Eurographics Symposium on Geometry Processing (2014)
Thomas Funkhouser and Shi-Min Hu (Guest Editors)

Statistical Mesh Shape Analysis with Nonlandmark Nonrigid Registration

J. Dupej[1,2] and V. Krajíček[1] and J. Velemínská[2] and J. Pelikán[1]

[1]Faculty of Mathematics and Physics, Charles University in Prague, Czech Republic
[2]Faculty of Science, Charles University in Prague, Czech Republic

Abstract

The analysis of shape represented as surface meshes is an important tool in anthropology and biomedicine for the study of aging, post-treatment development or sexual dimorphism. Most approaches rely on nonrigid registration using manually placed homologous landmarks, it is however often the case that some regions cannot be landmarked due to the lack of clear anatomical features. We therefore present a method of analyzing and visualizing the variability of a set of surface models that does not rely on landmarks for feature matching and uses coherent point drift (CPD), a nonrigid registration algorithm, instead. Our approach is based on the topology transfer of one arbitrarily selected base mesh to all other meshes with the use of CPD. The procedure ensures the identical meanings of corresponding vertices across the sample and allows the use of multivariate statistics even with shapes that would be difficult to process with methods that rely on landmarks for feature-matching.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

1. Introduction

A quantitative statistical analysis of mesh shape applied to large sets of specimens is of interest in scientific fields such as biology, anthropology and medicine. Due to the increasing availability of surface scanners in the last two decades, the methods for analyzing this modality of data have been in demand. Point clouds, and the subsequently reconstructed triangle meshes, are the output of these scanners.

The point clouds represent the sampled surfaces in points that are unrelated across the sample, i.e. the vertices with identical indices generally do not represent the same feature. Due to this fact, several methods of finding corresponding points in pairs of meshes have been developed, such as [HBH01] or [MJS*06]. These algorithms resample one surface in new loci that have equivalent (homologous) meaning to the other mesh's respective vertices. After that step, individual specimens (i.e. their resampled meshes) can be represented simply as long vectors constructed by the concatenation of their vertices' coordinates. These vectors are the observations of the same random variable and thus multivariate statistical methods can be applied.

Correspondence search is unquestionably the most challenging aspect of the shape analysis. That is because each pair of corresponding vertices must be homologous; however the point cloud itself offers few direct clues to finding such vertex pairs. This problem is generally reduced to the search for closest point in a point cloud if a non-rigid registration algorithm is used to bring the corresponding areas of the both meshes close together. Many approaches rely on fitting a deformation function to the pairs of landmarks that are placed in anatomically equivalent locations by an expert. Unfortunately, in many biological studies the surfaces cannot be landmarked densely enough, moreover smooth areas with very low curvature (e.g. forehead, cheeks) provide no clear features that could be reliably described with landmarks. This often results in the deformation not aligning such areas correctly. We therefore present an enhancement to DCA, an algorithm for statistical mesh shape analysis [HBH01] that we base on coherent point drift (CPD), a non-rigid registration algorithm by [MS10].

2. Our Method

The processing pipeline of our algorithm consists of rigid registration, mesh resampling, statistical analysis and optional shape prediction. The first step, rigid registration re-

quires the knowledge of a small count of homologous landmarks on each mesh. Generalized Procrustes analysis (GPA) [Gow75] is used to find the transformations that align each landmark set to the mean shape. The alignment of meshes is then performed using the transformations obtained from GPA. Please note that landmarks are not used beyond this point.

The next phase performs vertex-to-point matching and remeshing. Before that can be done, a *base mesh* has to be chosen as its topology will be used as a template for all other (*floating*) meshes. That mesh is chosen arbitrarily. The topology transfer is performed by nonrigidly registering the base mesh to a floating mesh using CPD configured for low motion smoothing ($\beta = 2.0$) and moderate regularization ($\lambda = 3.0$). As the registration will probably not be ideally tight, we project these vertices to the floating mesh. In many approaches this is done by finding the closest point on the floating mesh to each deformed base vertex $\tilde{x}_i$. We find the closest points y_i^d but also calculate the surface normals n_i of the deformed base mesh and find the closest intersections y_i^n of the normal and the floating surface. If the intersection y_i^n exists and $\|\tilde{x}_i - y_i^n\| < 3\|\tilde{x}_i - y_i^d\|$, y_i^n is used as the coordinates of i-th remeshed vertex, otherwise y_i^d is used.

At this point, the statistical analysis could be performed. It is, however often the case that there are vertices present that produce unwanted variability. That is primarily caused by inconsistent trimming of the studied meshes in their manual preparation or noise occurring on the eyeballs or regions with hair, which tends to manifest itself in one of the leading principal components. These phenomena result in overcompaction or overstretching of triangles which are removed by this stage of the pipeline In each specimen, the areas of each triangle a_{res} are compared to areas of the corresponding triangles on the base mesh a_{base}. If $a_{res}/a_{base} < 0.1$ or $a_{res}/a_{base} > 10$, all vertices of that triangle are marked defective in the particular specimen. If a vertex is defective in at least $k_{bad} = 0.05n$ of the specimens, it is removed from the subsequent statistical analysis. That threshold has been determined empirically and should be increased if the meshes contain a lot of high-frequency content.

After the removal of the unwanted vertices, high-dimensional principal component analysis (PCA) is performed [Bis06, p. 570]. If regression is used to model the PC scores and shape prediction is desired, it should be noted that the vertices that were discarded in the previous phase will not be predicted and must be approximated. To that end we find the vertices that lie on the boundaries of excluded regions. A thin-plate-spline (TPS) [Boo89] is fitted to the locations of those vertices in the base mesh and predicted mesh. That deformation is then used to transform the excluded vertices on the base mesh, producing an approximation of their location on the predicted surface.

This entire pipeline has been implemented and tested in Morphome3cs [Mor14].

3. Discussion

We have developed an algorithm for statistical analysis of sets of meshes. Analogous to other approaches, it performs mesh resampling in a way that produces homologous vertices across the sample and thus allows the use of multivariate statistics. Unlike other popular algorithms, CPD, a nonlandmark nonrigid registration algorithm, is used for feature matching which removes the need to rely on landmarks at this stage. Furthermore, vertices that may cause unwanted variability are removed before entering PCA for increased robustness. Finally, in the shape prediction phase a TPS-based approach is used to fill the holes that were created by excluding vertices.

We have tested our approach on two data sets, one consisting of 194 facial scans, the other comprising 56 scans of palatal casts. In the first set, we modeled aging of the human face from 20 to 80 years of age using quadratic regression. In sex discrimination using support vector machines [Bis06, p. 325] we outperformed [HBH01] with a cross-validation success rate of 92.8 % over 87.1 %. The palatal scans were of children with cleft lip and palate before receiving a corrective surgery and approximately 10 months after it. On that set, we modeled the growth of the palate after surgery.

To conclude, our mesh set analysis algorithm has many potential uses for the study of shape and form in the fields of anthropology, biology and medicine. It provides an advantage over other methods that rely on an even coverage of homologous landmarks on the studied meshes as this may not be possible to achieve with many real-life datasets.

Acknowledgments We would like to thank Miriama Mydlová and Eva Hoffmannová for kindly allowing us the use of their facial and palatal scan datasets in our tests. This research has been supported by GAUK as project #1388213.

References

[Bis06] BISHOP C. M.: *Pattern Recognition and Machine Learning*. Springer Science+Business Media, New York, 2006. 2

[Boo89] BOOKSTEIN F.: Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence 11*, 6 (1989), 567–585. 2

[Gow75] GOWER J. C.: Generalized procrustes analysis. *Psychometrika 40* (1975). 2

[HBH01] HUTTON T., BUXTON B., HAMMOND P.: Dense surface point distribution models of the human face. *Proceedings IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA 2001)* (2001), 153–160. 1, 2

[MJS*06] MAO Z., JU X., SIEBERT J. P., COCKSHOTT W. P., AYOUB A.: Constructing dense correspondences for the analysis of 3D facial morphology. *Pattern Recognition Letters 27*, 6 (Apr. 2006), 597–608. 1

[Mor14] Morphome3cs II. `http://cgg.mff.cuni.cz/Morpho/`, 2010–2014. 2

[MS10] MYRONENKO A., SONG X.: Point set registration: coherent point drift. *IEEE transactions on pattern analysis and machine intelligence 32*, 12 (2010), 2262–75. 1