A Unified Approach to Shape Model Fitting and Non-rigid Registration

Marcel Lüthi, Christoph Jud, and Thomas Vetter

Department of Mathematics and Computer Science, University of Basel

Abstract. Non-rigid registration and shape model fitting are the central problems in any shape modeling pipeline. Even though the goal is in both problems to establishing point-to-point correspondence between two objects, their algorithmic treatment is usually very different. In this paper we present an approach that allows us to treat both problems in a unified algorithmic framework. We use the well known formulation of non-rigid registration as the problem of fitting a Gaussian process model, whose covariance function favors smooth deformations. We compute a low rank approximation of the Gaussian process using the Nyström method, which allows us to formulate it as a parametric fitting problem of the same form as shape model fitting. Besides simplifying the modeling pipeline, our approach also lets us naturally combine shape model fitting and non-rigid registration, in order to reduce the bias in statistical model fitting, or to make registration more robust. As our experiments on 3D surfaces and 3D CT images show, the method leads to a registration accuracy that is comparable to standard registration methods.

1 Introduction

Statistical shape and deformation models are a well established part of many computer vision and medical image analysis pipelines. Both in building the statistical model, and in its application, the central problem is to find point-to-point correspondence between a reference (i.e. an image or surface) and a given image, such that the new image can be explained in terms of the reference. In the case where the reference is represented as an image, this is solved using image registration. The goal is to find a deformation field u^* from a space of deformations $\mathcal{F}$, which maps the corresponding points from the reference image I_R to a target images I_T. Formally, this is written as an optimization problem:

$$u^* := \arg\min_{u \in \mathcal{F}} \mathcal{D}[I_R, I_T, u] + \eta \mathcal{R}[u], \tag{1}$$

where $\mathcal{D}$ measures image similarity and the regularizer $\mathcal{R}$ how well the solution matches our prior assumptions. Given a statistical model, i.e. a generative model of the form $\mathcal{M}[\alpha](x) = \mu(x) + \sum_i \alpha_i \phi_i(x)$, which models the space of deformations in terms of a set of (learned) basis function ϕ_i, the optimization problem (1) becomes the parametric problem

$$\alpha^* := \arg\min_{\alpha \in \mathbb{R}^d} \mathcal{D}[I_R, I_T, \mathcal{M}[\alpha]] + \eta \mathcal{R}[\alpha]. \tag{2}$$

G. Wu et al. (Eds.): MLMI 2013, LNCS 8184, pp. 66–73, 2013.

It can be minimized using standard optimization techniques. The general non-rigid registration problem, is harder to solve, as there is no explicit model for the deformations available. In this case a variational approach is often employed, where the admissible deformations are specified by the regularization term $\mathcal{R}[u]$ which most often takes the form of a differential operator. A solution to the problem is obtained by solving a non-linear partial differential equation.

For building statistical models this mismatch of methodologies is unfortunate and adds considerable complexity to the modeling pipeline. In this paper we propose to unify both problems by constructing a parametric model for the general registration problem. The idea is to model the deformations as a Gaussian process $\mathcal{GP}(\mu, k)$, with mean function μ and covariance function (or kernel function) k. While this model is in general non-parametric, we can obtain a parametric approximation if we assume that the modeled deformation fields are sufficiently smooth. This is done by computing a low-rank approximation $\tilde{k}$ of the covariance function in terms of the first leading terms of its Mercer expansion $k(x, y) = \sum_i^n \lambda_i \phi_i(x) \phi_i(y)$ [11]. Under the new model $\mathcal{GP}(\mu, \tilde{k})$, each deformation can be written as

$$\mathcal{M}[\alpha](x) = \mu(x) + \sum_i \alpha_i \lambda_i \phi_i(x). \tag{3}$$

Thus, we can formulate non-rigid registration in the parametric form (2).

A main advantage of using Gaussian processes to model the deformations, is its flexibility. We can estimate its mean and covariance function from examples shapes to obtain a statistical model $\mathcal{GP}(\mu_{SM}, k_{SM})$ that incorporates shape constraints [1,12]. If we choose a zero mean and covariance function k_g that favors smooth functions, the resulting model $\mathcal{GP}(0, k_g)$ is generic and similar to models obtained by using differential operators as a regularizer [11]. The different mean and covariance functions can be combined, to construct a new Gaussian process $\mathcal{GP}(\mu_{SM}, k_D + k_{SM})$ that combines the characteristics of both models. Depending on the point of view, the resulting optimization problem can either be interpreted as a registration, which incorporates prior shape knowledge (see e.g. [17,18]) or as shape model fitting, which reduces the model bias [3,16]. In fact, this solution can be seen as an extension of the approach proposed by Wang et al. [16] for active shape models to statistical models with dense correspondence.

The use of Gaussian process models for non-rigid registration is not new. It has been extensively studied in the 90s by Grenander et al. (see the overview article [4] and references therein). Steinke et al. [13] later approached surface registration from a machine learning perspective, which led to a similar algorithm based on kernel methods. This approach was extended by Lüthi et al. [7] to a hybrid registration approach using Gaussian Process regression. The main novelty of our work is the use of the Nyström approximation to obtain a low-rank approximation of the Gaussian process. This allows us to derive an efficient numerical methods, for any covariance function that is sufficiently smooth, without requiring that the eigenfunctions are known analytically. In particular, this makes it possible to combine covariance functions for shape model fitting and non-rigid registration, which don't admit such an analytic form.

2 Background

2.1 Gaussian Processes

Stochastic processes allow us to define a probability distribution over a function space. Formally, a stochastic process is a collection of random variables $f(x), x \in \Omega$ where Ω is an index set. A Gaussian process is a stochastic process with the property that for any finite number of observations, $x_1, \ldots, x_n \in \Omega$ the values $f(x_1), \ldots, f(x_n)$ are jointly normally distributed [11]. A Gaussian process $\mathcal{GP}(\mu, k)$ is completely defined by its mean $\mu : \Omega \to \mathbb{R}$ and a covariance function $k : \Omega \times \Omega \to \mathbb{R}$. The covariance function $k(x, y)$ specifies for each pair of points x, y their covariance $E[f(x)f(y)]$. By specifying k, we define which functions are likely under the given process. Many known covariance functions imply a strong correlation between nearby values, which makes smooth functions more likely. Gaussian processes can also be used to model vector-valued functions. In this case, the covariance function becomes a matrix valued function $k(x, y) : \Omega \times \Omega \to \mathbb{R}^{d\times d}$, with $k(x, y) = E[f(x)f(y)^T]$. The most simple case of matrix-valued covariance function arise when we assume that the output dimensions are uncorrelated. In this case, we can construct a matrix-valued covariance function k from a scalar-valued covariance-function l by setting

$$\mathbf{k}(x, y) = \mathcal{I}_{d\times d}\, l(x, y),$$

where $\mathcal{I}_{d\times d}$ is the identity matrix. While vector-valued Gaussian processes seem like an extension of the theory, it can be shown that it can be reduced to the scalar case [5]. Thus all known results for real-valued Gaussian processes carry over to this more general setting.

2.2 Mercer's Expansion and Reproducing Kernel Hilbert Spaces

Closely related to a Gaussian process $\mathcal{GP}(\mu, k)$ is the reproducing kernel Hilbert space (RKHS) spanned by its covariance function k. An easy way to construct this space is to start from the eigenfunction expansion of k. According to Mercer's theorem (see e.g. [11]), a kernel k has an expansion in terms of a orthonormal set of basis functions:

$$k(x, y) = \sum_{i=1}^{\infty} \lambda_i \phi_i(x) \phi_i(y)^T, \tag{4}$$

where (λ_i, ϕ_i) are the eigenvalue/eigenfunctions pairs of the integral operator $\mathcal{T}_k f(\cdot) := \int_\Omega k(x, \cdot) f(x)\, dx$. We can define a Hilbert space by taking linear combinations of these eigenfunctions: $f(x) = \sum_{i=1}^{N} \phi_i(x)\alpha_i$ with $\sum_{i=1}^{n} \alpha_i^2/\lambda_i < \infty$. The inner product between two functions $f = \sum_{i=1} \alpha_i \phi_i$ and $g = \sum_{j=1} \beta_j \phi_j$ is defined by $\langle f, g\rangle_k = \sum_{i=1}^{\infty} \frac{\alpha_i \beta_i}{\lambda_i}$. Consequently, the norm becomes

$$\|f\|_k^2 = \langle f, f\rangle_k = \sum_{i=1}^{\infty} \frac{\alpha_i^2}{\lambda_i} \tag{5}$$

Note that the RKHS norm penalizes the eigenfunction components corresponding to small eigenvalues particularly strongly, a fact that we will use in Section 3.

2.3 The Nyström Approximation

To compute the eigenfunctions ϕ in the Mercer expansion (4), we use the Nyström approximation [11]. We randomly sample points $X = \{x_1, \ldots, x_N\}, x_l \in \Omega$ and perform a Monte Carlo integration of the eigenvalue equation:

$$\lambda_i \phi_i(x') = \int_\Omega k(x, x') \phi_i(x)\, dx \approx \frac{1}{N} \sum_{l=1}^{N} k(x_l, x') \phi_i(x_l),$$

which results in a matrix eigenvalue problem $K u_i = \lambda_i^{mat} u_i$. Here, $K_{il} = k(x_i, x_l)$ is the kernel matrix, u_i denotes the $i-$th eigenvector and λ_i^{mat} the corresponding eigenvalue. The eigenvalue λ_i^{mat} can be used as an approximation for λ_i. The eigenfunction ϕ_i in turn can be approximated using

$$\tilde{\phi}_i(x) = \frac{\sqrt{n}}{\lambda_i^{mat}} k_X(x) u_i \approx \phi_i(x)$$

with $k_X(x) = (k(x_1, x), \ldots, k(x_n, x))$.[1] In a practical implementation for image registration, it is computationally infeasible to explicitly compute $k_X(x')$ in every evaluation of ϕ_i. A suitable strategy, which we use in our method, is to pre-compute $\tilde{\phi}_i(x)$ for the points of a (possibly lower-resolution) image grid, and to use standard image interpolation to extend the values to the full image domain.

3 Registration Using a Low-Rank GP Model

The starting point for our method is the probabilistic formulation of the registration problem in [2]. The registration problem (1) is interpreted as the MAP estimation problem:

$$\arg\max_u p(u) p(I_T | I_R, u), \tag{6}$$

where $p(u) \propto \exp(-\mathcal{R}[u])$ is a Gaussian process prior over the admissible deformation fields and $p(I_T | I_R, u) \propto \exp(\eta^{-1} \mathcal{D}[I_R, I_T, u])$ is the likelihood. However, instead of specifying a regularization operator to define a Gaussian process, we model the mean μ and covariance function k directly. A MAP solution to (6) can be found by solving a minimization problem in the RKHS $\mathcal{F}_k$ defined by k (see e.g. [15] for details):

$$\arg\min_{u \in \mathcal{F}_k} \mathcal{D}[I_R, I_T, u] + \eta \|u\|_k^2, \tag{7}$$

where $\|\cdot\|_k$ denotes the RKHS norm. In the next step we construct a low-rank approximation defining an approximate kernel $\tilde{k}(x, x') = \sum_{i=1}^n \lambda_i \phi_i(x) \otimes \phi_i(x')$

[1] For the case of matrix-valued kernels, $k : \Omega \times \Omega \to \mathbb{R}^{d \times d}$, the matrices K and k_X become block matrices: $K \in \mathbb{R}^{nd \times nd}$ and $k_X \in \mathbb{R}^{nd \times d}$.

obtained as an eigenfunction expansion using the n largest eigenvalues. From (5) we see that the RKHS norm strongly penalizes components whose corresponding eigenvalue is small. Therefore, leaving out these components will have a negligible effect to the solution if the eigenvalues of the kernel k are quickly decreasing. Each deformation in the space modeled by the Gaussian process $\mathcal{GP}(\mu, \tilde{k})$ can now be written as the finite sum $u(x) = \mu(x) + \sum_{i=1}^{n} \alpha_i \phi_i(x)$. Thus, we can restate the problem in the parametric form

$$\arg\min_{\alpha_1,\ldots,\alpha_n} \mathcal{D}[I_R, I_T, \mu + \sum_{i=1}^{n} \alpha_i \phi_i] + \eta \sum_{i=1}^{n} \frac{\alpha_i^2}{\lambda_i}, \tag{8}$$

which can be minimized using any optimization algorithm.

3.1 Surface Registration

So far we have presented our method in the context of image registration. However, the approach is more general and the Gaussian process that defines the deformation model can be defined on arbitrary domains. Thus, we can define an algorithm for surface registration, by specifying the deformation model $\mathcal{GP}(\mu, k)$ on a reference surface $\Gamma_R \subset \mathbb{R}^d$. A simple formulation of the surface registration problem, which we use in this paper, is

$$\arg\min_{\alpha_1,\ldots,\alpha_n} \sum_{x_j \in \Gamma_R} D_T(x_j + \mu(x_j) + \sum_{i=0}^{n} \alpha_i \phi_i(x_j))^2 + \eta \sum_{i=1}^{n} \frac{\alpha_i^2}{\lambda_i}, \tag{9}$$

where D_T is a distance map defined for Γ_T.

4 Results

In this section we illustrate how our approach can be used to reduce the bias in statistical shape model fitting and show its feasibility for the use in 3D image registration. Our implementation uses the Statismo framework [8] for representing the Gaussian processes. Surface Registration is done using ITK[2] while for image registration we use Elastix [6]. Our implementation is freely available as part of Statismo[3]. For all the examples, we use a sum of squares distance metric and an LBFGS optimizer.

Shape Model Fitting. For our first experiment we used random face surfaces, which were generated using the Basel face model [10]. We sampled 60 faces as a training set to build a statistical shape model, and used 40 additional samples as a test set. We compared the fitting performance for four different models: 1) A shape model from all the samples, which we use as a ground truth, 2) A shape model obtained from the training data, which we denote by

[2] "The Insight Segmentation and Registration Toolkit" - `http://www.itk.org`
[3] "Statismo" - `http://www.statismo.org`

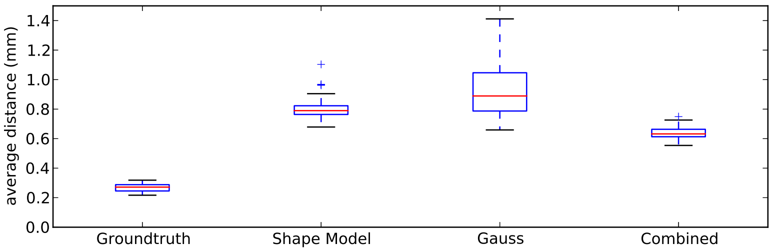

Fig. 1. Fitting results of different models to a set of 40 test surfaces. The shape model has a considerable bias. Combining it with a Gaussian with large bandwidth (which itself does not give good fitting results) helps to reduce the bias.

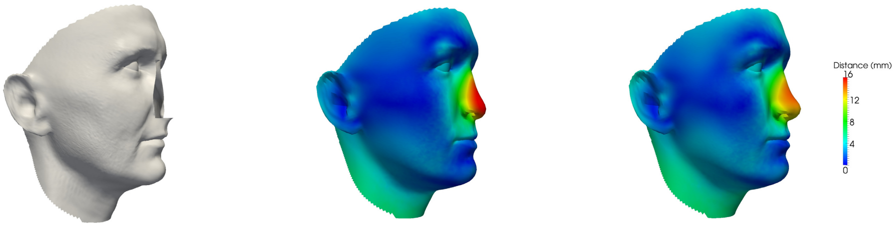

Fig. 2. A fitting result of a Gaussian model (middle) and a combined model (right) to a dataset with artifacts. Due to the shape constraint, the combined model preserves the nose shape better than the Gaussian model.

$\mathcal{GP}_{SM}(\mu_{SM}, k_{SM})$, 3) a model $\mathcal{GP}_G(0, k_G)$, with a Gaussian kernel $k_g(x, x') = \exp(\|x - x'/\sigma^2\|^2)$ ($\sigma = 100$), which is used to model the bias, 4) a combined model $\mathcal{GP}_C(\mu_{SM}, k_{SM} + k_G)$. Figure 1 shows the result of the four different fits obtained by minimizing (9). We observe that the shape model that was built from the training data only is biased and cannot accurately represent the faces. Due to the large bandwidth of the Gaussian kernel, also the Gaussian model cannot represent the faces accurately. However, combining them clearly reduces the bias of the shape model. In the second experiment we show how a registration method can be made more robust by including shape information. To simulate an artifact, we cut off the nose of one of the test faces (Figure 2a). We fit a Gaussian model (with $\sigma = 50$) and a combination of the shape model with this Gaussian model. Figure 2 shows that the combination of both models yields a smaller error around the nose, compared to using only the Gaussian model.

Registration of CT Data. In this experiment we use our method for the registration of CT images of dry femur bones, with a resolution of $176 \times 163 \times 622$. We select a reference image and perform a registration to 27 test images. As a deformation model, we used a zero mean Gaussian process with a Gaussian kernel ($\sigma^2 = 100$). We computed a low-rank approximation using the first 300 eigenfunctions. We compare the registration performance with the standard Demons

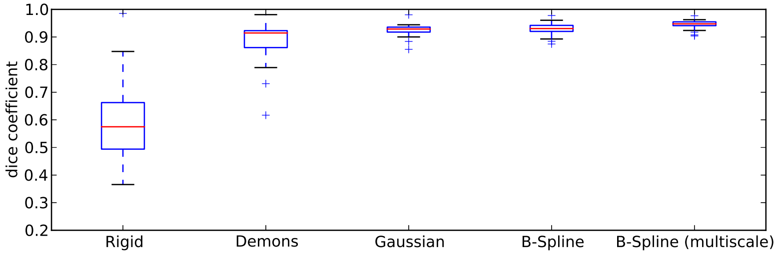

Fig. 3. Registration result achieved on a set of 27 femur images for different methods. The registration performance is measured by computing the dice coefficient of a ground-truth segmentation. As a baseline, we use a simple rigid registration (left). Our method (Gaussian) yields comparable performance to a standard B-spline and Demons registration.

algorithm [14] and a B-spline registration [12]. To have a fair comparison, we tested different parameters for each algorithm and used the best one in our comparison. As a performance measure, we use a dice coefficient, which is computed on manual segmentations of the images. Figure 3 shows the results for the different algorithms: We see that our approach outperforms the Demons algorithm and performs on par with the B-Spline registration. B-Spline registration perform, however, slightly better when a multi-scale strategy is used. This is expected, as the simple Gaussian kernel that we use only models deformations on a single scale. In order to obtain more flexible representation, more sophisticated kernels, such as e.g. the multi-scale kernel proposed in [9], can be used.

5 Conclusion

We have presented a unified approach to non-rigid registration and statistical model fitting. This is achieved by modeling the admissible deformations as a Gaussian process, which is fitted to the data. We compute a low-rank approximation using the Nyström method, and formulate registration as a parametric optimization problem. This makes our method computationally feasible even for large 3D images. We have shown that by combining kernels for non-rigid registration and shape modeling, we can reduce the bias of statistical shape models, or make non-rigid registration more robust. The main strength of our approach is that it makes it possible to use any sufficiently smooth kernel function to specify the admissible deformations. This gives us enormous flexibility to model our prior assumptions, while the algorithmic implementation remains the same. We believe that using more sophisticated prior models, such as a kernel for hybrid landmark and image registration [7], or the multi-scale kernel [9], we can obtain very powerful methods for non-rigid registration and model-fitting.

References

1. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3D faces. In: SIGGRAPH 1999: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, pp. 187–194. ACM Press (1999)
2. Christensen, G.E., Miller, M.I., Vannier, M.W., Grenander, U.: Individualizing neuro-anatomical atlases using a massively parallel computer. Computer 29(1), 32–38 (1996)
3. Cootes, T.F., Taylor, C.J.: Combining point distribution models with shape models based on finite element analysis. Image and Vision Computing 13(5) (1995)
4. Grenander, U., Miller, M.I.: Computational anatomy: An emerging discipline. Quarterly of Applied Mathematics 56(4), 617–694 (1998)
5. Hein, M., Bousquet, O.: Kernels, associated structures and generalizations. Max-Planck-Institut fuer biologische Kybernetik, Technical Report (2004)
6. Klein, S., Staring, M., Pluim, J.P.: Evaluation of optimization methods for non-rigid medical image registration using mutual information and b-splines. IEEE Transactions on Image Processing 16(12), 2879–2890 (2007)
7. Lüthi, M., Jud, C., Vetter, T.: Using landmarks as a deformation prior for hybrid image registration. Pattern Recognition, 196–205 (2011)
8. Lüthi, M., Blanc, R., Albrecht, T., Gass, T., Goksel, O., Büchler, P., Kistler, M., Bousleiman, H., Reyes, M., Cattin, P.C., et al.: Statismo-a framework for pca based statistical models (2012)
9. Opfer, R.: Multiscale kernels. Advances in Computational Mathematics 25(4), 357–380 (2006)
10. Paysan, P., Knothe, R., Amberg, B., Romdhani, S., Vetter, T.: A 3D face model for pose and illumination invariant face recognition. In: Advanced Video and Signal Based Surveillance 2009, pp. 296–301 (2009)
11. Rasmussen, C.E., Williams, C.K.: Gaussian processes for machine learning. Springer (2006)
12. Rueckert, D., Frangi, A.F., Schnabel, J.A.: Automatic construction of 3D statistical deformation models using non-rigid registration. In: Niessen, W.J., Viergever, M.A. (eds.) MICCAI 2001. LNCS, vol. 2208, pp. 77–84. Springer, Heidelberg (2001)
13. Schölkopf, B., Steinke, F., Blanz, V.: Object correspondence as a machine learning problem. In: ICML 2005: Proceedings of the 22nd International Conference on Machine Learning, pp. 776–783. ACM Press, New York (2005)
14. Thirion, J.P.: Image matching as a diffusion process: an analogy with Maxwell's demons. Medical Image Analysis 2(3), 243–260 (1998)
15. Wahba, G.: Spline models for observational data. Society for Industrial Mathematics (1990)
16. Wang, Y., Staib, L.H.: Boundary finding with prior shape and smoothness models. IEEE Transactions on Pattern Analysis and Machine Intelligence 22(7) (2000)
17. Wang, Y., Staib, L.H.: Physical model-based non-rigid registration incorporating statistical shape information. Medical Image Analysis 4(1), 7–20 (2000)
18. Xue, Z., Shen, D., Davatzikos, C.: Statistical representation of high-dimensional deformation fields with application to statistically constrained 3D warping. Medical Image Analysis 10(5), 740–751 (2006)