

Shape Modelling Using Markov Random Field Restoration of Point Correspondences

Rasmus R. Paulsen^{1,2} and Klaus B. Hilger²

¹ Oticon Research Centre Eriksholm
Kongevej 243, DK-3070 Snekkersten, Denmark
<http://www.oticon.com/>

² Informatics and Mathematical Modelling, Technical University of Denmark
IMM, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, Denmark
{rrp,kbh}@imm.dtu.dk, <http://www.imm.dtu.dk/>

Abstract. A method for building statistical point distribution models is proposed. The novelty in this paper is the adaption of Markov random field regularization of the correspondence field over the set of shapes. The new approach leads to a generative model that produces highly homogeneous polygonized shapes and improves the capability of reconstruction of the training data. Furthermore, the method leads to an overall reduction in the total variance of the point distribution model. Thus, it finds correspondence between semi-landmarks that are highly correlated in the shape tangent space. The method is demonstrated on a set of human ear canals extracted from 3D-laser scans.

1 Introduction

Point distribution models (PDMs) are widely used in modeling biological shape variability over a set of annotated training data [1, 2]. The generative models are highly dependent on the initial labeling of corresponding point sets which is typically a tedious task. Moreover, the labeling is often erroneous and sparse. A good representation of the training data is particularly hard to obtain in three dimensions. Finding a basis of homologous points is thus a fundamental issue that comes before generalized Procrustes alignment [3] and decomposition [4] in the shape tangent space.

A method for building a statistical shape model of the human ear canal is presented in [5]. An extension to this method is proposed in this paper using Markov Random Field (MRF) regularization for improving the initial set of point correspondences. The new approach leads to a more compact representation and improves the generative model by better reconstruction capabilities of the 3D training data. Related work include the application of Geometry Constrained Diffusion (GCD) [6, 7] and Brownian Warps [8] for non-rigid registration. A more compact model is obtained, since the shape tangent space residuals of the new representation have increased correlation. It thus indicates that a better correspondence field is obtained between the 3D semi-landmarks. Related work




Fig. 1. Left: An example of a surface representation of an ear canal with the anatomical landmarks and the separating planes that defines the region of interest. The thin tubular structure in the top is the actual canal. The larger lower section is the concha, of which only the upper part is of interest. A cutoff plane through the concha is therefore defined. Right: The model mesh, shown by a wireframe, fitted to a target shape using Thin Plate Spline warping.

on obtaining a minimum description length of PDMs is proposed in [9, 10] based on information theoretic criteria.

The data consists of 29 3D ear canal surfaces extracted from laser scans of ear impressions. The local surface geometry of the ear canals varies much from one individual to another. Therefore, only very few ridges and extremal points are stable when comparing groups of ear canals. A set of 18 anatomical landmarks of varying confidence are placed on each ear canal, and constitute a sparse correspondence between the surfaces of the ear canals in the training set. The surfaces of the ear canals are not closed due to the opening of the ear canal and because the ear impressions are terminated in front of the ear drum. It is therefore necessary to identify the region of interest of each ear canal. Hence, planes are defined, which separates the valid parts of the surface from the invalid parts. In Fig. 1, left, an ear canal with the anatomical landmarks and separating planes is shown.

The remaining paper is organized in three additional sections. Section 2 describes the proposed statistical method for improving the point correspondences. Section 3 presents the results of applying the extended algorithm. In Section 4 we summarize and give some concluding remarks.

2 Methods

2.1 Surface Correspondence Using Thin Plate Spline Warping

The anatomical landmarks do not constitute an exhaustive description of the surface of the ear canal. It is therefore necessary to generate a more dense set of landmarks describing the shape. For that purpose a model mesh is constructed and fitted to all shapes in the training set. The model mesh is chosen as a decimated version of a natural well-formed ear canal labeled with the anatomical landmarks. The model mesh is fitted to each of the shapes in the training set

using a Thin Plate Spline (TPS) warp based on the corresponding anatomical landmarks. TPS is a warp function that minimizes the bending energy [11]. Since the TPS transform is exact only for the anatomical landmark locations, the vertices of the model mesh will not lie on the surface of the target shape, see Fig. 1, right. Projecting each vertex in the warped model mesh to the closest point on the target surface produces a non-rigid deformation field and generates a dense correspondence. However, using the Point to Surface Projection (PSP) introduces a critical risk of inversions, where the vertices of the model mesh shift place and cause folds in the mesh. Another secondary artifact is the non-uniformity of the correspondence vector field shown in Fig. 2a,b giving rise to poor reconstruction of the target shape. In order to improve the correspondence vector field and avoid the problems inherent in applying point to surface projection a regularization must be included. Lorenz and Krahnstöver [12] propose a method for relaxing a polygonization into a more homogeneous representation, however, such methods are not suited when the polygonization is constrained to an underlying correspondence field. We propose to relax the problem by using a stochastic approach described in the following.

2.2 Markov Random Field Regularization

To obtain better reconstruction and correspondences we cast the problem of finding the deformation vector field into a Bayesian framework of MRF restoration. We thus follow the four successive stages of the Bayesian paradigm.

- 1: Construction of a prior probability distribution $p(\mathbf{d})$ for the deformation field \mathbf{D} matching the source shape \mathbf{S}_s onto the target shape \mathbf{S}_t .
- 2: Formulation of an observation model $p(\mathbf{y}|\mathbf{d})$ that describes the distribution of the observed shapes \mathbf{Y} given any particular realization of the prior distribution.
- 3: Combination of the prior and the observation model into the posterior distribution by Bayes theorem

$$p(\mathbf{d}|\mathbf{y}) = p(\mathbf{y}|\mathbf{d})p(\mathbf{d})/p(\mathbf{y}). \quad (1)$$

- 4: Drawing inference based on the posterior distribution.

We start by some useful definitions from graph theory in order to describe a probability distribution on a spatial arrangement of points.

Given a graph of n connected sites $S = \{s_i\}_{i=1}^n$. A neighborhood system $\mathcal{N} = \{N_s, s \in S\}$ is any collection of subsets of S for which i) $s \notin N_s$, and ii) $r \in N_s \Leftrightarrow s \in N_r$, then N_s are the neighbors of s . A clique C is a subset of sites S for which every pair of sites are neighbors. We use $i \sim j$ to denote that i and j are neighbors. Given a neighborhood system \mathcal{N} on the set of sites S we now consider the probability distribution of any family of random variables indexed by S , i.e. $\mathbf{D} = \{D_s | s \in S\}$. For simplicity we first consider a finite state space $A = 1, \dots, L$ of \mathbf{D} but later generalize to continuous distributions. Let Ω denote the set of all possible configurations $\Omega = \{\mathbf{d} = \{d_i\}_{i=1}^n \mid d_i \in A\}$. A random field

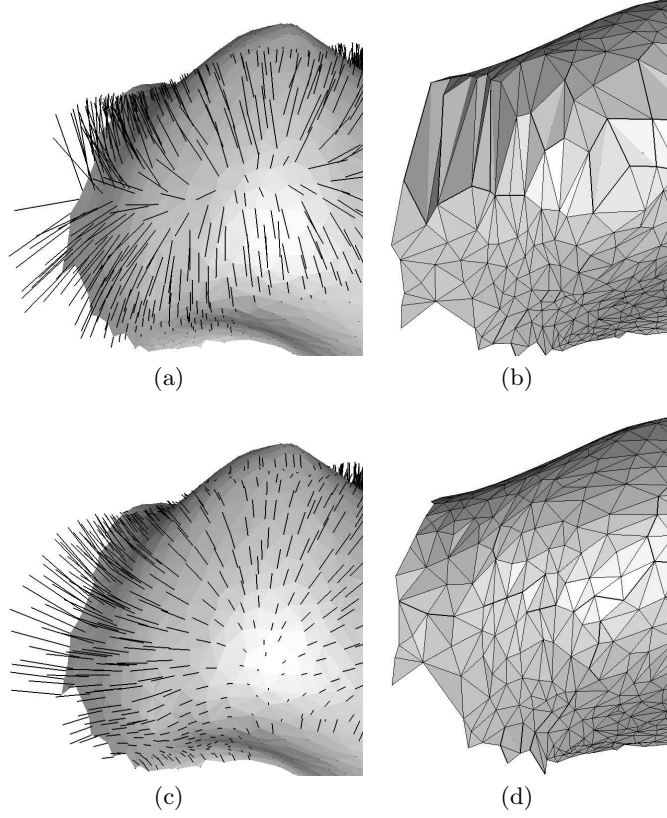


Fig. 2. **a)** The correspondence vector field derived using point to surface projection for moving the vertices of the source to the target shape. **b)** The resulting dense mesh representation of the target shape. **c)** The correspondence vector field derived using the Markov random field restoration of the deformation field for moving the vertices of the source to the target shape. **d)** The improved dense mesh representation of the target shape.

\mathcal{D} is a Markov Random Field (MRF) with respect to \mathcal{N} iff i) $p(\mathbf{d}) > 0 \forall \mathbf{d} \in \Omega$, and ii) $p(d_s | d_r, r \neq s) = p(d_s | d_r, r \in N_s) \forall s \in S, \mathbf{d} \in \Omega$. The first constraint is the positivity condition and can be satisfied by specifying a neighborhood large enough to encompass the Markovianity condition in the second constraint. Although the second condition is on the state of neighboring sites only, it does not exclude long range correlations in the probability distribution over the entire graph. Given a neighborhood system $\mathcal{N} = \{N_s\}$ let all cliques be denoted by \mathcal{C} . For all $C \in \mathcal{C}$ we assume that we have a family of potential functions V_C . We may now define an energy function of any given configuration of \mathbf{d} i.e. $U(\mathbf{d}) = \sum_{C \in \mathcal{C}} V_C$. This leads to the definition of the Gibbs measure. The Gibbs measure induced by the energy function $U(\mathbf{d})$ is $p(\mathbf{d}) = \frac{1}{Z} \exp(-U(\mathbf{d})/T)$, where Z is the partition function and T is a parameter referred to as temperature. The

Gibbs measure maximizes entropy (uncertainty) among all distributions with the same expected energy. The temperature controls the “peaking” of the density function. The normalizing constant may be impossible to obtain due to the curse of dimensionality but often we need only ratios of probabilities and the constant cancels out. The Hammersley-Clifford theorem gives the relation between MRF and Gibbs random fields and states that \mathbf{D} is a Markov random field with respect to \mathcal{N} iif $p(\mathbf{d})$ is a Gibbs distribution with respect to \mathcal{N} [13, 14]. Thus the task is to specify potentials that induce the Gibbs measure in order encompass MRF properties of \mathbf{D} on the graph.

So far the description only encompasses a one-dimensional finite state space. However, it generalizes to multivariate distributions since any high dimensional process may be recast into a single state space with $\prod_i L_i$ states, where L_i is the cardinality of the i th variable. Furthermore, the description generalizes to the case of continuous distributions in which case $\exp(-U(\mathbf{d})/T)$ must be integrable. Since we wish to model correspondence between \mathbf{S}_s and \mathbf{S}_t the displacements are bound to the surfaces, in effect only posing constraints on the length of the three dimensional displacements at the individual sites. In practice the constraint may be enforced by projection of the displacements onto the closest point of the target surface in every site update of the MRF relaxation.

2.3 Prior Distributions

Similar to pixel priors [15] we construct energy functions based on differences between neighboring sites. Extending to the multivariate case we get the general expression of the energy governing the site-priors

$$U_{site}(\mathbf{d}) = \sum_{i \sim j} \|\mathbf{d}_i - \mathbf{d}_j\|_p^p \quad (2)$$

where $\|\cdot\|_p$ is the p -norm, $1 \leq p \leq 2$, and \mathbf{d}_i represents the multivariate displacement in the i th site.

With $p = 2$ the energy function induces a Gaussian prior on the deformation field. Neglecting regions with strong surface dynamics the local optimization becomes convex and the maximum likelihood (ML) estimate of the displacement at the i th site is taken as the mean of the neighboring displacements. By applying a weighted average

$$\hat{\mathbf{d}}_i = \sum_{j \in N_i} w_i \mathbf{d}_j / \sum_{j \in N_i} w_j \quad (3)$$

and using Gaussian weights, derived from a fixed kernel size, the maximum a-posteriori (MAP) state-estimate of the MRF is similar to the steady state of the algorithm for geometry constrained diffusion (GCD). GCD of $\mathbf{D} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ mapping the surface \mathbf{S}_s onto the surface \mathbf{S}_t is given in [6] by

$$\partial_t \mathbf{D} = \begin{cases} \Delta \mathbf{D} - \mathbf{n}_{S_t} \frac{\mathbf{n}_{S_t}^T \Delta \mathbf{D}}{\|\mathbf{n}_{S_t}\|^2} & \text{if } \mathbf{x} \in \mathbf{S}_s \\ \Delta \mathbf{D} & \text{if } \mathbf{x} \notin \mathbf{S}_s \end{cases} \quad (4)$$

where \mathbf{n}_{S_t} is the unit surface normal of $\mathbf{S}_t(\mathbf{D}(\mathbf{x}) + \mathbf{x})$. Thus, GCD is numerical scheme for solving a space and time discretized version of the heat equation on the deformation field with certain boundary conditions. Notice that in the MRF formulation we explicitly constrain the correspondence problem on the source and target surfaces, whereas the GCD implementation works on volume-voxel diffusion.

Abandoning homogeneity and isotropy of the MRF non-global kernels may be introduced. Thus, adaptive Gaussian smoothing may be applied, e.g. by setting the standard deviation of the kernel to the square-root of the edge length of the closest neighbor of site i on the graph. Moreover, using the $p = 1$ norm induces a median prior, with the ML estimate being the median of the displacements at the weighted neighboring sites. This property makes the MRF attractive for correspondence fields with discontinuities, thus avoiding the smearing of edges attained by the Gaussian prior.

2.4 Observation Models

Given a realization of the prior distribution, the observation model $p(\mathbf{y}|\mathbf{d})$ describes the conditional distribution of the observed data \mathbf{Y} . By specifying an observation model we may favor a mapping that establish correspondences between regions of similar surface properties. The similarity measures may include derived features of the observed data such as curvature, orientation of the surface normals, or even texture.

The simple dot product between the normals may form the basis for specifying a governing energy function that favors correspondence between regions of similar orientation by

$$U_{norm}(\mathbf{y}|\mathbf{d}) = \sum_i \|\mathbf{n}_{S_s,i}^T \mathbf{n}_{S_t,i} - 1\|^q, \quad (5)$$

where $\mathbf{n}_{S_s,i}$ is the surface normal at location \mathbf{x}_i on the source \mathbf{S}_s , and $\mathbf{n}_{S_t,i}$ is the normal of the target surface \mathbf{S}_t at the coordinate $\mathbf{x}_i + \mathbf{d}_i$. The parameter $q > 0$ controls the sensitivity of the energy function.

2.5 Maximum a Posteriori Estimates

Normalization of the energy terms from the different prior and observation models is typically chosen such that they operate on the same domain. However, the data analyst may choose to favor some terms over others, e.g. by relaxing the smoothness conditions in favor of correspondences between regions of similar curvature orientation of the surface normals.

The posteriori conditional probability distribution is given by

$$p(\mathbf{d}|\mathbf{y}) \propto \exp(-U_{total}/T), \quad (6)$$

where we use $U_{total} = (1 - \alpha)U_{norm} + \alpha U_{site}$, in which $\alpha \in [0 : 1]$ weighs the influence of the model terms. In searching for the MAP estimate

$$\hat{\mathbf{d}} = \operatorname{argmax}_{\mathbf{d}} p(\mathbf{d}|\mathbf{y}) \quad (7)$$

The Iterative Conditional Modes (ICM) method is a typical choice of optimization if the objective functional is convex. However, this is often only the case for simple MRFs and ML estimates are not always available. More advanced optimization can be done e.g. by simulated annealing using Gibbs sampling or the Metropolis-Hastings (MH) algorithm, followed by averaging or application of ICM in search of the most optimal state of the random field.

When applying simulated annealing the a-posteriori probability distribution is linked to the prior and the observation model by

$$p(\mathbf{d}|\mathbf{y}) \propto (p(\mathbf{y}|\mathbf{d})p(\mathbf{d}))^{1/T}, \quad (8)$$

where T is the temperature governing the process. At high temperatures all states are equally likely, however, decreasing the temperature increases the influence of the model terms. If the temperature is decreased slowly enough the algorithm will converge to the MAP estimate [16]. See [17, 18] for decreasing temperature schemes.

3 Results

Markov random field restoration using the Gaussian site-prior is applied to the training data after the TPS deformation of the model mesh using the PSP for initialization. In Fig. 2c,d we show a correspondence field after the MRF relaxation and the resulting reconstruction of the target shape. The figure is to be compared to Fig. 2a,b using the point to surface projection.

Problems in the registration field using PSP are removed by applying the MRF restoration. This is the case in respect to both the regularity of the polygonization, and the reconstruction error in representing the target shape by the deformed model surface. To obtain a measure of the uniformity of the polygonization of the target shape we examine the regularity of its triangular structure. By measuring the coefficient of variance of the edge lengths we obtain a standardized measure of the width of the underlying distribution. Results are shown in Fig. 3 for all subjects. The left plot shows the coefficients before and after MRF restoration of the correspondence field, and the right figure shows a histogram of the reductions in the coefficients of variance. A rank test shows the significance of the MRF regularization since a reduction in the coefficient is obtained for all subjects. The improvement in shape reconstruction is shown in Table 1. Applying the observation model is performed with $\alpha = 0.5$. This parameter should be chosen using cross-validation in a more exhaustive search for an optimal deformation field. However, since the shapes are relatively smooth and regular the results shows no significant improvement in the reconstruction error by introducing the observational term. In Fig. 4 the reconstruction error of the target shape of subject 1 is shown using PSP and MRF restoration based on the Gaussian site-prior. Notice the improved reconstruction using MRF.

When the model mesh is warped to another shape, it occurs that some correspondences are placed outside the region of interest on the target shape. Therefore, the model mesh is pruned to contain only the points that are warped to




Fig. 3. Left: Comparison between the point to surface projection (upper curve) and the MRF regularization (lower curve) by evaluating the coefficient of variance of the edge lengths of the polygonization of the target surface. Right: A histogram of the reduction in coefficient of variance over the training data.

valid areas for all shapes in the training set. The model mesh contains approximately 3000 vertices after pruning. Having established a dense correspondence field it is now possible to dispose of the anatomical landmarks as well as the original meshes of the training set. The set of meshes with dense correspondence is applied in the following statistical shape analysis. The shapes are aligned by a generalized Procrustes analysis [19]. The pure shape model is built using a similarity transformation in the Procrustes alignment while a rigid-body transformation is used to build the size-and-shape model [20]. An Active Shape Model (ASM) [2] is constructed based on a Principal Component Analysis (PCA) of the Procrustes aligned shapes. Let each aligned shape be represented as a vec-




Fig. 4. The reconstruction error [mm] for subject one using the point to surface projection (left) and the MRF correspondence restoration (right).

Subject	Registration Method			
	PSP	MRF _{p=2}	MRF _{p=1}	MRF _{p=2,q=1}
1	0.048 ± 0.013	0.044 ± 0.013	0.049 ± 0.014	0.043 ± 0.013
2	0.046 ± 0.013	0.042 ± 0.013	0.043 ± 0.012	0.040 ± 0.012
3	0.048 ± 0.014	0.042 ± 0.013	0.043 ± 0.013	0.040 ± 0.012
4	0.044 ± 0.012	0.038 ± 0.011	0.040 ± 0.011	0.038 ± 0.012
5	0.045 ± 0.013	0.042 ± 0.012	0.043 ± 0.012	0.040 ± 0.012
6	0.045 ± 0.014	0.046 ± 0.015	0.045 ± 0.015	0.043 ± 0.014
7	0.047 ± 0.014	0.046 ± 0.014	0.046 ± 0.014	0.046 ± 0.015
8	0.040 ± 0.011	0.038 ± 0.011	0.039 ± 0.011	0.050 ± 0.013
9	0.041 ± 0.011	0.039 ± 0.011	0.039 ± 0.011	0.038 ± 0.011
10	0.049 ± 0.015	0.044 ± 0.013	0.045 ± 0.013	0.043 ± 0.013
11	0.046 ± 0.013	0.046 ± 0.014	0.045 ± 0.013	0.055 ± 0.014
12	0.050 ± 0.014	0.043 ± 0.013	0.044 ± 0.013	0.041 ± 0.012
13	0.042 ± 0.010	0.037 ± 0.009	0.039 ± 0.009	0.041 ± 0.009
14	0.048 ± 0.013	0.040 ± 0.011	0.042 ± 0.012	0.040 ± 0.011
15	0.043 ± 0.012	0.041 ± 0.012	0.040 ± 0.012	0.038 ± 0.011
16	0.049 ± 0.013	0.043 ± 0.012	0.044 ± 0.012	0.052 ± 0.013
17	0.064 ± 0.019	0.049 ± 0.014	0.059 ± 0.018	0.064 ± 0.016
18	0.051 ± 0.015	0.042 ± 0.012	0.048 ± 0.013	0.053 ± 0.013
19	0.064 ± 0.020	0.052 ± 0.015	0.058 ± 0.017	0.049 ± 0.015
20	0.053 ± 0.015	0.049 ± 0.015	0.050 ± 0.015	0.050 ± 0.013
21	0.049 ± 0.013	0.041 ± 0.011	0.045 ± 0.012	0.039 ± 0.010
22	0.048 ± 0.014	0.042 ± 0.012	0.044 ± 0.013	0.048 ± 0.014
23	0.040 ± 0.011	0.037 ± 0.011	0.038 ± 0.011	0.042 ± 0.011
24	0.043 ± 0.013	0.041 ± 0.013	0.042 ± 0.013	0.048 ± 0.014
25	0.044 ± 0.013	0.037 ± 0.011	0.039 ± 0.011	0.046 ± 0.012
26	0.056 ± 0.014	0.046 ± 0.011	0.052 ± 0.012	0.058 ± 0.013
27	0.042 ± 0.011	0.039 ± 0.011	0.040 ± 0.011	0.039 ± 0.012
28	0.049 ± 0.013	0.041 ± 0.011	0.045 ± 0.013	0.047 ± 0.013
29	0.048 ± 0.014	0.045 ± 0.014	0.045 ± 0.013	0.047 ± 0.013
Average	0.048 ± 0.013	0.042 ± 0.012	0.045 ± 0.013	0.045 ± 0.013

Table 1. Reconstruction errors [mm] using PSP and MRF regularization. The mean \pm one std. is shown for each method. The site-prior is governed by the p -norm and q controls the sensitivity of the observational energy term dependent on the surface normals.

tor of concatenated x , y and z coordinates $\mathbf{x}_i = [x_{i1}, y_{i1}, z_{i1}, \dots, x_{in}, y_{in}, z_{in}]^T$, $i = 1, \dots, s$, where n is the number of vertices and s is the number of shapes. The PCA is performed on the shape matrix $\mathbf{D} = [(\mathbf{x}_1 - \bar{\mathbf{x}}) | \dots | (\mathbf{x}_s - \bar{\mathbf{x}})]$, where $\bar{\mathbf{x}}$ is the average shape. A new shape exhibiting the variance observed in the training set is constructed by adding a linear combination of eigenvectors to the average shape $\mathbf{x}_{new} = \bar{\mathbf{x}} + \Phi \mathbf{b}$, where \mathbf{b} is a vector of weights controlling the modes of shape variation and $\Phi = [\phi_1 | \phi_2 | \dots | \phi_t]$ is the matrix of the first t eigenvectors of $\mathbf{D}\mathbf{D}^T$. The three first modes of variation of the size-and-shape model derived using Gaussian MRF regularization are shown in Fig. 5. All the generated shapes look like natural ear canals with no deformations or folds in the mesh. Mode 1 consists of a bending of the canal and a flattening of the concha part. Mode 2 explains some of the shape variation observed in the inner part of the




Fig. 5. Size-and-shape shape model. The first three modes of variation shown at $+3$ (top) and -3 (bottom) standard deviations from the mean shape.

ear canal. Mode 3 is a combination of a flattening and twisting of the inner part of the ear canal and a general shape change of the concha. The distribution of the modes against each other is examined using pairwise plots and no obvious abnormalities were found (results not shown). In comparing the effect of the MRF regularization over the PSP method in the shape tangent space we find a reduction of more than 4% of the total variance of the resulting point distribution model. In Fig. 6 the variance contained in each principal component is shown together with the pct. reduction of the variance in each subspace. The average reduction of variance in each subspace is approximately 8% and the pct. reduction generally increases for higher dimensions.

4 Summary and Conclusions

A method is proposed for building statistical shape models based on a training set with an initial sparse annotation of corresponding landmarks of varying confidence. A model mesh is aligned to all shapes in the training data using the Thin Plate Spline transformation based on the anatomical landmarks. From the deformed model mesh and a target shape we derive a dense registration field of point correspondences. Applying the Markov Random Field restoration we obtain a dense, continuous, invertible registration field (i.e. a homeomorphism). The stochastic restoration acts as a relaxation on the TPS constrained model mesh with respect to the biological landmarks. The landmarks are identified with varying confidence and the MRF relaxation allows for a data driven




Fig. 6. Left: the variance contained in each principal component, the dotted line using point to surface projection and the solid line applying the MRF regularization step. Right: the reduction in the variance as a function of dimensionality of the model. The average reduction in each subspace is approximately 7% and the reduction of the total variance in the shape tangent space more than 4%.

enhancement of the object correspondences. Using the site-prior, the algorithm converges to the most simple deformation field which creates a tendency to match points of similar geometry since the field otherwise must be more complex. Moreover, inclusion of observational models could compensate further where the prior fails in more complex regions. In the present case study of smooth and regular shapes no significant benefit of applying more complex MRF were obtained. In comparison to applying point to surface projection the MRF regularization provides i) improved homogeneity of the target shape polygonization free of surface folds, ii) better reconstruction capabilities, and iii) a more compact Active Shape Model description of all the training data. The point to surface projection performs reasonably well in representing the target shape over most regions of the ear canals. However, it fails in regions with strong surface dynamics and when the source and target surfaces are too far apart. The fact that the MRF regularization produces a reduction of more than 4% of the total variance contained in shape tangent space is noteworthy. The reduction is explained by increased collinearity between semi-landmarks distributed over the entire shape. It indicates an improvement in the shape representation in terms of homologous point correlation and thus constitutes a better basis for generative modeling.

Acknowledgments

The authors would like to thank Dr. Rasmus Larsen, IMM, DTU, for valuable discussions on MRFs, and Audiology Technician Claus Nielsen, Oticon Research Centre Eriksholm, for annotating the ear canals.

References

1. Cootes, T.F., Taylor, G.J., Cooper, D.H., Graham, J.: Training models of shape from sets of examples. In: British Machine Vision Conference: Selected Papers 1992, Berlin, Springer-Verlag (1992)
2. Cootes, T., Cooper, D., Taylor, C., Graham, J.: Active shape models - their training and application. *Comp. Vision and Image Understanding* **61** (1995) 38–59
3. Gower, J.: Generalized Procrustes analysis. *Psychometrika* **40** (1975) 33–51
4. Larsen, R., Hilger, K.B.: Statistical 2D and 3D shape analysis using non-euclidean metrics (to appear). *Medical Image Analysis* (2003)
5. Paulsen, R.R., Larsen, R., Laugesen, S., Nielsen, C., Ersbøll, B.K.: Building and testing a statistical shape model of the human ear canal. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI*, Springer (2002) 373–380
6. Andresen, P.R., Nielsen, M.: Non-rigid registration by geometry-constrained diffusion. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI*, Springer (1999) 533–543
7. Andresen, P.R., Bookstein, F.L., Conradsen, K., Ersbøll, B.K., Marsh, J., Kreiborg, S.: Surface-bounded growth modeling applied to human mandibles. *IEEE Transactions on Medical Imaging*, **19** (2000) 1053–1063
8. Nielsen, M., Johansen, P., Jackson, A., Lautrup, B.: Brownian warps: A least committed prior for non-rigid registration. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI*, Springer (2002) 557–564
9. Davies, R., Cootes, T., Twining, C., Taylor, C.: An information theoretic approach to statistical shape modelling. In: *British Machine Vision Conference*. (2001) 3–11
10. R.H.Davies, C.J.Twining, T.F.Cootes, J.C.Waterton, C.J.Taylor: 3D statistical shape models using direct optimisation of description length. In: *Proc. ECCV*. Volume 3. (2002) 3–20
11. Bookstein, F.L.: Shape and the information in medical images: A decade of the morphometric synthesis. *Comp. Vision and Image Understanding* **66** (1997) 97–118
12. Lorenz, C., Krahnstöver, N.: Generation of point-based 3D statistical shape models for anatomical objects. *Comp. Vision and Image Understanding* **77** (2000) 175–191
13. Besag, J.: Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society, Series B* **36** (1974) 192–236
14. Geman, D.: Random fields and inverse problems in imaging. In: *Saint-Flour lectures 1988. Lecture Notes in Mathematics*. Springer-Verlag (1990) 113–193
15. Besag, J.: Towards Bayesian image analysis. *Journal of Applied Statistics* **16** (1989) 395–407
16. Geman, S., Geman, D.: Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **6** (1984) 721–741
17. Vidal, R.V.V.: Applied simulated annealing. In: *Lect. Notes in Econom. and Math.Syst.* Volume 396. Springer Verlag, Berlin (1993)
18. Cohn, H., Fielding, M.: Simulated annealing: searching for an optimal temperature schedule. *SIAM Journal of Optimization* **9** (1999) 779–802
19. Hilger, K.B.: *Exploratory Analysis of Multivariate Data*. PhD thesis, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby (2001)
20. Dryden, I., Mardia, K.: *Statistical Shape Analysis*. Wiley, Chichester (1997)