

Nonparametric joint shape learning for customized shape modeling

Gozde Unal*

Sabanci University, Faculty of Engineering and Natural Sciences, Tuzla 34956, Istanbul, Turkey

ARTICLE INFO

Article history:

Received 23 April 2009

Received in revised form

23 November 2009

Accepted 1 December 2009

Keywords:

Shape estimation

Variational shape optimization

Customized shape modeling

Nonparametric shape density

Joint shape prior

Hearing aid design

Pre-operative and intra-operative shape modeling

ABSTRACT

We present a shape optimization approach to compute patient-specific models in customized prototyping applications. We design a coupled shape prior to model the transformation between a related pair of surfaces, using a nonparametric joint probability density estimation. The coupled shape prior forces with the help of application-specific data forces and smoothness forces drive a surface deformation towards a desired output surface. We demonstrate the usefulness of the method for generating customized shape models in applications of hearing aid design and pre-operative to intra-operative anatomic surface estimation.

1. Introduction

Customized shape modeling is one of the most important tasks in virtual prototyping of implants and anatomical modeling of organs. Some applications in customized prototyping are realized for providing patient-specific models for bones, teeth, and hearing aid devices. Estimating geometry of organs before, during, and after surgery and modeling bilateral relationships between the symmetric structures in the brain are other applications. In problems of customized shape design, the first step involves acquiring a rough raw data of the structure to be modeled. For instance, impressions of the teeth or the ear canal can be acquired via a laser scan; or 3-dimensional (3D) computed tomography (CT), magnetic resonance (MR) image volumes of the patient are obtained and 3D models of the organs are extracted from these images. In the next step, certain operations and rules are applied to the model with constraints from the input patient data, which is at this point an unprocessed 3D surface model. If one would like to predict the intra-operative shape of a particular organ, for instance prostate, a natural approach would be to acquire a pre-operative CT or MR scan of the patient and extract the prostate surface, and apply a learned transformation, to the pre-operative surface to generate the intra-operative

geometry of the prostate. A similar problem arises in hearing aid design: one obtains an impression of the ear geometry, and estimates the final hearing aid device, which should comfortably fit to a patient's ear as well as satisfy other constraints in terms of the electronics components, ventilation vents, and so on. An example problem is depicted in Fig. 1 for illustration, where a complete tooth model is shown along with separate root and crown parts: in dental implants, the decayed or fractured crown part of the tooth is covered with a prosthetic crown, whose shape should be customized for a good fit. We achieve a solution to such problems via a shape optimization method that jointly models multiple shapes either consisting of different parts of one structure or different phases of a given structure.

1.1. Related work

One of the pioneers of morphometrics, the study of statistical variations on the geometry of biological forms and shapes, is Bookstein [1]. In active shape models of Cootes and Taylor [2], shapes are represented by a point distribution model through the position of the landmarks described by a Gaussian probability density. The problem of modeling shape spaces as a Gaussian probability density via a principal component analysis (PCA) appeared in the context of image segmentation [3] and [4], which were followed by a kernel PCA in [5], and later adopted with nonparametric shape probability distributions in [6]. In this context, the estimated shape probability density from a training set is used as a shape prior term for an image segmentation.

* Corresponding author. Tel.: +90 216 483 9553; fax: +90 216 483 9550.

E-mail address: gozdeunal@sabanciuniv.edu.

URL: <http://vpa.sabanciuniv.edu/people/gozdeunal/>.




Fig. 1. Shape estimation problem exemplified for crown or complete tooth estimation given the root of the tooth.

In the context of object estimation, a multivariate statistical approach was taken in Rao et al. [7] where statistical variation of bilateral structures within the brain were studied via canonical correlation analysis and partial least squares regression. In another work by Davis et al. [8], shape changes in the brain are analyzed through nonparametric regression and exemplified for local shape changes regressed as a function of age. In the context of simultaneous segmentation of multiple objects from images, the problem of joint prior density estimation appears, and is frequently based on a parametric density to capture dependencies among multiple but coupled shapes to improve segmentation performance by providing complementary information as in [9,10], and statistical multi-object modeling via medial shape representations was presented in [11]. Nonparametric shape priors were utilized in [12] for segmentation of multiple objects in images.

We will present a coupled learning approach for the problem of customized shape modeling in order to estimate a customized output surface given as input an anatomical structure of a patient. Our work was inspired by [13], which presented a multivariate regression estimation framework between two related classes of shapes, where the underlying probability of the shape spaces were assumed to be Gaussian. The authors used a covariance analysis for dimensionality reduction, and worked in a reduced space by retaining a small number of principal components of shape variation for both of the paired shapes to estimate the transformation between them in the reduced space. However, anatomic shape variations are not guaranteed to lie on a space modeled by a Gaussian probability distribution, i.e. on a linear manifold where the shapes are represented by adding a number of principal modes of variations to an average shape. Therefore, in this paper, we propose using a non-parametric joint shape probability density to model the spaces of paired shapes coupled through anatomical and patient-specific constraints.

1.2. Our contribution

Our main contribution is the development of building a coupled and non-parametric density estimation between two related shape classes and given an input shape from one of the classes, use the estimated joint density as a shape prior in deforming a template shape towards an expected output shape in the second shape class. For this purpose, our method constructs a conditional probability density to model relationship between two classes of shapes and a surface propagation equation is derived as the gradient of an energy in a variational formulation. Starting from an initial template shape, an output shape is generated from one of the classes related to an input shape from the other class. Such an automatic shape estimation or deformation method can be used in various applications like customized design of anatomical parts.

Due to an expected flexibility in modeling more complex and arbitrary shape deformations, a non-parametric multivariate density estimator is utilized. This approach estimates the shape distribution on a nonlinear manifold rather than a linear one, and

depending on a kernel width parameter and the number of shape samples, constrains it to the manifold defined by the shapes. In addition, we present an application specific data term to constrain deformation of a template surface towards a desired output surface. We exemplify our methodology via two interesting applications: hearing aid design for patients, and pre-operative to intra-operative prostate surface estimation, however, the method can be applied to similar problems in customized anatomic part design, or for symmetry analysis, for instance where structures on right or left part of the brain can be estimated from that of the other side.

2. Methodology

Our problem can be categorized in the patient population analysis framework, where for each patient, a given pair of related shapes are anatomic structures, and for further analysis, pose variation is discarded by aligning the anatomic shapes in a given patient data population. For the initial alignment steps both within the pair of shapes and among all the shapes, we followed the registration method presented in [14]. We utilized an initial rigid registration with 3D rotation and 3D translation parameters, where each shape in the training set was first registered to an arbitrarily chosen shape in the set. After the first step, an average template shape was computed, and all the shapes were registered to the average shape. After the initial alignment step, all the pose variation is removed, and we proceed to the shape deformation step that includes the main components of our approach. The shapes are represented in an Eulerian framework with signed distance functions (SDFs): $\Phi : \mathbb{R}^3 \rightarrow \mathbb{R}$ defined in Eq. (2).¹

In a typical variational setting, we define a cost functional:

$$E(\mathbf{S}) = E_{\text{data}} + E_{\text{coupled shape}} + E_{\text{regularizer}} \quad (1)$$

where E_{data} is a data faithfulness term, $E_{\text{coupled shape}}$ is a coupled shape prior term, and $E_{\text{regularizer}}$ is a regularizer to obtain a feasible and smooth surface in the space of admissible surfaces,² and \mathbf{S} is the unknown shape surface to be estimated, which is embedded in \mathbb{R}^3 . The surface \mathbf{S} is defined as an implicit surface embedded in a signed distance function Φ :

$$\mathbf{S} = \{\mathbf{x} = (x, y, z) \in \Omega \subset \mathbb{R}^3 \mid \Phi(x, y, z) = 0\} \quad (2)$$

with the convention that $\{\Phi(x, y, z) \leq 0\}$ inside and on the surface \mathbf{S} , and $\{\Phi(x, y, z) > 0\}$ outside the surface \mathbf{S} .

2.1. Coupled shape prior term

In scenarios where parametric modeling of density of shapes such as with a Gaussian density is not satisfactory, a classical approach is based on a kernel density estimation, known as non-parametric density estimation [15]. Given a set of N training shapes $\{\Phi_k\}_{k=1, \dots, N}$, one can define a probability density on the space of signed distance functions as follows:

$$P(\Phi) = \frac{1}{N\sigma} \sum_{k=1}^N \exp\left(-\frac{1}{2\sigma^2} d^2(\Phi, \Phi_k)\right) \quad (3)$$

which is a well-studied and widely used kernel density estimator, also known as Parzen window method, that usually entails a

¹ We note that the shape surface and its embedding function, the SDF, is used interchangeably to describe a shape, and should be clear from the context.

² Mathematically speaking, an *admissible* set of functions (here surfaces) is a set on which an extremum problem is defined, here via the energy (cost) functional $E(\mathbf{S})$. A function (surface) that satisfies all the functional constraints imposed by the problem (here the data, the coupled shape and the regularizer constraints) is said to be *feasible*.

Gaussian kernel with a kernel width σ , as is adopted here. It can be shown that this probability estimate converges toward the true probability with $N \rightarrow \infty$ and $\sigma \rightarrow 0$ [16]. The choice of distance d in this work is the L_2 distance among the SDF shape representations:

$$d^2(\Phi, \Phi_k) = \int_{\Omega} (\Phi - \Phi_k)^2 d\Omega. \quad (4)$$

where the SDFs are defined over the domain $\Omega \in \mathbb{R}^3$.

We assume that the shapes are related by a transformation whose general nature are either anatomic deformations, and constraints, and/or rules, which are unknown. Usually, the operations that constitute the transformation are typically operator-dependent, and patient-specific. In order to mathematically describe the resulting anatomic relation between the shapes, one can build a shape prior that captures the coupled relation between shapes under these different possible scenarios. To do so, we utilize the joint nonparametric density estimator [17]:

$$P(\Phi) = \frac{1}{N} \sum_{k=1}^N \prod_{l=1}^L \frac{1}{\sigma_l} \exp\left(-\frac{1}{2\sigma_l^2} d^2(\Phi^l, \Phi_k^l)\right) \quad (5)$$

where Φ_k^l denotes the k 'th training shape for the shape class l , and the multidimensional kernel is the product of uni-dimensional kernels with corresponding kernel widths σ_l . Intuitively, the kernel width σ_l can be selected proportional to the average distance between the shapes in the training set, so that the training shapes are within a multiple of the standard deviation of the underlying shape space [15].

Suppose we have two shape classes which are paired due to a certain relation that exists between them. Then for the specific case of $L = 2$, we have the following joint probability density function of the paired shape (Φ^1, Φ^2) :

$$E_{\text{coupled shape}} = P(\Phi^1, \Phi^2) = \frac{1}{N\sigma_1\sigma_2} \sum_{k=1}^N \exp\left(-\frac{1}{2\sigma_1^2} d^2(\Phi^1, \Phi_k^1)\right) \times \exp\left(-\frac{1}{2\sigma_2^2} d^2(\Phi^2, \Phi_k^2)\right) \quad (6)$$

Utilizing this joint distribution and Bayes theorem, one can write the conditional probability density function for the probability of the surface Φ^2 given a surface Φ^1 . We then define the negative log-likelihood of the conditional probability as our coupled shape prior:

$$P(\Phi^2|\Phi^1) = -\log \frac{P(\Phi^2, \Phi^1)}{P(\Phi^1)} \quad (7)$$

Here, the marginal probability density $P(\Phi^1)$ is assumed to be uniform since any surface to be given is equally likely in the space of the first shape class. Therefore, it counts as a normalizing constant, say η . If prior knowledge about the input shape class exists, e.g. size distribution and so on, this certainly can be incorporated as a prior probability as well.

Introducing a marching parameter, $t \geq 0$, we denote the unknown shape as $\Phi^2(\mathbf{x}, t)$ which is to be initialized with an initial shape $\Phi^2(\mathbf{x}, t = 0)$. Taking the derivative of the negative log-likelihood in Eq. (7) w.r.t. the surface Φ^2 , the propagation term for $\Phi^2(\mathbf{x}, t)$ can be obtained as follows:

$$\frac{\partial \Phi^2(\mathbf{x}, t)}{\partial t} = \sum_{k=1}^N \bar{\gamma}_k(t) (\Phi^2(\mathbf{x}, t) - \Phi_k^2(\mathbf{x})) \quad (8)$$

where $\mathbf{x} \in \Omega$ is a coordinate variable in \mathbb{R}^3 , and the weighting term $\bar{\gamma}_k$ is defined by:

$$\gamma_k(t) = \frac{1}{N\sigma_1\sigma_2\eta} \frac{\exp(-d^2(\Phi^1, \Phi_k^1)/2\sigma_1^2) \exp(-d^2(\Phi^2(t), \Phi_k^2)/2\sigma_2^2)}{\sigma_2^2} \quad (9)$$

$$\bar{\gamma}_k(t) = \frac{\gamma_k(t)}{N} \quad (10)$$

$$\sum_{k=1}^N \gamma_k(t)$$

In the propagation of the unknown surface Φ^2 , the shapes in the second class' data set, i.e. Φ_k^2 are weighted according to their distance to the current estimate of Φ^2 , and an additional coupling comes from the distances between the input shape Φ^1 and the shapes in the first class data set, i.e. Φ_k^1 's. These weights are prescribed by the coefficients $\bar{\gamma}_k(t)$.

2.2. Data term

Data term in an optimization problem generally entails a design based on the given application usually in the form of a similarity or dissimilarity measure between given input data and a model that is to be estimated. We note that in this work, the input data also takes the form of a surface which is a raw unprocessed shape. For instance, given an input raw shape Φ^1 and the estimate of the processed shape Φ^2 that is expected to resemble Φ^1 in certain aspects, we define the following energy functional:

$$E_{\text{data}}(\Phi^2) = \int_{\Omega} f(\Phi^1(\mathbf{x})) H_{\epsilon}(-\Phi^2(\mathbf{x})) d\Omega \quad (11)$$

where $f(\cdot)$ provides a data faithfulness constraint, and H_{ϵ} is a regularized Heaviside function:

$$H_{\epsilon}(x) = \frac{1}{2} \left(1 + \frac{2}{\pi} \arctan\left(\frac{x}{\epsilon}\right) \right), \quad (12)$$

which acts as a smoothed indicator function for the inside of the closed surface of Φ^2 (for the inside region: $H_{\epsilon}(-\Phi^2)$ is used). Here, the ϵ parameter adjusts the steepness of the step to provide a smoothed discontinuity in H . A typical value for this parameter is $\epsilon = 0.5$.

A possible choice for f is that $f(\Phi^1) = \text{sign}(\Phi^1(\mathbf{x}))$ so that the shape Φ^2 minimizes the total sign or maximizes its interior overlap with shape Φ^1 over the domain Ω . Equivalently, the zero level sets of the two surfaces are to be lined up: hence the surface Φ^2 is completely morphed to the surface Φ^1 . Another similar choice for the $f(\cdot)$ function is given by [18] where the target SDF Φ^1 was directly used, i.e. $f(x) = x$.

The gradient descent on a typical energy functional as in Eq. (11) leads to:

$$\frac{\partial \Phi^2(\mathbf{x}, t)}{\partial t} = -f(\Phi^1(\mathbf{x})) \delta_{\epsilon}(-\Phi^2(\mathbf{x}, t)) \quad (13)$$

where δ_{ϵ} is the derivative of the regularized Heaviside function H_{ϵ} in Eq. (12). Another possibility is to replace $\delta_{\epsilon}(\Phi^2)$ by $|\nabla \Phi^2|$ in order to extend the evolution to all level sets of the SDF Φ^2 :

$$\frac{\partial \Phi^2(\mathbf{x}, t)}{\partial t} = f(\Phi^1(\mathbf{x})) |\nabla \Phi^2(\mathbf{x}, t)| \quad (14)$$

Note that $|\nabla \Phi^2(\mathbf{x}, t)|$, which is the L_2 norm of the gradient of Φ^2 , is induced from a geometric flow: $(\partial \mathbf{S} / \partial t) = f \mathbf{N}$, where \mathbf{S} is a parametric surface and \mathbf{N} is the unit normal to the surface, and can be also expressed in terms of the SDF: $\mathbf{N} = (\nabla \Phi^2 / |\nabla \Phi^2|)$. The equivalence between the generic surface evolution equations described mathematically either on \mathbf{S} or the implicit surface Φ naturally follows [19,20]. Here, f denotes the speed of the propagating front Φ .

Next, we construct a more general data speed $g(\Phi^1, \Phi^2)$, which is a function of both the fixed surface Φ^1 , and the time-varying surface Φ^2 , and define the energy functional for this data constraint:

$$E_{\text{data}}(\Phi^2) = \int_{\Omega} g(\Phi^1(\mathbf{x}), \Phi^2(\mathbf{x})) H_{\epsilon}(-\Phi^2(\mathbf{x})) d\Omega, \quad (15)$$

where g is defined as:

$$g(\Phi^1, \Phi^2) = -\Phi^2(\mathbf{x}) \text{sign}(\Phi^1(\mathbf{x})). \quad (16)$$

The Euler-Lagrange equation for (15) then leads to the gradient descent equation:

$$\frac{\partial \Phi^2(\mathbf{x}, t)}{\partial t} = -\text{sign}(\Phi^1(\mathbf{x})) H_{\epsilon}(-\Phi^2(\mathbf{x})) + \Phi^2(\mathbf{x}) \text{sign}(\Phi^1) \delta_{\epsilon}(-\Phi^2(\mathbf{x})). \quad (17)$$

For instance, in the hearing aid shell design application, the processed output shape Φ^2 should fit to the input shape, i.e. to the patient's anatomy comfortably. Therefore, the output shape should stay inside the input shape in certain anatomical regions, such as the ear canal. The second term in (17) is adequate for our purpose, i.e., to prevent the Φ^2 surface to propagate outside Φ^1 surface. Therefore, the designed evolution term becomes:

$$\frac{\partial \Phi^2(\mathbf{x}, t)}{\partial t} = \Phi^2(\mathbf{x}) \text{sign}(\Phi^1(\mathbf{x}, t)) \delta_{\epsilon}(-\Phi^2(\mathbf{x})). \quad (18)$$

Inserting the data speed notation, g in (16), and using the geometric flow term $|\nabla \Phi^2|$ instead of the $\delta_{\epsilon}(-\Phi^2)$ term, we obtain:

$$\frac{\partial \Phi^2(\mathbf{x}, t)}{\partial t} = g(\Phi^1(\mathbf{x}), \Phi^2(\mathbf{x}, t)) |\nabla \Phi^2(\mathbf{x}, t)|. \quad (19)$$

In the speed g , a regularized version of the sign function is utilized: $\text{sign}(\Phi^1(\mathbf{x})) = H_{\epsilon}(\Phi^1(\mathbf{x})) - H_{\epsilon}(-\Phi^1(\mathbf{x}))$.

To further improve the data constraint, more sophisticated rules that depend on more complex features can be incorporated. For instance, an anatomical surface atlas, which partitions the shape surface into different anatomical labels can be employed as a spatial mask over the deformation of the unknown shape Φ^2 .

2.3. Smoothing term

As a regularizer, standard surface area penalty is utilized so that the deforming surface does not grow its area unnecessarily large and become irregular and rugged, which is undesirable for most of the anatomic surfaces. The smoothing term in this case is typically based on the mean curvature function κ of the surface as its speed function:

$$\frac{\partial \Phi^2(\mathbf{x}, t)}{\partial t} = \kappa(\mathbf{x}) |\nabla \Phi^2(\mathbf{x}, t)|. \quad (20)$$

Note that in the absence of shape priors in a variational setting, using curvature function as the speed function of the surface can be considered as the most basic shape prior for a surface, as an anatomic structure is generally not expected to have an uneven and nonuniform surface.

2.4. Overall deformation

Finally, for our given problem of estimating the processed shape surface given the raw shape surface, the paired shape combination (Φ^1, Φ^2) becomes (Φ^r, Φ^p) , where Φ^r is a raw unprocessed input

surface and Φ^p is the processed version of that surface. Combining the different terms of the surface propagation equations obtained above, we can write the overall shape deformation equation as:

$$\begin{aligned} \frac{\partial \Phi^p(\mathbf{x}, t)}{\partial t} &= \alpha \text{ Coupled Shape Prior} + \beta \text{ Data term} \\ &\quad + \zeta \text{ Smoothness term} \\ \frac{\partial \Phi^p(\mathbf{x}, t)}{\partial t} &= \alpha \sum_{k=1}^N \tilde{\gamma}_k(t, \Phi^r, \Phi^p) (\Phi^p(\mathbf{x}, t) - \Phi_k^p(\mathbf{x})) \\ &\quad + [\beta g(\Phi^p(\mathbf{x}, t), \Phi^r(\mathbf{x})) + \zeta \kappa(\mathbf{x}, t)] |\nabla \Phi^p(\mathbf{x}, t)| \end{aligned} \quad (21)$$

$$\Phi^p(\mathbf{x}, t=0) = \Phi_{\text{mean}}^p(\mathbf{x}) \quad (22)$$

where α, β, ζ are the weights corresponding to the coupled shape prior term, the data term, and the smoothness term, respectively, and $\tilde{\gamma}_k$ are given by

$$\tilde{\gamma}_k(t, \Phi^r, \Phi^p) = \frac{1}{N} \frac{\exp(-d^2(\Phi^r, \Phi_k^r)/2\sigma_r^2) \exp(-d^2(\Phi^p(t), \Phi_k^p)/2\sigma_p^2)}{\sum_{k=1}^N \exp(-d^2(\Phi^r, \Phi_k^r)/2\sigma_r^2) \exp(-d^2(\Phi^p(t), \Phi_k^p)/2\sigma_p^2)}. \quad (23)$$

Eq. (21) is solved until it reaches a steady state with the initial condition (22), which starts with the mean SDF surface of the processed output shape class:

$$\Phi_{\text{mean}}^p = \frac{1}{N} \sum_{k=1}^N \Phi_k^p, \quad (24)$$

where Φ_k^p are the corresponding training shapes. The initial template is deformed and converges to an estimate of the unknown final surface Φ^p .

The PDE in (21) provides an update equation for the SDF Φ^p of the surface, and in order to retain the properties of a signed distance function, it should be re-initialized periodically. This can be achieved either by re-computing the distance function around the zero-level set of the current SDF, or by the initial value PDE suggested in [19]:

$$\frac{\partial \Phi^p(\mathbf{x}, t)}{\partial t} = \text{sign}(\Phi^p(\mathbf{x})) (1 - |\nabla \Phi^p(\mathbf{x}, t)|). \quad (25)$$

2.4.1. Implementation

Finite differences are used for solving the initial value partial differential equation (PDE) (21) along with Neumann Boundary Conditions at the volumetric grid borders. The upwind differencing scheme [19] is used both for the coupled shape prior term, the first on the right hand side of the PDE (21), which is a source term, and for the data term, which is of advection form.

For the curvature term, central differences are used in discretizing the mean curvature function at each coordinate \mathbf{x} :

$$\kappa = \frac{(\Phi_{yy} + \Phi_{zz})\Phi_x^2 + (\Phi_{xx} + \Phi_{zz})\Phi_y^2 + (\Phi_{xx} + \Phi_{yy})\Phi_z^2 - 2\Phi_x\Phi_y\Phi_{xy} - 2\Phi_x\Phi_z\Phi_{xz} - 2\Phi_y\Phi_z\Phi_{yz}}{(\Phi_x^2 + \Phi_y^2 + \Phi_z^2)^{3/2}}. \quad (26)$$

The time step in discretization of the PDE (21) was set to $\Delta t = 0.25$. The maximum number of iterations was set to 100, which ensured that the shape updates have converged to the final steady state. The SDF was re-initialized every 20 iterations by re-computing the distance function around its zero-level set.

3. Results

The proposed shape modeling algorithm is tested on two different datasets. The first application is the estimation of a hearing aid shell from a patient raw ear impression. The second application is

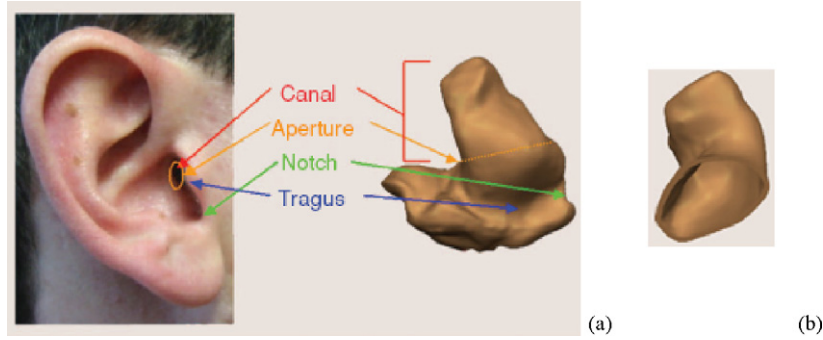


Fig. 2. (a) Raw ear impression surface; (b) hearing aid shell resulting from processing, i.e. detailing of the raw surface. (Pictures: adapted from [21].)




Fig. 3. (a) An input shape (raw ear impression) sample from the dataset; (b) corresponding detailed shape (ground truth); (c) (a and b) rendered together with transparency.

the estimation of the intra-operative geometry of an organ given the pre-operative surface, here exemplified for the prostate shape.

3.1. Hearing aid shell design

An example is provided in Fig. 2 where the initial surface is acquired originally through a laser scanner in the form of a point cloud and then triangulated to form a mesh model of the ear canal and surrounding outer regions of the ear.

In the hearing aid dataset, there are 43 input and output surface pairs, i.e. $N = 43$. The original data is saved in an stl format as triangulated meshes, and we converted them into a voxelized format to represent each surface as a signed distance function embedded in a $100 \times 100 \times 100$ voxel grid. The reason behind this choice is that partial differential equations we proposed in this paper involve local surface deformations, which are naturally solved in such an implicit representation avoiding the problems of possible mesh irregularity after updates of mesh vertices. A sample pair from the hearing aid dataset: an input shape and its corresponding output shape produced by a specialist are shown in Fig. 3. The output surface fits to mainly the canal part of the input surface. The goal is then given the input full surface model, to estimate the smaller output hearing aid shell which should house necessary components for the device as well as conform comfortably to a patient's ear.

In Fig. 4, we compare the influence of different terms in the deformation equation (21) for the estimation of the hearing aid shell from a patient's raw ear impression data. When only the data term is utilized as in (a), the resulting surface is not constrained by any shape prior forces therefore continues to deform towards the input model or the observation data, which is the complete input raw ear model, to produce a very distant result.

Note that as an initial condition for the PDE (21), a template shape was computed as the average of all ground truth output shapes in the training dataset in Eq. 24, which was slightly eroded as well (see Fig. 6a). On the other hand, the coupled shape prior term alone will constrain the deforming shape to the canal region and conform to similar surfaces in the training set due to the weighted averaging in the equation, however, when the distance between the initial surface and the expected result is large, it alone may not be enough to fully drive the surface to the true solution (Fig. 4b).

If there are samples in the data set, whose shapes are similar to the given test shape, the sole non-parametric prior would also produce finer details, however, as the existence of close models is usually unknown, the data term, which depends on the input data, i.e. the observation, ensures that the model is pulled to the specific fine details. When both the coupled shape prior term and the data term are utilized as in Fig. 5, the final estimated surface correctly mimics the ground truth surface generated by the specialist (in blue).




Fig. 4. Resulting surface (in red): (a) data term only (no coupled shape prior) rendered alone and with ground truth shape (in blue); (b) shape term only (no data term) rendered alone and with ground truth shape (in blue). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)




Fig. 5. Resulting surface with both data term and coupled shape prior term (red), shown together with the ground truth shape (blue) from two different render view points. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)




Fig. 6. (a) Template surface used as the initial condition for the deformation equation (21), which is computed as the average of all registered ground truth detailed shapes in the dataset; (b) initial condition surface visualized together with the deformation result (also rendered in Fig. 5).

Through all the experiments for the hearing aid application in this paper, we fixed the weights in Eq. (21) as follows: the coupled shape prior weight $\alpha = 1$, the data weight $\beta = 1$, the smoothness weight $\zeta = 1$. The Gaussian kernel widths σ_r^2 and σ_p^2 in Eq. (23) were both set to 1% of the average distance between all surfaces in input and output training sets, respectively.

3.1.1. Validation studies

For quantitative validation of the results, leave-one-out tests over 43 ear impression shape set are carried out. The averaged results for the Dice measure, which is a kind of overlap measure, and the median absolute distance between the estimated shape and the ground truth shape are displayed in Table 1. The Dice measure between two shapes A and B represented in the voxel domain is defined as $2\#(A \cap B) / (\#A + \#B)$, where $\#$ denotes the number of voxels on and inside a shape A. All shapes in our experiments are defined as signed distance functions over a $100 \times 100 \times 100$ grid. The median absolute distance between two shapes is multiplied by 0.45 mm which is the size of one voxel cube during triangulated mesh to voxel domain conversion. The Dice overlap measure is calculated as $85 \pm 4.9\%$, reported as mean \pm standard deviation values, over 43 separate leave-one-out experiments, where in each experiment one shape was left out as the test shape and the rest were used as the training shapes. In the same experiments, the median absolute distance is calculated as 0.47 ± 0.27 mm.

We show qualitative results for the experiments on 43 ear impression surfaces in Fig. 7, where the estimated hearing aid shells are observed as superimposed over the input surfaces and with the ground truth shapes. As demonstrated by the quantitative results in Table 1, and the visual observations, the estimated shells are in good agreement with the expected shells, particularly in the canal region, and the interaction between the coupled shape prior, the

Table 1

Mean \pm standard deviation values of the Dice measure and the median absolute distance between the estimated surface and the ground truth surface for the hearing aid shell dataset.

Hearing aid shell dataset	Dice overlap (%)	Median absolute distance (mm)
43 leave-one-out tests	85.01 ± 4.90	0.47 ± 0.27

Table 2

Mean \pm standard deviation values of the Dice measure and the median absolute distance between the estimated surface and the ground truth surface using the proposed approach, compared for different number of training shapes in the hearing aid shell dataset.

Number of training samples	Dice overlap (%)	Median absolute distance (mm)
Tests with $N = 33$ shapes	84.53 ± 6.29	0.53 ± 0.37
Tests with $N = 23$ shapes	84.26 ± 7.32	0.58 ± 0.43
Tests with $N = 13$ shapes	83.10 ± 7.43	0.67 ± 0.44

data, and the smoothness forces aid each other to drive the initial template surface to the correct solution.

3.1.2. Effect of training sample size

Theoretically, by using kernel density estimation functions (Eqs. (5) and (6)), a large number of samples are required. In real clinical applications, however, the number of training shape instances is usually limited. We present some results with different number of training samples to explore effect of number of training shapes over the proposed method. The validation studies were carried on the full data set with $N = 43$ samples in Section 3.1.1, therefore, we tested with the following three different sample sizes: $N = 33, 23$, and 13 in Eq. (21). The same performance measures of Table 1 are reported in Table 2. The results depicted show that when compared to those in Table 1, as the number of shapes in the sample set is reduced, a slight decrease in both the Dice measure and an increase in the distance between the shapes is observed, as expected.

One should cautiously interpret these results since the total number of shapes in our data set were limited, and with much larger datasets, substantial increase or decrease in the number of shapes may cause more significant performance changes. However, still, for a well-sampled data set, which contains a range of expected variations in a given anatomical shape class, the obtained results indicate that the coupled shape prior term, even with a smaller number of training shapes, helps constrain the data term adequately to converge to a reasonable expected output shape. Note that the fact that by construction, the initial surface will always completely deform towards the input raw surface, e.g. the full ear impression, by the data term without the coupled shape prior term, points to importance of employing the latter, even with a small number of available training shapes.

3.1.3. Comparison

We have implemented part of the technique by [13] to which we compare our method. Their method applies traditional statistical modeling approach to both shape spaces, i.e. the input surfaces and the output surfaces. After projecting the input and output shapes onto a shape space modeled with Gaussian probability densities, the method in [13] finds a transformation T between these two shape spaces. Given a new input shape, the transformation T is used to transfer the input shape into the desired output shape related to the input shape. Let us call this the direct shape estimate as a result of using the traditional statistical model. As this technique assumes an inherent Gaussian density, and the class of shapes are not guaranteed to lie on a linear space, i.e. cannot be exactly expressed by adding modes of variations around an average template shape, the authors have proposed additional steps by estimating an auxiliary class of shapes, referred to as the mask shapes, using the difference shape between the input and the output shape. This is followed by fitting planes to the estimated auxiliary shape and finally deforming the direct shape estimated in the first stage using the plane constraints. Here, for a fair comparison, we implement and use the direct shape estimate without further constraints to compare with our approach.




Fig. 7. Visualizations from ear impression/hearing aid surface experiments (left-to-right): input surface (green), ground truth (blue), surface estimate result (red), along with the ground truth, and with the input surface depict highly overlapping, close to true solutions. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

Table 3

Mean \pm standard deviation values of the Dice measure and the median absolute distance between the ground truth surface and the direct surface estimate using a more traditional statistical shape modeling approach [13] for the hearing aid shell dataset.

Traditional stat. approach	Dice overlap (%)	Median absolute distance (mm)
10 leave-one-out tests	75.45 ± 7.52	1.16 ± 0.50

Table 3 presents the results of 10 test shapes that were not used in the training of the traditional statistical approach. The computed Dice measure was 75.45 ± 7.52 (mean \pm std deviation), and the median absolute distance between the ground truth surface and the direct shape estimate was 1.16 ± 0.50 mm. As $N = 33$ shapes were used in the data set, the comparison was done against the corresponding experiment with the same sample size in the proposed method. It can be observed from Table 2 (first row) and Table 3 that the non-parametric probability density modeling of the input and output shape spaces achieves closer results to the expected shape. Fig. 8 depicts some visualizations from the results obtained using this approach, which qualitatively shows the degraded performance. The resulting estimated surfaces (red) display the expected character of a certain number of variations added around an average shape, and although they depict the global characteristics of the

ground truth, i.e. the desired output shape (gray), they fall slightly short of producing the local variations of the desired shapes. This is a natural behavior expected from using the Gaussian probability density modeling over the shape spaces. This was the very reason to use a nonparametric probability density modeling of the shape spaces as proposed in this paper.

3.2. Pre-op to intra-op anatomic shape estimation

In surgical planning and medical interventional guidance, pre-operative and intra-operative information and observations over a patient's anatomy are crucial. To improve the steps of surgical processes, usually high resolution pre-operative (pre-op) medical images of the patient are acquired to create a detailed planning model for the organ to be operated and its surrounding regions. During the surgery, fast intra-operative (intra-op) images of the patient are often acquired, and registered to and fused with the pre-op information to help with the surgical navigation and intervention. In situations where an intra-op model of the organ to be operated is not available (since an intra-op scan is not required in some operations), an analysis of the morphometric changes in an organ before and during surgery may be useful. In this paper, as a second application of the method we proposed, we model how the prostate surfaces deform from a pre-op to intra-op scenario.




Fig. 8. Shape estimation results using Gaussian density modeling of shape spaces, (in each box): input surface (blue); surface estimate result (red) overlaid transparently on ground truth surface (gray), as expected falls a little short of modeling nonlinear local variations in shapes as opposed to the nonparametric modeling approach proposed. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

In Fig. 9, samples from the prostate dataset which included pre-op/intra-op shape pairs: the pre-op surfaces (visualized in green), the intra-op surfaces (visualized in blue), and both surfaces superimposed are depicted. It can be observed that the pre-op surface has mainly blob-like features and the intra-op surface has mainly heart-like features with a more pointed-tip at the bottom and flat geometry at the top. It is anatomically known that prostate is a highly deformable organ, therefore modeling the deformations of the prostate and directly relating its pre-op state to its intra-op state is a challenging problem.

In this paper, we proposed mainly a probabilistic mathematical approach to describe the relation between inter-related shapes. For the prostate application, the coupled shape probability constraint we presented will produce a likely intra-op surface given a pre-op prostate surface. The prostate dataset originally obtained from Harvard’s Surgical Planning Laboratory included about 30 pre-op and intra-op triangulated meshes with 1026 vertices, which are pre-aligned. We voxelized the meshes and obtained signed distance function representations of the prostate surfaces over a $100 \times 100 \times 100$ grid. We then applied our proposed deformation equation in (21) with the weights in the differential equation fixed to: $\alpha = 1$, $\beta = 0$, $\zeta = 1$. The calculated mean surface of the intra-op surfaces was used as the initial template surface, which was propagated towards a final solution. The output shape, i.e. the intra-op surface, is not directly constrained by the input shape, i.e. the pre-op surface, but undergoes a form of expected deformation. Therefore, a natural and direct data constraint term as in the hearing aid application in Section 3.1 was not designed for the prostate shape application. The corresponding weight β was set to zero to lift any input data influence over the deformation. Here, as the shape term is the only influential force that drives an initial surface towards the desired result, the two parameters, the Gaussian kernel size in the kernel density estimator, i.e. the kernel widths in Eq. (23) gain importance. After experimenting with various values of kernel widths, they were set to 1% and 0.5% of the average distance between all surfaces in the dataset for the pre-op surface kernel width σ_p^2 and the intra-op surface kernel width σ_p^2 , respectively.

Fig. 10 shows example intra-op surfaces obtained as a result of the proposed deformation (21) rendered separately and superim-




Fig. 9. Input and output surfaces from prostate dataset (left-to-right): pre-op surface (green), intra-op (blue), and depicted together. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)




Fig. 10. Sample surface estimation results from prostate dataset (left-to-right): estimated intra-op surface (red), together with ground truth intra-op (blue) from two different render view points. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

posed with the ground truth, which is the true intra-op surface extracted as a result of manual segmentation from computed tomography image data. The final surfaces are mostly in agreement with the ground truth surfaces. Table 4 presents quantitative results for 23 leave-one-out tests over the dataset. We note that the two shape classes, i.e. the pre-op and intra-op already overlapped considerably compared to the previous hearing aid application: the Dice overlap and distance measures between the input and output shape classes were $84 \pm 4.19\%$, and 2.29 ± 0.71 , respectively. Furthermore, the initial template shape calculated as the mean surface of all intra-op shapes also overlapped at 87% on average with the true surfaces. As a result of applying our method, the calculated Dice overlap measure between the estimated and the true intra-op surface showed an increase to $92 \pm 2.51\%$, mean \pm standard deviation, and the median absolute distance was reduced to 1.25 ± 0.36 voxels.

3.3. Discussion

We show typical problems in the surface estimation of the hearing aid shells in Fig. 11. On the given typical samples, it can be observed that the surface has slightly propagated past the top and bottom ends of the canal (Fig. 11: second from the right). Most of the errors are concentrated around these regions where a tight data constraint does not exist but mainly a probabilistic shape constraint drives the solution. A solution to such a problem can be to incorporate further anatomical constraints through features

Table 4

Mean \pm standard deviation values of the Dice measure and the median absolute distance between the estimated surface and the ground truth surface for the prostate dataset.

Prostate pre-op/intra-op dataset	Dice overlap (%)	Median absolute distance (voxels)
Pre-op surface to intra-op	84 ± 4.19	2.29 ± 0.71
23 leave-one-out test results to intra-op	92 ± 2.51	1.25 ± 0.36




Fig. 11. Hearing aid data: typical problems in the surface estimates (in red) appear in the canal tip and bottom ends, rendered along with the ground truth (in blue) hearing aid shell surfaces, and input raw ear impression (in green). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)




Fig. 12. Prostate data: problems in the surface estimation: estimated intra-op surfaces (in red) rendered along with the ground truth intra-op surfaces (in blue) and pre-op surfaces (in green). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

extracted over the given anatomic shape. An anatomic region segmentation, a surface atlas, on the ear impression surface could provide restricted propagation and limit the amount of overflow from both sides. Nevertheless, we should note that the coupled shape prior term we presented generally restrained the deforming surface successfully to the canal region and produced a likely feasible shape as validated by the quantitative results in Section 3.1.1.

In the prostate application, problems arise due to the fact that a regularly defined relation between the pre-op and intra-op surfaces of the prostate does not always exist, as prostate can arbitrarily deform in some situations. Some problems in the surface estimation of the intra-op prostate are exemplified in Fig. 12. On the first row, the intra-op surface (blue) does not exhibit the pointed-tip feature at the bottom of the expected heart-like shape, however, the shape prior term drives it past the true geometry (blue) expecting such a pattern. The situation on the second row is worse, as the true intra-op surface has an upward bulging unlike most of the other intra-op surfaces in the dataset. The coupled shape prior term drives the template towards again an expected heart-like geometry, which is flat at the top and pointed at the bottom (surface in red).

Another limitation here is that in both applications, the voxelization operation we carried out has decreased the resolution of the shapes to a 100^3 voxel grid, which optimized our memory and computation constraints. A typical estimation through our defor-

mation equation implemented in MatlabTM currently takes a few minutes. The resolution and computational speed can be increased with more optimized coding and hardware.

On a final remark, we note that our gradient descent solution to the energy functional, which is composed of a coupled shape prior, a data term, and a smoothness term, is non-convex, hence produces locally optimal solutions. Furthermore, the energy terms are combined in a weighted fashion, hence there is an inter-play among the different constraints. For the hearing aid experiments, we fixed the weights in Eq. (21): the coupled shape prior weight $\alpha = 1$, the data weight $\beta = 1$, the smoothness weight $\zeta = 1$, and for the prostate shape application the weights were fixed to: $\alpha = 1$, $\beta = 0$, $\zeta = 1$. The two other parameters, the kernel widths, were also fixed for both applications as reported in Sections 3.1 and 3.2. One can of course explore the space of parameters to further fine tune the parameters for a specific application.

4. Conclusion

This paper presents a method that uses a coupled shape modeling for automatic generation of a surface that belongs to an output shape class given a surface from a related input shape class. A kernel density estimator approach is utilized for modeling the shape spaces via a joint non-parametric probability density estimator, and a conditional probability density function is derived for the prediction problem also constrained by the data terms and other

regularization terms such as the mean curvature flow. The main idea in this paper can be expanded or augmented by adding other prior knowledge about the process or anatomic information such as features extracted on a pair of input and output surfaces, which can improve the performance. Further large scale testing for rapid prototyping applications are required for an extensive validation of the methodology, however, our studies have demonstrated the potential of our method to be successfully used in such applications.

Acknowledgements

The author thanks Dr. Tong Fang and Alex Zouhar at Siemens Corporate Research (SCR), Princeton NJ, USA for providing the hearing aid dataset. She also thanks Dr. Steve Haker at Harvard's Surgical Planning Laboratory, MA, USA, for the prostate dataset provided through the grant U41RR019703.

References

- [1] Bookstein F. Size and shape spaces for landmark data in two dimensions. *Statistical Science* 1986;1(2):181–242.
- [2] Cootes T, Taylor C, Cooper D, Graham J. Active shape models—their training and application. *Computer Vision and Image Understanding* 1995;61(January 1):38–59.
- [3] Leventon M, Grimson W, Faugeras O. Statistical shape influence in geodesic active contours. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 1. 2000.
- [4] Tsai A, Yezzi A, Wells W, Tempany C, Tucker D, Fan A, et al. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging* 2003;22(February 2):137–54.
- [5] Rathi Y, Dambreville S, Tannenbaum A. Statistical shape analysis using kernel PCA. In: *IS&T, SPIE Symposium on Electronic Imaging*. 2006.
- [6] Rousson M, Cremers D. Efficient kernel density estimation of shape and intensity priors for level set segmentation. In: *Proc. MICCAI—Medical Image Computing and Computer Assisted Intervention*. 2005. p. 757–64.
- [7] Rao A, Aljabar P, Rueckert D. Hierarchical statistical shape analysis and prediction of sub-cortical brain structures. In: *Medical Image Analysis*, vol. 12. 2008.
- [8] Davis B, Fletcher P, Bullitt E, Joshi S. Populations shape regression from random design data. In: *IEEE Int. Conf. Computer Vision*. 2007.
- [9] Tsai A, Wells W, Tempany C, Grimson W, Willsky A. Mutual information in coupled multi-shape model for medical image segmentation. *Medical Image Analysis* 2004;8(4):429–45.
- [10] Yang J, Duncan J. Joint prior models of neighboring objects for 3-d image segmentation. In: *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, vol. 1. 2004. p. 314–9.
- [11] Lu C, Pizer S, Joshi S, Jeong J-Y. Statistical multi-object shape models. *International Journal of Computer Vision* 2007;75(3):387–404.
- [12] Uzunbas G, Cetin M, Unal G, Ercil A. Coupled nonparametric shape priors for segmentation of multiple basal ganglia structures. In: *IEEE Int. Symposium on Biomedical Imaging*. 2008.
- [13] Unal G, Nain D, Slabaugh G, Fang T. Customized hearing aid design using statistical shape learning. In: *MICCAI—Medical Image Computing and Computer Assisted Intervention*. 2008.
- [14] Zouhar A, Fang T, Unal G, Slabaugh G, Xie H, McBagonluri F. Anatomically-aware, automatic and fast registration of 3d ear impression models. In: *Third International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT)*. 2006.
- [15] Silverman B. *Density estimation for statistics and data analysis*. London: Chapman and Hall; 1992.
- [16] Duda R, Hart P. *Pattern classification and scene analysis*. New York: Wiley; 1973.
- [17] Erdogmus D, Jenssen R, Rao Y, Principe J. Multivariate density estimation with optimal marginal parzen density estimation and gaussianization. In: *IEEE Workshop on Machine Learning for Signal Processing*. 2004.
- [18] Whitaker RT, Breen DE. Level-set models for the deformation of solid objects. In: *InEurographics/Siggraph, Proceedings of Implicit Surfaces*. 1998. p. 19–35.
- [19] Sethian J. *Level set methods and fast marching methods*. Cambridge University Press; 1999.
- [20] Chan T, Vese L. Active contours without edges. *IEEE Transactions on Image Processing* 2001;10(2):266–77.
- [21] Slabaugh G, Fang T, McBagonluri F, Zouhar A, Melkisetoglu R, Xie H, et al. 3d shape modeling for hearing aid design. *IEEE Signal Processing Magazine* 2008;(September):98–102.

Gozde Unal received her Ph.D. degree in electrical and computer engineering from North Carolina State University, Raleigh, NC, USA, in 2002, and then was a post-doctoral fellow at the Georgia Institute of Technology, Atlanta, GA, USA, for a year. Between 2003 and 2007, she worked as a research scientist at Siemens Corporate Research, Princeton, NJ, USA. She joined the faculty of Sabanci University, Istanbul, Turkey in Fall 2007, where she is currently an assistant professor. Her current research is focused on medical image analysis, segmentation, registration, and shape analysis techniques with applications to clinically relevant problems in MR, CT, US, and intravascular images. She is a Senior Member of the IEEE, and an Associate Editor for *IEEE Transactions on Information Technology on Biomedicine*.