Automatic Alignment of Shape Collections

Melinos Averkiou and Niloy Mitra

{m.averkiou,n.mitra}@cs.ucl.ac.uk

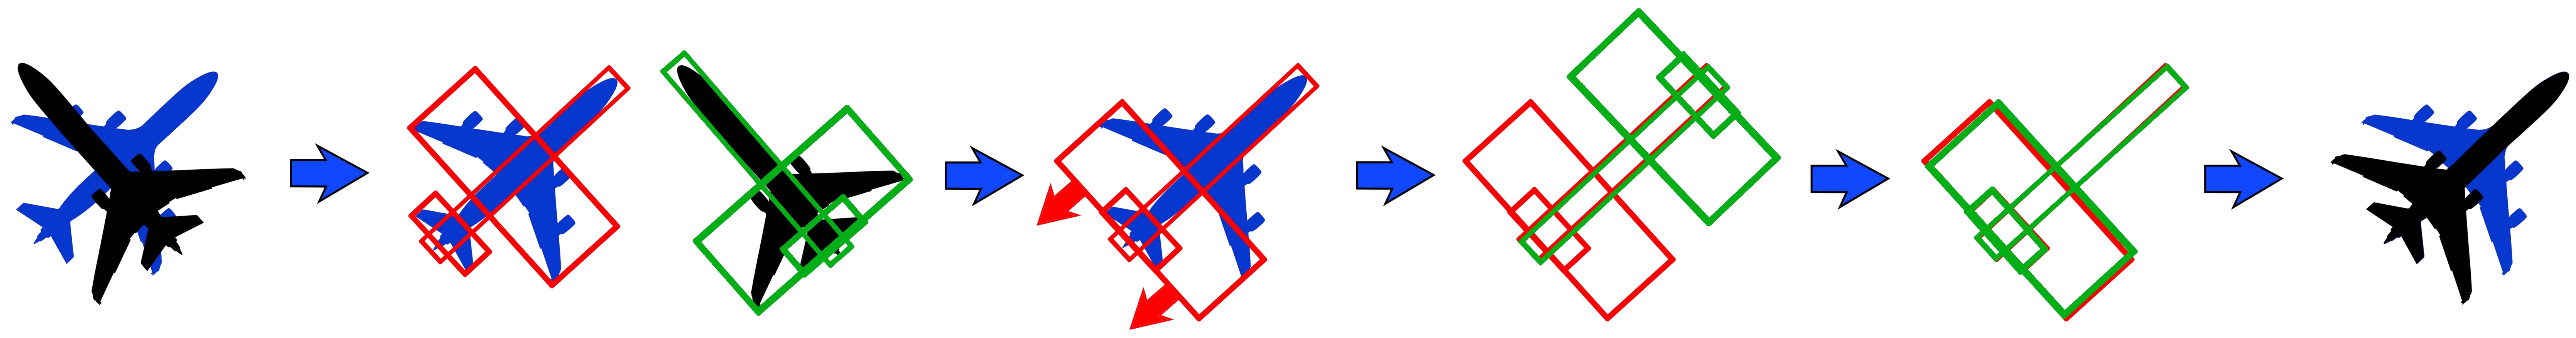

(a) Input: $\mathcal{S}_i$ (black) will be aligned to $\mathcal{S}_0$ (blue).

(b) Extract a template shape $\mathbb{S}_0$ (red) from $\mathcal{S}_0$ and a proxy shape $\mathbb{S}_i$ (green) from $\mathcal{S}_i$.

(c) Deform $\mathbb{S}_0$ to match shape $\mathcal{S}_i$.

(d) Align $\mathbb{S}_0$ and $\mathbb{S}_i$ using PCA.

(e) Correct misalignments.

(f) Output: $\mathcal{S}_i$ is aligned to $\mathcal{S}_0$.

Figure 1: Our algorithm. We start from a collection of similar shapes in arbitrary poses and align them using deformation information.

1. Motivation

Shape collections have become ubiquitous. Many interesting applications, ranging from shape search engines to assembly-based 3D modeling [FKS*04] depend on such collections. These applications typically assume the input shapes to be in alignment. Manually aligning these shapes is tedious and often produces imprecise results.

Some previous approaches for aligning pairs of shapes have relied on normalization of shape poses, typically by using Principal Component Analysis (PCA) [ETA02], or on exhaustive search over the space of rotations. None of these were designed for aligning shape collections.

We present a method for the automatic alignment of a collection of similar shapes. Similar means they belong to the same semantic class, for example a class of cars.

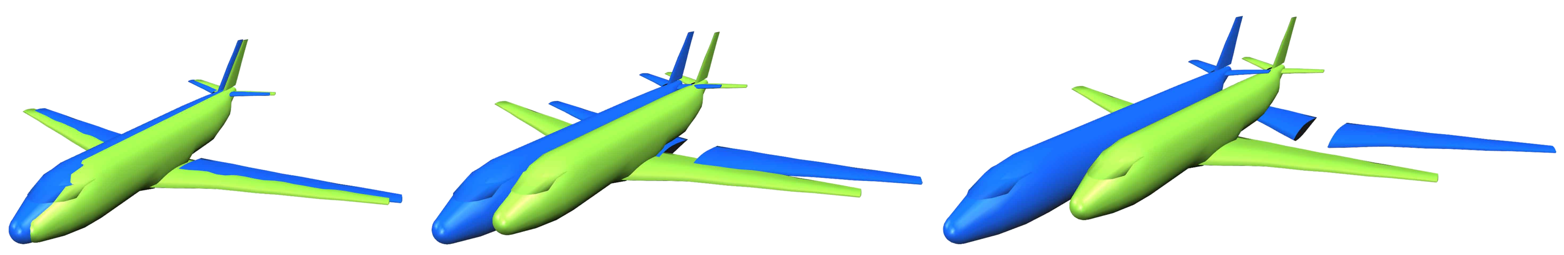

Figure 2: Alignment results for three pairs of shapes from the synthetic dataset we used.

2. Method

Our method takes as input a collection $\mathcal{S}$ of similar shapes. One shape $\mathcal{S}_0$ is chosen so that all other shapes are aligned to it. The main steps are (see Figure 1):

Step 1: 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 $\mathcal{S}_i, i \in 1 \ldots N$ where N is the size of the shape collection, using connected component analysis. Each template shape is essentially a collection of bounding boxes.

Step 2: 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 1, the template's component corresponding to the wings of the aeroplane moves along the fuselage.

Step 3: 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 $\mathbb{S}_i$.

Step 4: Misalignment correction.

Due to eigenvector ambiguity, it is possible that $\mathbb{S}_0$ and $\mathbb{S}_i$ are aligned wrongly in step 3. We go through the 24 possible alignments and select the one that gives the smallest distance between $\mathbb{S}_0$ and $\mathbb{S}_i$. The resulting transformation aligns $\mathcal{S}_i$ to $\mathcal{S}_0$.

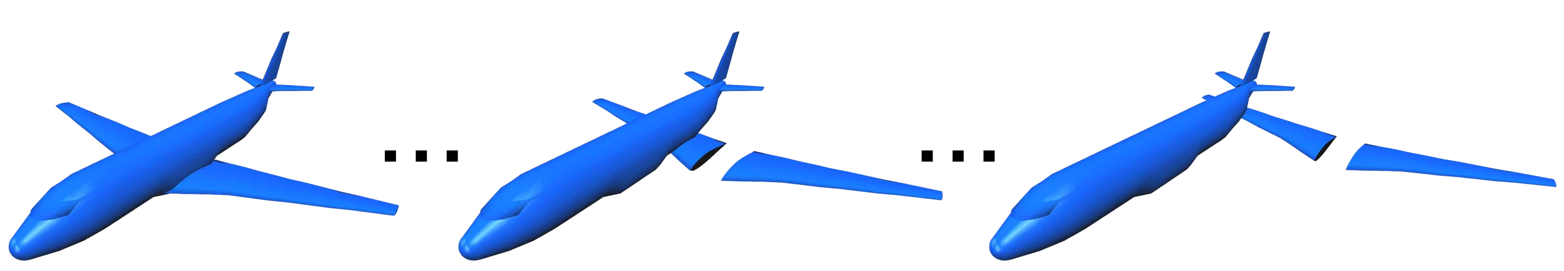

Figure 3: The synthetic dataset we used for testing our approach.

3. Results

We created 100 synthetic aeroplanes by moving the wings of an aeroplane $\mathcal{S}_0$ along the fuselage and to its side (for examples of aeroplanes see Figure 3). 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 of the following approaches:

1. Our method
2. Direct PCA alignment
3. 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 4 illustrates the Frobernius distance to the ground truth for each pair of shapes $\mathcal{S}_0$ and $\mathcal{S}_i, i \in 1 \ldots 99$. As illustrated in the plot, our method clearly outperforms the other two methods, providing alignments very near the ground truth for all shape pairs.

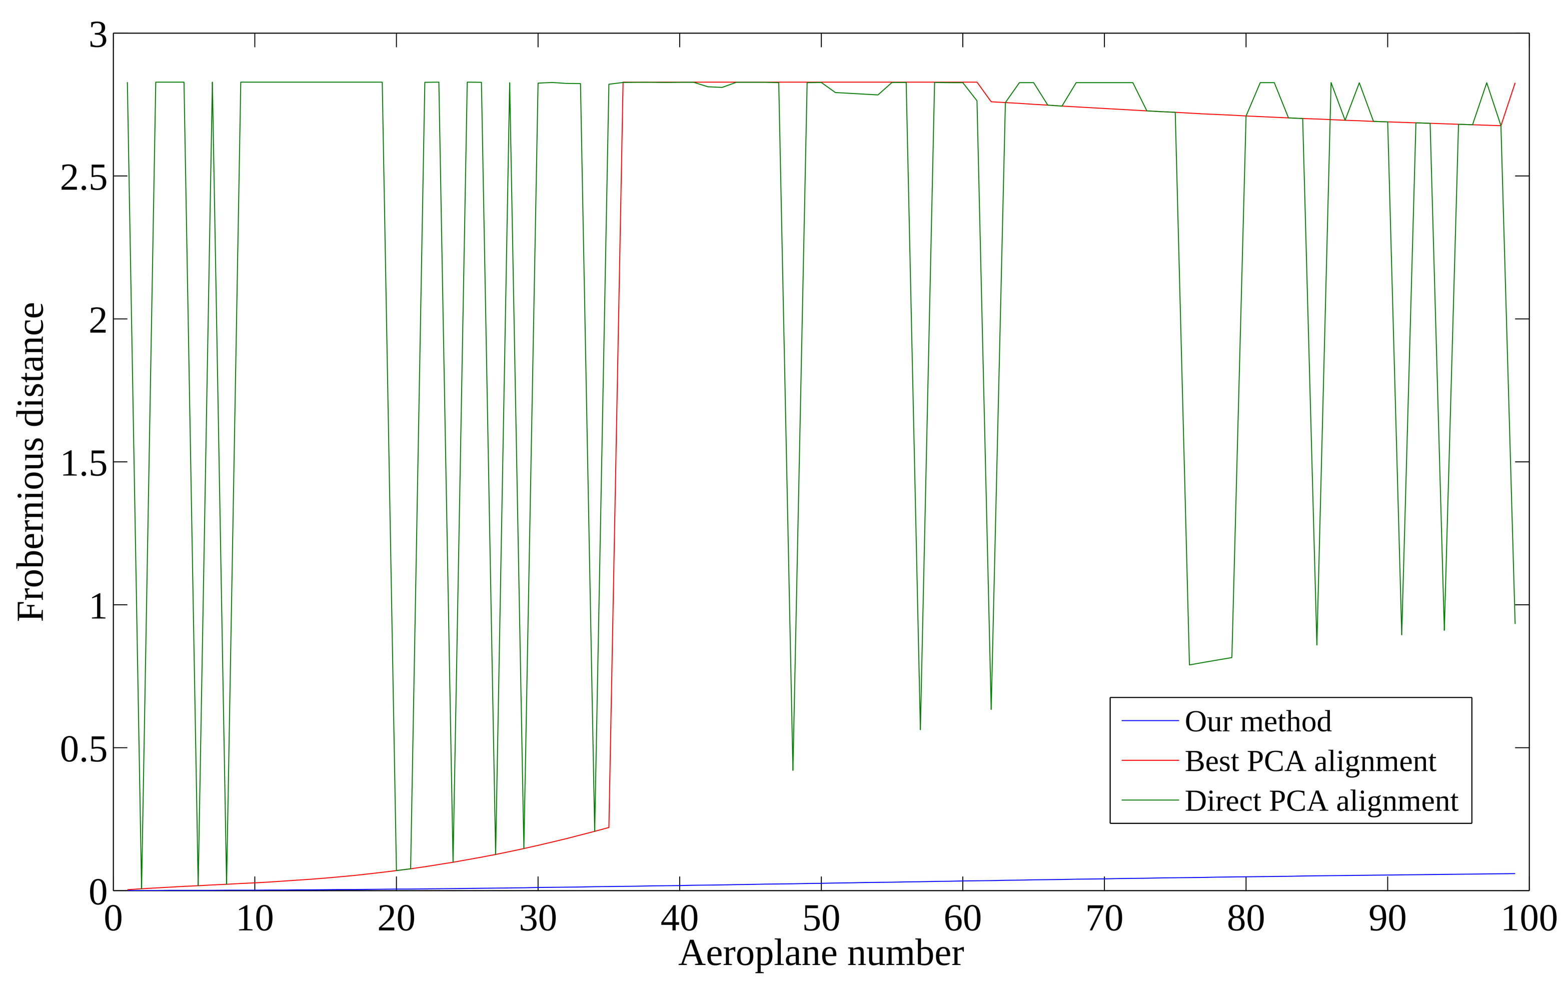

Figure 4: Accuracy comparison of the three methods in terms of Frobernius distance to the ground truth. Our method gives near ground truth alignments.

4. Conclusions and Future Work

Our alignment method for collections of similar shapes 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 on this synthetic dataset.

References

[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.

[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.

[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.

UCL Engineering Doctorate
Virtual Environments Imaging & Visualisation