Customized Design of Hearing Aids Using Statistical Shape Learning

Gozde Unal[1], Delphine Nain[2], Greg Slabaugh[3], and Tong Fang[3]

[1] Faculty of Engineering and Natural Sciences, Sabanci University, Turkey
gozdeunal@sabanciuniv.edu
[2] McKinsey& Company, USA
[3] Siemens Corporate Research, Princeton NJ, USA⋆

Abstract. 3D shape modeling is a crucial component of rapid prototyping systems that customize shapes of implants and prosthetic devices to a patient's anatomy. In this paper, we present a solution to the problem of customized 3D shape modeling using a statistical shape analysis framework. We design a novel method to learn the relationship between two classes of shapes, which are related by certain operations or transformation. The two associated shape classes are represented in a lower dimensional manifold, and the reduced set of parameters obtained in this subspace is utilized in an estimation, which is exemplified by a multivariate regression in this paper. We demonstrate our method with a felicitous application to estimation of customized hearing aid devices.

1 Introduction

3D shape modeling and estimation is a crucial task in custom design of anatomical shapes. A sample shape estimation problem in systems for rapid prototyping of hearing aid devices is depicted in Figure 1. For a comfortable fit, it is important that the shape of the hearing aid match the patient's ear geometry. The two classes of shapes, here patients' 3D raw ear impressions and the output hearing aid shapes, are normally related by certain operations or a transformation R. Current practice involves mainly a manual design (even in an electronic environment), and the goal is to automate this process for increasing efficiency, patient comfort, repeatability, and throughput in audiologist offices.

Our work is similar in spirit to the image analogy problem [1], where a new painting D is produced by the input photograph C, by copying matching patches from a prior painting/photo pair A/B. The explicit relation between the pairs is not learned however, which is our aim in this work. The celebrated active shape models work [2] developed a compact description of the variation of shapes in a class using statistical methods. In [3], a statistical shape model is built for the human ear canal (as point clouds), where the correspondences are obtained by warping a template onto shapes, which are annotated with 18 landmarks by a specialist. In [4] a smoother dense mesh is obtained by a Markov field regularization of the correspondence field. In these works, the ear canal model is used for analysis of gender differences in its shape, and for its deformation
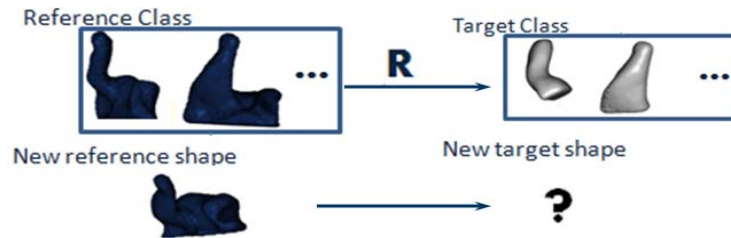

Fig. 1. Shape Estimation Problem

by mandibular movement [5]. Manual marking of landmarks is not suitable for rapid prototyping systems, moreover, finding stable feature points in all shapes is difficult due to individual variations. The correspondence problem was alleviated in an Eulerian shape representation framework in [6], which used PCA to build a shape prior to guide the segmentation of objects in images. Variations on PCA such as kernel PCA [7] and principal factor analysis [8] were employed for statistical shape analysis, although PCA is preferred here for its optimality in dimensionality reduction. In a recent work [9], regression techniques are utilized to investigate degrees of correlation and dependence variation between shapes of different structures within the brain. Our method is different in mathematical details and includes deformation of predicted shapes via predicted difference masks and fitted planes, with a focus on shape generation for prosthetic devices. [10] took a machine learning approach for image segmentation, and [11] used Frèchet expectation to generalize univariate regression to manifold-valued data to study the effect of aging on brain shape in patient populations.

Our main contribution is the development of an automatic shape transformation method to be used in various applications like customized design of anatomical parts. Our system learns the relationship between two classes of shapes and generates a shape from one class when given as an example a shape from another class. We design an asymmetric registration for shapes that significantly differ in geometry in Sec 2.1. We then propose a novel shape generation technique based on multivariate regression in Sec 2.2 and 2.3.

2   Method

There are various styles of hearing aids such as canal, and in-the-ear [12], and the design process starts with a rough mold of the patient's ear, so called undetailed shell (or shape), that is then detailed by a specialist. The detailing process includes cutting unused parts based on the desired shell style and the geometry of the patient's mold (Fig 2), rounding edges and other operations needed to fit the electronics in the shell. This is a time-consuming process that is based on the skills and experience of the specialist. Alternatively, the specialist now can carry out the detailing on digitized ear shells using CAD software systems, which are still not-fully automatic. Our work aims at removing this bottleneck in rapid prototyping systems for hearing aid devices.

We obtained 90 digitized undetailed molds and their corresponding detailed molds from a specialist. We define a "Ground truth" (GT) shape as the shape, which is detailed
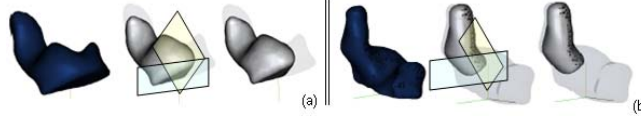
Fig. 2. Ear impressions;(a) Half shell, (b) Canal shell; undetailed shapes (left); detailed shapes with (middle) and without cutting planes (right)

by the audiologist, and is later digitized. Ear shells were digitized by a 3D laser scanner, a 3D point cloud is obtained and triangulated to build a polygonal mesh surface, which is then converted to a voxel and signed distance function (SDF) representation [6]. The dataset consisted of 41 molds of Half-Shell type (Figure 2-a), and 49 molds of Canal Shell type (2-b). The undetailed shell is made of two main parts: the long and thin structure on the top is the ear canal and the round bowl-like structure is the concha which funnels sound to the ear canal [13]. We see that the detailed half shells occupy most of the ear canal and a large part of the concha, and the detailed canal shells occupy mostly the ear canal.

2.1 Asymmetric Shape Registration

The reference (undetailed) shapes must be registered with their corresponding target (detailed) shapes in the training set. In a variational registration setting, we propose an asymmetric distance restricted to a band around both the detailed shape Φ^d and undetailed shape Φ^u, but constrained mainly by the smaller of the two shapes. This provides a new asymmetric rigid registration differential equation, which is derived from sum of squared distances energy functional between Φ^u and Φ^d:

$$\frac{\partial g_i}{\partial t} = \int_\Omega \mathcal{X}_\beta\left(\Phi^u(\boldsymbol{X}), \Phi^d(g(\boldsymbol{X}))\right) \left[\Phi^u(\boldsymbol{X}) - \Phi^d(g(\boldsymbol{X}))\right] \left\langle \nabla\Phi^d(g\boldsymbol{X}), \frac{\partial g(\boldsymbol{X})}{\partial g_i} \right\rangle d\boldsymbol{X}$$

(1)

where g is a rigid transformation $g(\boldsymbol{X}) = R\boldsymbol{X} + \boldsymbol{T}$, $\boldsymbol{X} \in \mathbb{R}^3$, with parameters g_i of 3D rotation matrix R, and 3D translation $\boldsymbol{T}$. Φ^u and Φ^d are the undetailed and detailed shell SDFs defined over the domain Ω. The new characteristic function takes the form

$$\mathcal{X}_\beta(\Phi_u, \Phi_d) = \left\{ \begin{array}{l} 0, \max(|\Phi_u|, |\Phi_d|) > \beta \\ 1, \max(|\Phi_u|, |\Phi_d|) < \beta \end{array} \right.$$

(2)

Figure 3 depicts an application to ear shell registration, where the symmetric one [14] in (b) fails because the detailed shape Φ_d is significantly smaller than the undetailed shape Φ_u and parts of the undetailed shape that do not exist in the detailed shell still influence the registration. In the asymmetric version, the alignment was successful as shown in (c).

Note that in addition, all undetailed shapes in the training are aligned with the symmetric registration so that the variation in the data is due to the geometry and not the pose.
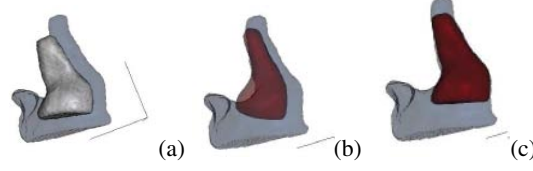
Fig. 3. Registration of detailed ear impressions to undetailed: (a) shapes before registration; (b) shape (red) after symmetric registration; (c) shape (red) after registration by Eq.(1)
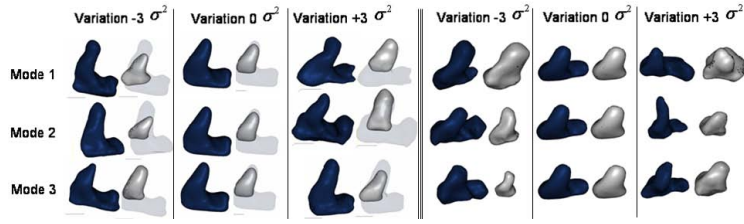

Fig. 4. The first 3 modes for both the undetailed (dark) and detailed (light) shapes. The weights are varied from $-3\sigma^2 \leq w_i \leq 3\sigma^2$. Left: Canal shell, Right: Half shell.

2.2 Shape Estimation

After alignment of the shapes, we will represent them in a lower dimensional manifold and relate the undetailed shapes with the detailed shapes on this manifold as explained next.

We conduct a standard statistical analysis on a training set of N undetailed shapes to obtain a shape variability matrix $\boldsymbol{S^u} = [\tilde{\boldsymbol{\Phi}}_1^u \ \tilde{\boldsymbol{\Phi}}_2^u \ \cdots \tilde{\boldsymbol{\Phi}}_N^u]$ on which a PCA is carried out: $\frac{1}{N}\boldsymbol{S^u}\,\boldsymbol{S^u}^T = \boldsymbol{U^u}\,\boldsymbol{\Sigma^u}\,\boldsymbol{U^u}^T$. Here the columns $\boldsymbol{e}_i^u$ of the matrix $\boldsymbol{U^u}$ represent the orthogonal modes of variation in the undetailed shapes, called eigenshapes, and the diagonal matrix $\boldsymbol{\Sigma^u}$ contains the corresponding eigenvalues, $\boldsymbol{\sigma}_i^u$. A similar analysis is carried out for the detailed shapes. Each shape $\boldsymbol{\Phi}_i^u$ and $\boldsymbol{\Phi}_i^d$ in the undetailed and detailed classes then can be represented by a vector of weights as:

$$\boldsymbol{w}_i^u = \boldsymbol{U^u}^T(\boldsymbol{\Phi}_i^u - \boldsymbol{m^u}), \text{and } \boldsymbol{w}_i^d = \boldsymbol{U^d}^T(\boldsymbol{\Phi}_i^d - \boldsymbol{m^d}), \tag{3}$$

where $\boldsymbol{m^*}$'s are the mean shapes. A small number of principal modes, k, which is selected as the same for both the undetailed and detailed shape classes, explains a significant portion of the variability in the ear impression shape space (see Sec 3).

Figure 4 depicts the first 3 modes of variations for both the undetailed and detailed shape datasets. In the Canal Shell dataset, we see similar variations as observed in [3]. The mode 1 deformation corresponds to a bending of the canal and a flattening of the concha for the undetailed shape, and a widening at the base of the canal and a height shift for the detailed shape. Mode 2 corresponds to a thickening of the concha and a bending of the canal for the undetailed shape and a general size change for the detailed shape. Mode 3 corresponds to a flattening of the ear canal for both shapes. For the Half Shell variations, we see that mode 1 corresponds to a general size change of the concha
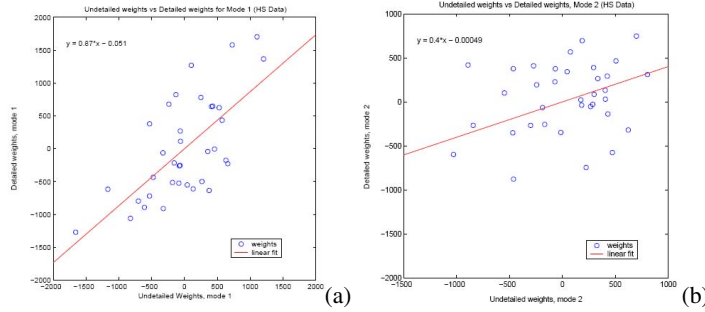
Fig. 5. The weights of undetailed vs. detailed shapes associated to 1^{st} (a), 2^{nd} mode (b)

for both detailed and undetailed shapes. Mode 2 corresponds to an overall size change with a flattening of the tip of the ear canal and the concha. Mode 3 corresponds to narrowing/widening of the concha and the ear canal.

For the undetailed set, we form the $N\mathrm{x}k$ weight matrix $\boldsymbol{W^u}$ where row i is the vector $\boldsymbol{w_i^u}$ representing the ith undetailed shape. Similarly, $\boldsymbol{W^d}$ represents the weight matrix for the detailed shapes. We would like to find a model that best describes the relation between the two shape classes in this highly reduced dimensional space, i.e., a mapping $\boldsymbol{R}$ between the two shape representations as:

$$\boldsymbol{R} \circ \boldsymbol{W^u} = \boldsymbol{W^d}.\tag{4}$$

We discuss the nature of this mapping we would like to find next. Figure 5 shows the first mode of variation plotted for each detailed (y-axis) and undetailed (x-axis) pair of shapes in the Half Shell Data. It can be observed that there exists a correlation that is close to linear between both sets of weights. As expected, the y-intercept is close to 0 since an undetailed weight of 0 on the major axis of variation means that the undetailed shape is very close to the mean and therefore the detailed shape should also be expected to be close to the mean. Similarly a linear relationship is observed for the second mode as shown in Fig. 5(b). We also examined other mode combinations such as first mode vs. second mode and third mode, and found close to linear correlations. Intuition gained from these experiments lead us to assume a linear relationship between the two classes of shapes. In order to find a general multivariate regression between all the weights, we construct a linear least squares optimization problem:

$$\boldsymbol{W^u}\,\boldsymbol{X^d} = \boldsymbol{W^d}\tag{5}$$

where $\boldsymbol{X^d}$ is a transformation matrix that encodes the detailing process: it transforms the undetailed shape class into the detailed shape class. To find $\boldsymbol{X^d}$, $\boldsymbol{W^u}$, which is of size $N\mathrm{x}k$, where $N > k$ (hence overdetermined problem), is inverted by an SVD decomposition [15]:

$$\boldsymbol{W^u} = \boldsymbol{U^w}\,\boldsymbol{D^w}\,\boldsymbol{V^w}\,^T,\tag{6}$$

then the least squares solution $\boldsymbol{X^d}$ is given by:

$$\boldsymbol{X^d} = \boldsymbol{V^w}\,(\boldsymbol{D^w})^{-1}\boldsymbol{U^w}\,^T\boldsymbol{W^d}\tag{7}$$

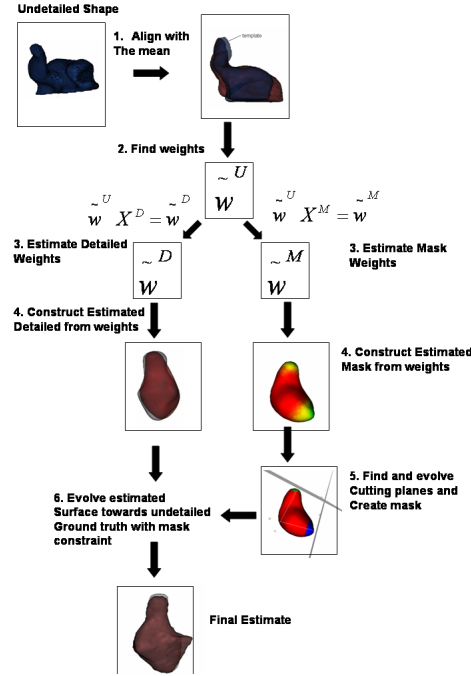where $(\boldsymbol{D^w})^{-1}$ is the simple inverse of a diagonal matrix with singular values.

Fig. 6. Flowchart describing the proposed automatic shape generation method

We build a statistical model for a class of auxiliary shapes, called mask shapes, to help deform the estimated model towards the patient ear canal anatomy for an exact fit. 3D mask shapes M_i are formed from the difference of shapes in the two training sets, i.e. $\Phi_i^u - \Phi_i^d$. This mask will be used to indicate regions in the estimated detailed shape that are allowed to propagate towards the undetailed shell, except at the parts that we detected as "cuts" by our algorithm during the final deformation phase. Hence along with the estimated matrix for the detailed shapes X^d, a second regression matrix X^m is estimated and stored during the training phase.

2.3 Automatic Shape Generation

After the training phase, a new undetailed shape is given as input to the system, and the expected output is a corresponding detailed shell similar to one that would have been produced by a specialist. We note that since different types of detailed shells exist, the training phase must be done on the type that is expected as an output of the estimation.

Flowchart in Figure 6 summarizes our automatic shape generation method: The new undetailed shape is registered to the mean shape from the training data as explained in Section 2.1. The weight vector for the new undetailed shape is computed and the weight vector for the detailed shape is estimated through the stored regression matrix (7)
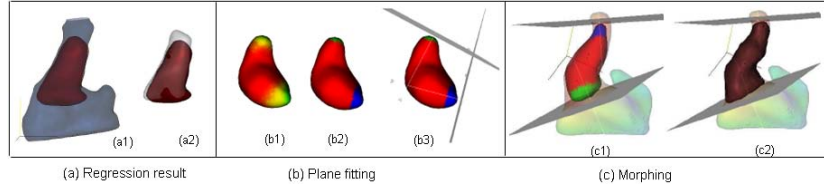
Fig. 7. (a) Detailed shape found through regression is depicted (in red/dark) on (a1) the undetailed shape; (a2) the GT detailed shape (in grey/light). **(b)** (b1) estimated mask; (b2) the resulting clustering with 2 clusters; (b3) the fitted planes. **(c)** Morph of an initial detailed surface (c1) towards the undetailed surface constrained by planes, final result in (c2).

(Fig 7-a). Next, difference mask weights are estimated via the regression matrix X^m to form the mask M, which is binarized with a threshold of 1. The regions of M that are 0 correspond to "cuts" that would have been made by the technician. Those cuts usually produce flat surfaces. Since the estimated detailed shapes do not always have flat surfaces where the mask is 0, first we cluster points on the mask surface by using the k-means algorithm, then we fit planes to these regions as shown in Figure 7-b. Finally, we morph the initially estimated detailed shape towards the undetailed shape where the evolution mask field is nonzero. It is allowed to evolve until it reaches the plane and flattens out as depicted in Figure 7-c. As for the surface deformation step, we utilized a simple morphing PDE as in [16] modified with our estimated mask and the planes.

3   Results

For both datasets (half-shell and canal shell), we randomly split the data so that 90% of the data was used for training and 10% used for testing. This operation is repeated three times for more robust validation. The number of modes needed to explain 95%, 97% and 99% variability in the data, were 12, 15, and 22, respectively. For a 95% variation, only 12 modes are retained, which is a significant decrease from 22 modes needed to explain 99% of the variability in the data. This also shows the huge dimensionality reduction given that the shapes in our training dataset were represented on a voxel grid of 80^3.

The sum of squared difference (SSD) in millimeters between the estimated shape $\tilde{\Phi}_d$ and the ground truth shape Φ_d is used as a validation measure. Table 1 shows the average SSD values obtained for both datasets. The overall SSD value between the estimated shape and the ground truth was smaller than $1.5mm$ in all cases.

Table 1. Average SSD (in mm) between the estimated shapes and the ground truth shapes in the Half-Shell (HS) and Canal (C) dataset for all three test sets and three variation values

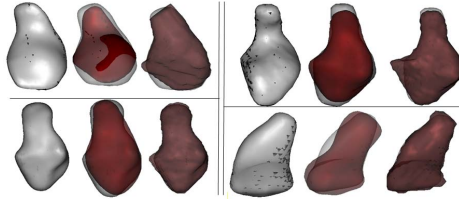| SSD (HS)	95% variation	97% variation	99% variation	SSD (C)	95% variation	97% variation	99% variation
test 1	1.16	1.19	1.17	test 1	1.51	1.47	1.40
test 2	1.07	1.17	0.97	test 2	1.41	1.40	1.35
test 3	1.37	1.35	1.32	test 3	1.32	1.32	1.31

Fig. 8. (Each Quadrant)Left: GT detailed shape; Middle: Estimated detailed shape before evolution; Right: Estimated Detailed Shape after evolution (all are superimposed on the GT shape)

We show qualitative results for the Half-Shell dataset in Figure 8, where the estimated shapes are observed to be in good agreement with the GT shapes, particularly in the canal region. The errors as expected are distributed around the cut points since the shape evolution is restricted to the cutting planes which are estimated through our shape learning process on a mask field. Our preliminary but extensive experiments have shown that our algorithm is robust to different shell types, and a low SSD value is obtained for all tests. As part of our future studies, further validation will be performed on our system.

4 Conclusions and Discussions

We presented a general framework to automatically generate a target shape from a reference shape via learning the relation of these two shape classes on a much lower dimensional manifold than the original shape space. As a specific application, our system learns how to detail a hearing aid shape by estimating a mapping from a patient's digitized ear mold to the detailed shell. Further refinement of the shape is achieved by deforming the estimated shell towards the undetailed shell in regions where the shapes should fit using an estimated auxiliary shape class. We learned the relation between two shape classes through a linear multivariate regression due to the underlying assumption that the two shape classes have a linear relation, which is demonstrated to be a reasonable assumption by our results.

Our proposed framework is quite general in that a mapping between the two classes of shapes can be estimated in a proper setting. This framework contains several known components such as variational registration, PCA in the distance transform space for shape analysis, and linear regression used in an interesting and novel system for automatic shape transformation. Some of these components can be replaced with other ones: for instance, using non-linear regression instead of linear regression (to account for more complex relations between associated classes of shapes), kernel PCA or manifold learning instead of PCA, Eulerian vs. Lagrangian shape representation, and others. While such changes may improve the results, they would not change the overall concept. However, we plan to apply our technique with different components to other applications in custom design of various anatomical parts such as dental implants and prosthetic hips.

References

1. Hertzmann, A., Jacobs, C., Oliver, N., Curless, B., Salesin, D.: Image analogies. In: SIG-GRAPH Conference Proceedings (2001)
2. Cootes, T., Taylor, C., Cooper, D., Graham, J.: Active shape models - their training and application. Computer Vision and Image Understanding 61(1), 38–59 (1995)
3. Paulsen, R., Larsen, R., Nielsen, C., Laugesen, S., Ersboll, B.: Building and testing a statistical shape model of the human ear canal. In: Dohi, T., Kikinis, R. (eds.) MICCAI 2002. LNCS, vol. 2489, pp. 373–380. Springer, Heidelberg (2002)
4. Paulsen, R.R., Hilger, K.B.: Shape modelling using markov random field restoration of point correspondences. In: Taylor, C.J., Noble, J.A. (eds.) IPMI 2003. LNCS, vol. 2732, pp. 1–12. Springer, Heidelberg (2003)
5. Darkner, S., Larsen, R., Paulsen, R.: Analysis of deformation of the human ear and canal caused by mandibular movement. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part II. LNCS, vol. 4792, pp. 801–808. Springer, Heidelberg (2007)
6. Leventon, M., Grimson, W., Faugeras, O.: Statistical shape influence in geodesic active contours. In: IEEE CVPR, Hilton Head, vol. 1 (2000)
7. Rathi, Y., Dambreville, S., Tannenbaum, A.: Statistical shape analysis using kernel pca. In: IS&T, SPIE Symposium on Electronic Imaging (2006)
8. Aguirre, M., Linguraru, M., Marias, K., Ayache, N., Nolte, L.P., Ballester, M.: Statistical shape analysis via principal factor analysis. In: IEEE Int. Symp. Biomedical Imaging (2007)
9. Rao, A., Aljabar, P., Rueckert, D.: Hierarchical statistical shape analysis and prediction of sub-cortical brain structures. Medical Image Analysis 12 (2008)
10. Zhou, K., Comaniciu, D.: Shape regression machine. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 13–25. Springer, Heidelberg (2007)
11. Davis, B., Fletcher, P., Bullitt, E., Joshi, S.: Populations shape regression from random design data. In: IEEE Int. Conf. Computer Vision (2007)
12. NIH-NICD: National institute on deafness and other commnunication disorders: Hearing aids information,
http://www.nidcd.nih.gov/health/hearing/hearingaid.asp
13. Gray, H.: Gray's anatomy of the human body (1918), www.bartleby.com/107/
14. Paragios, N., Rousson, M., Ramesh, V.: Distance tranforms for non-rigid registation. Computer Vision and Image Understanding (CVIU) 89(2-3), 142–165 (2003)
15. Press, W.H., Teukolsky, S.A., Vetterling, W., Flannery, B.P.: Numerical Recipes in C. Cambridge University Press, Cambridge (1992)
16. Whitaker, R.T., Breen, D.E.: Level-set models for the deformation of solid objects. In: Eurographics/Siggraph, Proceedings of Implicit Surfaces, pp. 19–35 (1998)