



Probabilistic Surface Reconstruction with Unknown Correspondence

Dennis Madsen^(✉), Thomas Vetter, and Marcel Lüthi

Department of Mathematics and Computer Science, University of Basel,
Basel, Switzerland

{dennis.madsen,thomas.vetter,marcel.luethi}@unibas.ch

Abstract. We frequently encounter the need to reconstruct the full 3D surface from a given part of a bone in areas such as orthopaedics and surgical planning. Once we establish correspondence between the partial surface and a Statistical Shape Model (SSM), the problem has an appealing solution. The most likely reconstruction, as well as the full posterior distribution of all possible surface completions, can be obtained in closed form with an SSM. In this paper, we argue that assuming known correspondence is unjustified for long bones. We show that this can lead to reconstructions, which greatly underestimate the uncertainty. Even worse, the ground truth solution is often deemed very unlikely under the posterior. Our main contribution is a method which allows us to estimate the posterior distribution of surfaces given partial surface information without making any assumptions about the correspondence. To this end, we use the Metropolis-Hastings algorithm to sample reconstructions with unknown pose and correspondence from the posterior distribution. We introduce a projection-proposal to propose shape and pose updates to the Markov-Chain, which lets us explore the posterior distribution much more efficiently than a standard random-walk proposal. We use less than 1% of the samples needed by a random-walk to explore the posterior. We compare our method with the analytically computed posterior distribution, which assumes fixed correspondence. The comparison shows that our method leads to much more realistic posterior estimates when only small fragments of the bones are available.

Keywords: Statistical Shape Model · Posterior estimation · Surface prediction · Surface uncertainty · Metropolis-Hastings

1 Introduction

Surface reconstruction is encountered in many different areas. The reconstructed surface can be used to guide the design of patient-specific implants in the medical area, or estimate the sex and ethnicity of an individual in forensic investigations [10, 11]. When only incomplete data is available, SSMs can be used to determine the most likely complete surface [2, 12, 13]. Nevertheless, the reconstruction becomes wrong and overconfident if correspondence cannot be obtained.




Fig. 1. A femur SSM is first registered to the partial surface to establish point-to-point correspondence. If we assume fixed correspondence, the posterior is computed analytically. With our method, we assume unknown correspondence and compute the posterior as a Markov-Chain. Note the correspondence difference to the available data visualised with the yellow markers on the long and short bone reconstructions. The coloured bones show the uncertainty, computed as the sum of the variances at each landmark with the different methods. Our method clearly shows more variability (red) far away from the partial surface, but at the same time has low variability (blue) at the known area. (Color figure online)

In cases where the surfaces are free of distinctive features, such as e.g. the shaft of a femur bone, there might even be multiple equally likely reconstructions with different lengths. In medical applications, a certainty estimate for a reconstruction is often required. This estimate can be computed as the likelihood of the chosen reconstruction in the distribution of all other possible reconstructions. Such a full distribution of surfaces given partial knowledge about the solution is known in the Bayesian setting as the posterior distribution. Since an SSM is formulated as a distribution over shapes, it is possible to derive a posterior model if only part of a surface is given [1] or if knowledge such as weight, sex, or age of a patient is known [3]. Current methods compute the posterior distribution analytically by assuming both fixed pose and fixed point-to-point correspondence [1, 3]. Furthermore, the analytical-posterior requires an initial dataset alignment before it can be computed. In absence of exact point-to-point correspondence, those conditions are impossible to fulfil.

We present a method to estimate the posterior distribution from partial surface knowledge. A similar method has previously been used for fitting an active shape model to a target image [9] to compute the most likely solution. In contrast, we compute the full distribution of possible reconstructions.

In Fig. 1, we show how our method compares to the analytical method. We use the Metropolis-Hastings (MH) algorithm to compute the Markov-Chain posterior distribution. We will be referring to our method as the sampling-posterior.

The default random-walk in MH takes a long time to converge. As the SSM should stay fixed around the given part of the surface, we have to use very small shape and pose updates. Informed sampling approaches overcome this problem by including knowledge from the current state into its proposal [4, 7]. We introduce a new projection-proposal, which keeps the known part of the model fixed and only varies the pose and shape in the unknown part. In [8], an Iterative Closest Point (ICP)-like proposal is introduced for surface registration, whereas our projection-proposal explores the space of possible surface reconstructions of partial surfaces and includes the variability in pose difference.

We apply the sampling-posterior to estimate the posterior distribution of partial femurs. The femur bone is used as an example as the full shape of the femur (as well as other long bones) is inherently difficult to estimate. Thus, there is almost no correlation between the shape of the upper and lower part and the length. We show the limitations of the current method by comparing our sampling-posterior to the analytical-posterior distributions of different partial femur bones. This paper makes the following contributions:

- We show how to compute the estimated posterior distribution of a partial surface without assuming a fixed point-to-point correspondence in Sect. 3.1.
- We introduce a new MH proposal to create independent samples and thereby speed up the posterior estimation process in Sect. 3.2.
- We show the limitations of the current analytical-posterior [1] in Sect. 2.1 and based on experimental results in Sect. 4.

2 Statistical Shape Models

SSMs learn the shape variability from examples shapes. When working with a shape M_i , we usually work with the vector representation $\mathbf{s}_i = (p_{1_x}^i, p_{1_y}^i, p_{1_z}^i, \dots, p_{n_x}^i, p_{n_y}^i, p_{n_z}^i)$ where $p \in \mathbb{R}^3$ is a landmark and n is the number of landmarks in the shape. A compact representation can be found by performing a Principal Component Analysis (PCA). The covariance matrix can then be represented by using $N - 1$ basis functions. In matrix format, the shapes are represented as $\mathbf{s} = \boldsymbol{\mu} + U\mathbf{D}\boldsymbol{\alpha} = \boldsymbol{\mu} + Q\boldsymbol{\alpha}$, with $\boldsymbol{\mu}$ being the mean shape, U being the matrix containing all the eigenvectors and D containing the square-root of the eigenvalues of the covariance matrix Σ . Each shape M_i can then be determined by an $\boldsymbol{\alpha}$ vector. The pose of the model can also be changed with both a translation vector $\mathbf{t} = (t_x, t_y, t_z)^T \in \mathbb{R}^3$ and a rotation matrix parameterised by the Euler angles $R(\phi, \psi, \rho) \in SO(3)$. All parameters are concatenated into one vector $\boldsymbol{\theta} = (\alpha_0, \dots, \alpha_{N-1}, \phi, \psi, \rho, t_x, t_y, t_z)^T$ and we use the notation $M[\boldsymbol{\theta}]$ to refer to the triangulated surface M defined by the parameter vector $\boldsymbol{\theta}$. The scale is directly incorporated in our construction of the SSM, as it would otherwise be difficult to obtain a correct statistical size measures if the size of the SSM can be scaled arbitrarily.

2.1 Analytical Posterior Models

We compare the sampling-posterior method to the analytical-posterior described in [1]. The given part of a shape is described as $\mathbf{s}_g \in \mathbb{R}^{3q}$ with q being the number of landmarks. In our model, this becomes $\mathbf{s}_g = \boldsymbol{\mu}_g + Q_g \boldsymbol{\alpha} + I_{3q} \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma^2)$ being the noise term of each landmark observation. The difficulty with the analytical-posterior is that point-to-point correspondence needs to be obtained before the \mathbf{s}_g vector can be defined. Furthermore, the rigid alignment needs to be fixed, resulting in the posterior distribution only containing shape variance.

In [1], the authors mention that all training shapes have to be aligned with respect to the subset of points available in \mathbf{s}_g in order to have a meaningful result. In the following, we will refer to the analytical-posterior computed without aligning according to the \mathbf{s}_g dataset as the naive-posterior and the analytical-posterior with the dataset alignment as the aligned-posterior.

3 Method

Now we explain how to compute the posterior distribution without assuming a fixed correspondence between the given data and the SSM. We define a probabilistic model over possible surface reconstructions (determined by $\boldsymbol{\theta}$) given partial surface information \mathbf{s}_g ,

$$P(\boldsymbol{\theta}|\mathbf{s}_g) = \frac{P(\boldsymbol{\theta})P(\mathbf{s}_g|\boldsymbol{\theta})}{\int P(\boldsymbol{\theta})P(\mathbf{s}_g|\boldsymbol{\theta})d\boldsymbol{\theta}}. \quad (1)$$

The shape prior $P(\boldsymbol{\theta})$ penalises unlikely shapes. The combined likelihood over all the points in the given surface \mathbf{s}_g is

$$P(\mathbf{s}_g|\boldsymbol{\theta}) = \prod_{i=1}^q \mathcal{N}(d_i(\boldsymbol{\theta}, \mathbf{s}_g); 0, \sigma^2), \quad (2)$$

where d_i is the Euclidean distance between the point p_i in the partial surface \mathbf{s}_g to the closest point on the surface of $M[\boldsymbol{\theta}]$. A similar likelihood function was used in [9] to measure the distance between an SSM and expert annotation in images. We define $\sigma^2 = 1.0 \text{ mm}^2$ and the same for ϵ in the analytical-posterior in order for the posterior distributions to be comparable. Note that the distance likelihood assumes that no pathologies exist in the partial surface.

3.1 Approximating the Probabilistic Model

Unfortunately, the full distribution of surfaces given the partial surface, as in Eq. (2), cannot be obtained analytically. Instead, it is possible to compute the unnormalised density for any surface described with $\boldsymbol{\theta}$. This allows us to use the MH algorithm to estimate the full posterior distribution. We use a random-walk to explore the shape space and have independent proposals for the translation $Q(\mathbf{t}'|\mathbf{t})$, rotation $Q(\mathbf{R}'|\mathbf{R})$, and shape $Q(\boldsymbol{\alpha}'|\boldsymbol{\alpha})$ parameters. As scaling is directly incorporated in our SSM, a scaling proposal is not used.




Fig. 2. Visualisation of the projection-proposal steps as described in Sect. 3.2. (Color figure online)

3.2 Projection-Proposal

Ideally, we would like to keep the known part of the shape model fixed around \mathbf{s}_g , as we are interested in the posterior distribution given partial surface information. With the random-walk proposals, we have to use very small shape and pose update steps. As a consequence many samples need to be taken before independent samples are found. Therefore, we suggest a projection-proposal to keep the shape at the known part of the model fixed and only vary the unknown part. The projection-proposal makes use of the analytical-posterior as described in Sect. 2.1. Before computing the analytical-posterior, we make a random rotation or translation proposal and compute the posterior distribution based on the initial position of the model. When computing the analytical-posterior, an anisotropic noise term ϵ is used. To simulate correspondence shift along the surface, we model it as a multivariate normal distribution with a low variance along the normal and a higher variance along the surface. The variance at each point p_k in $M[\boldsymbol{\theta}]$ is computed by

$$\Sigma_{p_k} = [\mathbf{n}, \mathbf{v}_1, \mathbf{v}_2] \begin{bmatrix} \sigma_n^2 & 0 & 0 \\ 0 & \sigma_v^2 & 0 \\ 0 & 0 & \sigma_v^2 \end{bmatrix} [\mathbf{n}, \mathbf{v}_1, \mathbf{v}_2]^T \quad (3)$$

where \mathbf{n} is the normal vector at the vertex p_k in the surface and \mathbf{v}_1 and \mathbf{v}_2 are perpendicular vectors to the normal. The variance along each vectors is set as $\sigma_n^2 = 0.1 \text{ mm}^2$ and $\sigma_v^2 = 5.0 \text{ mm}^2$. The projection-proposal can be described in 5 steps with Fig. 2 as visualisation for each step:

1. Compute corresponding points by taking the closest points from the partial surface (red) \mathbf{s}_g to the current surface $M[\boldsymbol{\theta}]$ (black). We compute \mathbf{s}_{g^*} as the points in the SSM corresponding to the partial surface.
2. Propose a random pose update from $Q(\mathbf{t}'|\mathbf{t}) + Q(\mathbf{R}'|\mathbf{R})$, while keeping the current shape parameters $\boldsymbol{\alpha}$ fixed, such that a new $\boldsymbol{\theta}'$ is computed ($M[\boldsymbol{\theta}']$ shown in blue).
3. Compute the analytical-posterior $p(\boldsymbol{\alpha}|\boldsymbol{\theta}', \mathbf{s}_{g^*})$ based on Sect. 2.1.
4. Draw a sample from the distribution $p(\boldsymbol{\alpha}|\boldsymbol{\theta}', \mathbf{s}_{g^*})$ by randomly setting the $\boldsymbol{\alpha}$ vector in the SSM (green shows the posterior mean).

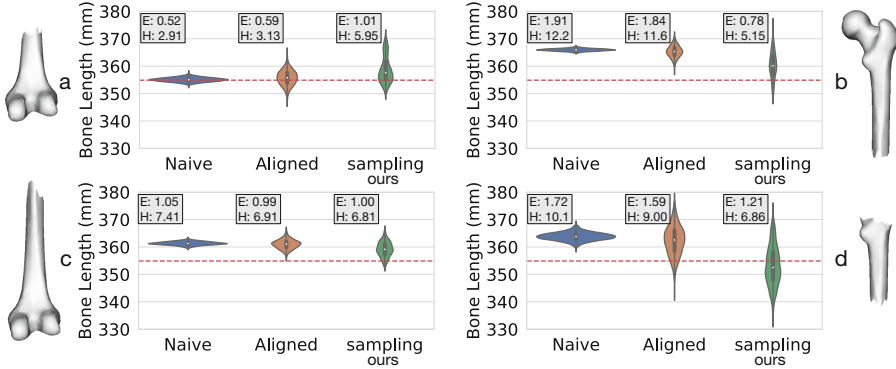


Fig. 3. Violin and box plots of bone length prediction in mm using the analytical-posterior and our sampling-posterior. All plots (a, b, c, d) concern the same ground-truth bone (length visualised with the red-dashed line), but differ in how much of the bone is given. The average Euclidean (E) and Hausdorff (H) distances from the ground-truth surface to the mean surface from the distributions are in mm. (Color figure online)

5. Compute the θ'' update based on the full SSM $p(\alpha)$ (the proposed sample $M[\theta'']$ is shown in green).

Unlike the random-walk proposal, this proposal is not symmetric. Therefore, to ensure convergence of the MH algorithm, we need to be able to compute the transition probabilities of going from θ to θ' as well as from θ' to θ [6]. The transition probability can be computed using the posterior distribution from step 3 as also shown in [8].

Projection-Proposal Importance. We need i.i.d. samples to compute the variance, which means that we need to find the number of samples to be discarded from the Markov-Chain before an independent sample is found. We compute the autocorrelation of the individual shape parameters and look for the number of samples needed to reach 0 correlation. We report 50 samples for the projection-proposal and 30×10^3 samples for the random-walk. While the random-walk requires 600 times more samples, the projection-proposal is only 6 times slower, making it overall 100 times faster.

We compute the bone length variance based on the distance between two landmarks. The length variation converges at 10^3 samples with the projection-proposal. For the random-walk, we need 500×10^3 samples to achieve the same length variance. With the projection-proposal, we can, therefore, explore similar variance numbers with less than 1% of the samples needed by the random-walk.

4 Evaluation

For the experiments, we use 61 femur meshes extracted from Computed Tomography (CT) images. We use 50 femurs for the femur SSM (femur lengths, mean:

372 mm, min: 322 mm, max: 437 mm) and 11 for the test-set (femur lengths, mean: 372 mm, min: 322 mm, max: 441 mm). The SSM contains a total of 1622 landmarks. Each test femur is divided into several partial meshes from where the posteriors are estimated. In Fig. 3, a subset of the cuts are shown¹.

Experimental Setup. We compare the sampling-posterior with the naive-posterior and the aligned-posterior. For the aligned-posterior, we need to estimate the observed points in the SSM. This is the same procedure that was done in step 1 of the projection-proposal. We perform a registration with the SSM and take the closest points to it from s_g . For the registration, we use the method from [9]. Alternatively, the non-rigid ICP algorithm can be used [5].

In the overview image of our method (Fig. 1), the posterior variability of the different methods is visualised with colours on the full femur bone. Very little variance is maintained in the naive-posterior, which highlights the importance of dataset alignment when computing the analytical-posterior. The sampling-posterior contains 2 to 3 times more variability than the aligned-posterior, suggesting that the full variability cannot be obtained using a fixed correspondence.

Length Estimation of Partial Bones. We compare the mean and the variance of bone lengths from the different posterior estimation methods. A landmark is placed at each end of the femur bones and the variability of the distance between the two landmarks is reported. For the analytical-posteriors, we randomly sample 10^3 shapes from the posterior models to be used for the estimate. For the projection-proposal we take 10^3 samples with 50 sample spacing in between. The bone length results for test femur 1 are shown in Fig. 3. Notice the difference between the results for partial bone *a* and *c*. More data is available in *c*, which results in a more narrow distribution, whereas the correspondence used in *c* is worse, making the ground-truth surface very unlikely under its distribution. The sampling-posterior results of the remaining test femurs are shown in Fig. 5.

We observe that both of the analytical-posterior methods sometimes fail to estimate the ground-truth length within their variability for most of the cuts. In contrast, the sampling-posterior can explain the shape length accurately.

Importance of Correct Correspondence. The quality of a surface reconstruction can be measured with the average Euclidean or Hausdorff distance to the ground-truth. These measures are, however, not a good indicator for the registration quality when large uncertainty exists in the correspondence. In Fig. 4 we show the same bone length experiment as in Fig. 3, but only for the aligned-posterior computed using different correspondences. The different correspondences have been found by initialising the SSM either as a very short, medium or long bone. We see that a close to perfect reconstruction can be found if the ground-truth correspondence is known, but at the same time can be extremely over

¹ All experiments are performed with the open source-library <https://scalismo.org>.




Fig. 4. Violin and box plots of bone length prediction in mm. All plots concern the same ground-truth bone (length visualised with the red-dashed line). The posteriors are computed with the aligned-analytical-posterior method and differ only in the correspondence which has been used. The average Euclidean (E) and Hausdorff (H) distances from the ground-truth surface to the mean surface from the distributions are in mm. (a) and (b) refers to the same partial shapes as in Fig. 3. (Color figure online)




Fig. 5. Violin and box plots of bone length prediction in mm. All plots of different ground-truth bones. The posteriors are computed with our sampling method.

or underestimate the bone length if wrong correspondences is used. The average Euclidean distance from the partial surface to the reconstruction is, in all the cases, less than 0.25 mm, which suggests that the model represents the surface well in the available part.

5 Conclusion

It is difficult to infer the full shape from a bone fragment. This is due to missing exact point-to-point correspondence. Previous methods deterministically find a set of correspondences to estimate the posterior. This can result in overconfident posterior estimates if incorrect correspondences are used. We have shown how previous methods even fail to explain the ground truth solution in an experimental setup with synthetic data. Our main contribution is a sampling approach that estimates the posterior distribution without relying on a fixed set of correspondences. We use the MH algorithm to obtain the variability in shape and pose reconstruction of partial surfaces. We have shown that the sampling-posterior, in contrast to the analytical methods, robustly is able to explain ground-truth data under its posterior. We also presented a technical contribution to the sampling-posterior in the form of a projection-proposal. This proposal is able to explore the posterior distribution more efficiently. With our sampling-posterior approach, both correspondence and reconstruction estimates are more accurate than with the traditional analytical approach. We are also able to more reliably estimate the uncertainty of the reconstruction results.

Acknowledgements. This research is sponsored by the Gebert Rűf Foundation under the project GRS-029/17.

References

1. Albrecht, T., Lüthi, M., Gerig, T., Vetter, T.: Posterior shape models. *Med. Image Anal.* **17**(8), 959–973 (2013)
2. Bernard, F., et al.: Shape-aware surface reconstruction from sparse 3D point-clouds. *Med. Image Anal.* **38**, 77–89 (2017)
3. Blanc, R., Reyes, M., Seiler, C., Székely, G.: Conditional variability of statistical shape models based on surrogate variables. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) *MICCAI 2009. LNCS*, vol. 5762, pp. 84–91. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04271-3_11
4. Cui, T., Law, K.J., Marzouk, Y.M.: Dimension-independent likelihood-informed mcmc. *J. Comput. Phys.* **304**, 109–137 (2016)
5. Feldmar, J., Ayache, N.: Rigid, affine and locally affine registration of free-form surfaces. *Int. J. Comput. Vision* **18**(2), 99–119 (1996)
6. Hastings, W.K.: Monte Carlo sampling methods using Markov chains and their applications (1970)
7. Kortylewski, A., et al.: Informed MCMC with Bayesian neural networks for facial image analysis. arXiv preprint [arXiv:1811.07969](https://arxiv.org/abs/1811.07969) (2018)
8. Madsen, D., Morel-Forster, A., Kahr, P., Rahbani, D., Vetter, T., Lüthi, M.: A closest point proposal for MCMC-based probabilistic surface registration. arXiv preprint [arXiv:1907.01414](https://arxiv.org/abs/1907.01414) (2019)
9. Morel-Forster, A., Gerig, T., Lüthi, M., Vetter, T.: Probabilistic fitting of active shape models. In: Reuter, M., Wachinger, C., Lombaert, H., Paniagua, B., Lüthi, M., Egger, B. (eds.) *ShapeMI 2018. LNCS*, vol. 11167, pp. 137–146. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-04747-4_13
10. Purkait, R.: Triangle identified at the proximal end of femur: a new sex determinant. *Forensic Sci. Int.* **147**(2–3), 135–139 (2005)
11. Trotter, M., Gleser, G.C.: Estimation of stature from long bones of american whites and negroes. *Am. J. Phys. Anthropol.* **10**(4), 463–514 (1952)
12. Zheng, G., et al.: Accurate and robust reconstruction of a surface model of the proximal femur from sparse-point data and a dense-point distribution model for surgical navigation. *IEEE Trans. Biomed. Eng.* **54**(12), 2109–2122 (2007)
13. Zhu, Z., Li, G.: Construction of 3D human distal femoral surface models using a 3D statistical deformable model. *J. Biomech.* **44**(13), 2362–2368 (2011)