 1221–1229. 1

[OLGM11] OVSJANIKOV M., LI W., GUIBAS L., MITRA N. J.: Exploration of continuous variability in collections of 3d shapes. *ACM Trans. Graph. 30* (2011), 33:1–33:10. 1, 2