Spatially Varying Registration using Gaussian Processes.

Thomas Gerig[1], Kamal Shahim[2], Mauricio Reyes[2], Thomas Vetter[1], and Marcel Lüthi[1]

[1] University of Basel, Basel, Switzerland
{thomas.gerig,thomas.vetter,marcel.luethi}@unibas.ch
[2] University of Bern, Bern, Switzerland
{kamal.shahim,mauricio.reyes}@istb.unibe.ch

Abstract. In this paper we propose a new approach for spatially-varying registration using Gaussian process priors. The method is based on the idea of spectral tempering, i.e. the spectrum of the Gaussian process is modified depending on a user defined tempering function. The result is a non-stationary Gaussian process, which induces different amount of smoothness in different areas. In contrast to most other schemes for spatially-varying registration, our approach does not require any change in the registration algorithm itself, but only affects the prior model. Thus we can obtain spatially-varying versions of any registration method whose deformation prior can be formulated in terms of a Gaussian process. This includes for example most spline-based models, but also statistical shape or deformation models. We present results for the problem of atlas based skull-registration of cone beam CT images. These datasets are difficult to register as they contain a large amount of noise around the teeth. We show that with our method we can become robust against noise, but still obtain accurate correspondence where the data is clean.

1 Introduction

Methods for spatially-varying registration have recently gained a lot of attention. While traditional registration methods assume a deformation model that remains constant over the entire object domain, these methods allow for different regularization properties in different regions. Thus, they can be used to differentiate between tissue types or to regularize more strongly in areas where the data is noisy. In this work we propose a new scheme for spatially-varying registration based on non-stationary Gaussian process models. The basic method underlying our approach is a registration method proposed by Lüthi et al. [7]. In their work, all prior assumptions about the deformation fields u for the registration is specified in terms of a Gaussian process $u \sim \mathcal{GP}(\mu, k)$, with mean μ and covariance function (or kernel) k. This prior information is incorporated into the registration method by means of the reproducing kernel Hilbert space (RKHS) norm $\|\cdot\|_k$. The actual registration is performed by minimizing the functional:

$$\min_u \mathcal{D}[I_R, I_T, u] + \gamma \|u\|_k^2. \tag{1}$$

$\mathcal{D}$ is a distance measure between the reference and target image $I_R, I_T : \Omega \to \mathbb{R}$ with the dimension d, $\Omega \in \mathbb{R}^d$ and $\gamma \in \mathbb{R}^+$ is a regularization weight. The main contribution of our work is to obtain a flexible class of spatially-varying registration algorithms by using a spectral tempering scheme, proposed by Pintore and Holmes [8]. It is based on the simple idea, that for any given Gaussian process $\mathcal{GP}(\mu, k)$ we can obtain a new, non-stationary Gaussian process $\mathcal{GP}(\mu, k')$, by evolving the spectrum of k over space. This will effectively dampen or strengthen "high frequency" components in different areas, depending on a user defined tempering function η. Using the new kernel k' in (1) makes the registration spatially-varying, without affecting any other aspect of the registration method.

To show the versatility of our approach, we perform simple experiments on synthetic 2D images using different deformation models. We apply our method to the problem of establishing correspondence between a template skull and a set of target skulls which were extracted from CBCT data. These datasets exhibit a lot of noise and metal artifacts around the teeth and are difficult to register with standard methods. Our qualitative and quantitative evaluation shows that increasing the regularization strength in the noisy areas greatly improves the registration results.

Related work: Several different approaches to spatially-varying registration have been proposed in the literature [4, 1, 12, 3, 11]. The idea behind all of them is to make the regularization strength dependent on the location. This can be achieved by using a spatially-varying diffusion operator to smooth the deformation fields [4, 1] or by explicitly modeling spatially-varying deformations [12, 3, 11]. The difference to our method is that these all are incorporated directly into the registration algorithm, whereas our method is specified solely in terms of the prior model. Closest to our work is the work by [11], who also propose to achieve spatial variation by using a non-stationary kernel and to use a regularizer based on the RKHS norm. Their focus lies, however, on the algorithmic aspects of registration and they discuss only a very simple non-stationary kernel, which is based on partitioning the space into different regions, and assigning to each region a kernel with different smoothness properties.

2 Background

2.1 Gaussian Processes, Mercer's Expansion and Reproducing Kernel Hilbert Spaces

Stochastic processes allow us to define probability distributions over a function space. In the case of registration, we can use a stochastic process to model our prior assumptions about the deformations $u : \Omega \to \mathbb{R}^d$ in a registration task. A Gaussian process $GP(\mu, k)$ is a special stochastic process, which is completely defined by its mean function $\mu : \Omega \to \mathbb{R}^d$ and a covariance function (or kernel) $k : \Omega \times \Omega \to \mathbb{R}^{d \times d}$. Note that since we are modeling deformation fields, each value $k(x, x')$ is a $d \times d$ matrix, which specifies the correlation between all components of $u(x)$ and $u(x')$.

Closely related to a Gaussian process $\mathcal{GP}(\mu, k)$ is the reproducing kernel Hilbert space (RKHS) spanned by the kernel k. One way to construct this space is to start from the eigenfunction expansion of k. According to Mercer's theorem (see e.g. [9]), k has an expansion in terms of a orthonormal set of basis functions:

$$k(x, x') = \sum_{i=1}^{\infty} \lambda_i \phi_i(x) \phi_i(x')^T, \tag{2}$$

where (λ_i, ϕ_i) are the eigenvalue/eigenfunction pairs of the integral operator $\mathcal{T}_k f(\cdot) := \int_\Omega k(x, \cdot) f(x)\, dx$. The RKHS is spanned by linear combinations of these eigenfunctions: $f(x) = \sum_{i=1}^N \phi_i(x) \alpha_i$, with $\phi_i : \Omega \to \mathbb{R}^d$, $\alpha_i \in \mathbb{R}$ and $\sum_{i=1}^n \alpha_i^2 / \lambda_i < \infty$. The associated norm is defined by

$$\|f\|_k^2 = \langle f, f \rangle_k = \sum_{i=1}^{\infty} \frac{\alpha_i^2}{\lambda_i}. \tag{3}$$

Note that the components corresponding to small eigenvalues contribute particularly strongly to the RKHS norm. As these are usually associated to high frequency components, $\|f\|_k$ is a measure of smoothness for f. The exact notion of smoothness is defined by the kernel k.

2.2 Gaussian Process Registration using a Low-Rank GP Model

Gaussian process can be used to define a prior distribution $p(u) \sim GP(\mu, k)$ of possible deformation fields for a registration task. Defining a likelihood function $p(I_T|I_R, u) \propto \exp(-D[I_R, I_T, u])$, where $\mathcal{D}$ is a similarity measure between the images I_T, I_R, the registration problem can be cast as a MAP estimation problem [2]:

$$\arg\max_u p(u|I_R, I_T) = \arg\max_u p(u) p(I_T|I_R, u). \tag{4}$$

A MAP solution to (4) can be found by solving a minimization problem in the RKHS $\mathcal{F}_k$ defined by k (see e.g. [13] for details):

$$\arg\min_{u \in \mathcal{F}_k} \mathcal{D}[I_R, I_T, u] + \gamma \|u\|_k^2. \tag{5}$$

In order to minimize (5), Lüthi et al. [7] have proposed to perform a low rank approximation of k from the leading n eigenfunctions in (2). This allows them to obtain an approximate solution to (5) by solving the parametric problem

$$\arg\min_{\alpha_1, \ldots, \alpha_n} \mathcal{D}[I_R, I_T, \mu + \sum_{i=1}^n \alpha_i \phi_i] + \gamma \sum_{i=1}^n \frac{\alpha_i^2}{\lambda_i}. \tag{6}$$

by standard optimization procedures.

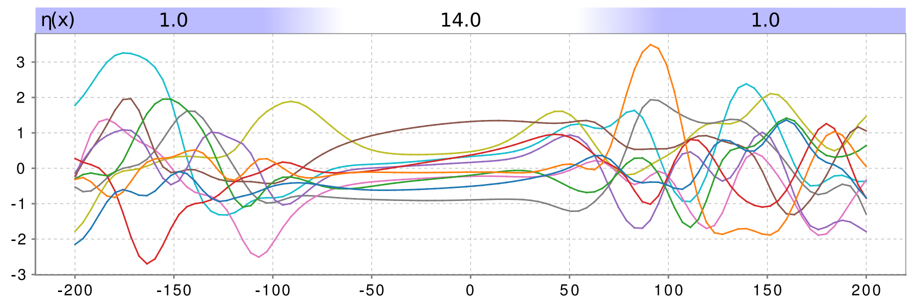

Fig. 1: Functions sampled from a Gaussian process ($GP(0, e^{-\|x-y\|^2/20})$) whose spectrum has been tempered using η. The functions are much smoother in areas where η is high.

3 Gaussian Processes for Spatially-Varying Registration

Using the registration approach discussed above, we can obtain a spatially-varying registration method, by making the deformation model a non-stationary Gaussian process. Many possibilities to define non-stationary kernels have been proposed in the statistics and machine learning literature. We use an approach proposed by Pintore and Holmes [8], which allows us to specify the degree of regularization using a function $\eta : \Omega \to \mathbb{R}$ defined over the image domain. The main idea behind this method is to compute a Mercer expansion of a kernel

$$k(x, x') = \sum_{i=1}^{n} \lambda_i \phi_i(x) \phi_i(x')^T$$

and to define a new kernel function

$$k'(x, x') = \sum_{i=1}^{n} \lambda_i^{\eta(x)/2} \lambda_i^{\eta(x')/2} \phi_i(x) \phi_i(x')^T \quad (7)$$

where the spectrum (i.e. the eigenvalues λ_i) is changed such that it varies over the space. By choosing $0 < \eta(x) < 1$, components corresponding to small eigenvalues will have more influence at x, while choosing $\eta(x) > 1$ will dampen their influence. Figure 1 illustrates this behavior for a simple 1D case. This tempering approach is quite natural for many registration models. If a translation-invariant kernel is used, such as the Gaussian kernel, the eigenfunctions are Fourier basis functions and hence the tempering approach will dampen high frequency components. When a kernel has been estimated from data using PCA, such as for statical shape or deformation models, this approach will dampen the components corresponding to small variations in the data. A problem arises, however, when this approach is used for registration: The tempering of the eigenvalues changes the scale of the deformations. We therefore normalize the eigenvalues in (7) by dividing by $\lambda_0^{\eta(x)/2} \lambda_0^{\eta(x')/2}$ with λ_0 as the largest eigenvalue.

The spatially-varying tempering function $\eta : \Omega \to \mathbb{R}$ was created by placing control points on the image domain Ω. In regions where we want to keep the global regularization term $\eta(x)$ has the value 1.0, which leaves the spectrum unchanged. We manually assigned higher values to regions where we want to damp the higher frequencies and therefore enforce low frequency deformations. Afterwards a smooth function was created by interpolating the points with Gaussian kernel regression and a B-spline kernel.

3.1 A Note on the Implementation

The implementation of this method is straight-forward. From a numerical point of view, the only challenge is to compute the eigenfunctions of the kernel. For this, we use the Nyström method (see e.g. [9], Chapter 4). As we are applying this method together with a registration approach that itself performs a low-rank approximation of the kernel, it is sufficient to compute the same number of eigenfunctions as are used for the registration. Usually, a few hundred eigenfunctions are sufficient to obtain a good approximation. The actual implementation is done using the statismo framework [6], which also contains an implementation of the Gaussian process registration method.

4 Results

We first illustrate our method using a simple toy example. The goal is to register the reference and target image shown in Figure 2. As a deformation model we use a zero-mean Gaussian process with a Gaussian Kernel. We see that a normal registration changes the shape of the inner rectangle. If we dampen the high frequencies in this inner rectangle, by choosing $\eta = 14$ the inner shape remains rectangular and the registration result is improved.

4.1 Model Based Segmentation of CBCT Skull Images

To evaluate the proposed approach on a medical application, four CBCT scans of patients undergoing cranio-maxillofacial (CMF) surgery were employed. Appro-

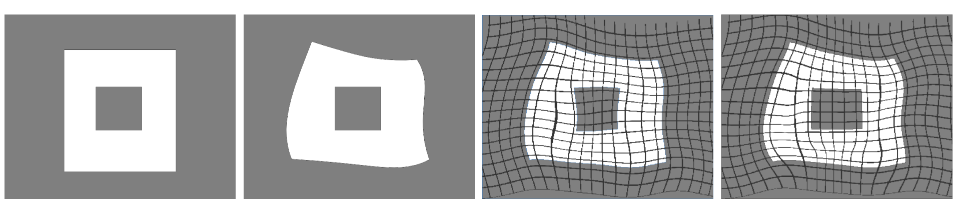

(a) Reference (b) Target (c) stationary (d) spatially-varying

Fig. 2: Toy example: Comparison of the results for a stationary and spatially-varying registration. The spatially-varying version is regularized more strongly inside the rectangle, and therefore preserves the shape much better.

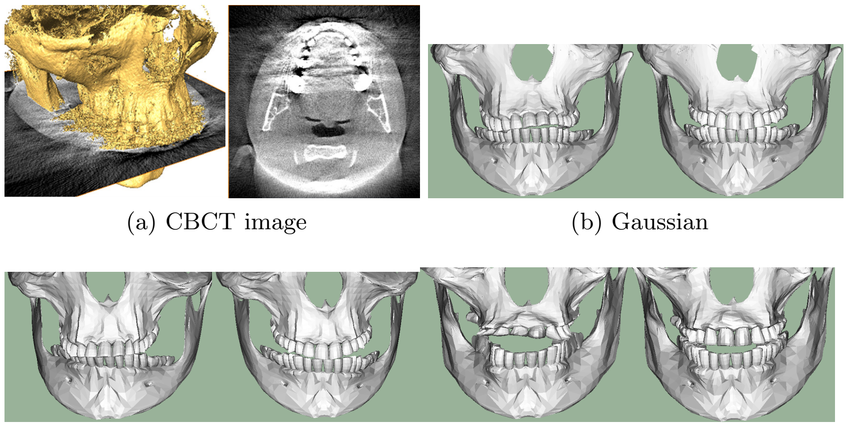

(a) CBCT image (b) Gaussian

(c) B-Spline (d) Statistical model

Fig. 3: Registration of skulls from CBCT data: (a) shows a slice through the image and a reconstruction of the surface that was obtained using threshold segmentation. (b) - (d) show registration results obtained using different deformation models. The left images show a normal registration, while in the right images a spatially-varying registration has been used.

priate segmentation of the maxilla, mandible and skull is a crucial step in CMF surgical planning and simulation. Typically for these patients metal artifacts surrounding the teeth area are a challenge, as a compromise between data quality fitting (risk of overfitting) and a-priori information (risk of over-regularization) needs to be made. In the following we show that the spatially-varying registration scheme is more robust towards outliers than the standard stationary registration approach. To perform model fitting on a new unseen image, we adopt the following approach: An initial estimate of the segmentation is generated via image thresholding. From this we compute a distance image to which we fit an atlas skull surface, using the Gaussian process registration approach [7]. We run the registration with three different deformation models. 1) A zero-mean Gaussian process with a Gaussian kernel, 2) a zero-mean Gaussian process with a B-spline kernel 3) A statistical shape model (PCA model) of the skull build from 48 example datasets. The left images in Figure 3 shows the result for these models. We see that for all three deformation models, the teeth are unnaturally distorted from the noise. The images on the right in Figure 3 show the same results with our spatially-varying registration (with $\eta = 1.5$ around the teeth and $\eta = 1$ otherwise). Clearly, the teeth look much more natural. In Figure 4a we see a numerical comparison of the mean error for all 4 cases. The error of the spatial adaptive method is lower for all the models. In Figure 4 the average error is depicted when the regularization weight γ in (1) is increased. As expected, the effect of the spatial adaption is only visible when the regularization weight is small. When the regularization weight is increased, both results will correspond more and more to the mean of the Gaussian process, and hence become very similar.

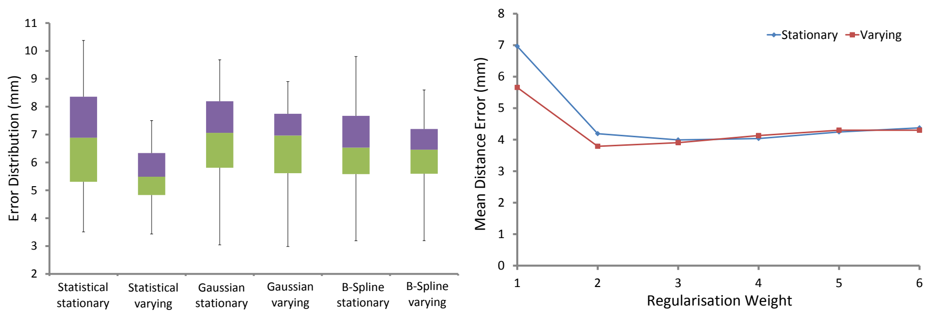

Fig. 4: A quantitative evaluation: In the left figure we see a numerical comparison of the mean error for all 3 cases. The mean error is an overall (i.e. entire skull surface) average distance error computed between the registered models and manually segmented results. The right plot shows the error for when the regularization weight is increased.

5 Discussion

In this paper we have shown how we can obtain a method for spatially-varying registration, by combining 1) A registration method that allows us to specify a deformation prior using a Gaussian process, 2) A method to make a given Gaussian process non-stationary by adapting its spectrum. The degree of smoothness that is induced by the Gaussian process can be specified using a tempering function. To define this tempering function we have proposed to specify the desired value at a few given points, and then to interpolate these points to define the function on the full registration domain. While this approach is easy, a more convenient approach could be implemented by using a learning scheme, similar to the ones proposed in [3, 12]. The tempering function either can be defined manually by an expert or also be inferred from data. In a first scenario the function is built by including a distance measure to the current target image and optimize control points of the function η with subject to local differences. [12] published a method to infer the regularization parameters directly from the data. A further possibility is the definition of the tempering function by taking systematic misregistrations into account. A spatially varying tempering function is a solution to compensate for the error [5]. In another possible scenario, tissue properties are learned from the data to include a variation of elasticity in the reference [10]. Our experiments have shown that our approach can significantly improve a registration, by enabling us to choose the smoothness based on the image characteristics. Indeed, we believe that the possibility to make a method spatially-varying is important for obtaining better correspondence in difficult registration tasks. As our method is able to work with many different deformation models, taking advantage of spatially-varying registration becomes feasible for many registration problems. To make its application easy, we have made the implementation available as part of the statismo framework [6].

References

1. Cahill, N.D., Noble, J.A., Hawkes, D.J.: A demons algorithm for image registration with locally adaptive regularization. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2009, Springer (2009)
2. Christensen, G.E., Miller, M.I., Vannier, M.W., Grenander, U.: Individualizing neuro-anatomical atlases using a massively parallel computer. Computer 29(1), (1996)
3. Davatzikos, C.: Spatial transformation and registration of brain images using elastically deformable models. Computer Vision and Image Understanding 66(2), (1997)
4. Freiman, M., Voss, S.D., Warfield, S.K.: Demons registration with local affine adaptive regularization: application to registration of abdominal structures. In: Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on. pp. 1219–1222. IEEE (2011)
5. Gee, A.H., Treece, G.M.: Systematic misregistration and the statistical analysis of surface data. Medical Image Analysis 18, 385–393 (2014)
6. 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. Insight Journal (2012)
7. Lüthi, M., Jud, C., Vetter, T.: A unified approach to shape model fitting and non-rigid registration. In: Machine Learning in Medical Imaging, pp. 66–73. Springer (2013)
8. Pintore, A., Holmes, C.: Spatially adaptive non-stationary covariance functions via spatially adaptive spectra. http:\\ www. stats. ox. ac. uk cholmes\ Reports\ spectral tempering. pdf (2004)
9. Rasmussen, C.E., Williams, C.K.: Gaussian processes for machine learning. Springer (2006)
10. Ruan, D., Fessler, J.a., Roberson, M., Balter, J., Kessler, M.: Nonrigid Registration Using Regularization that Accomodates Local Tissue Rigidity. In: Proc. Of SPIE. Medical Imaging: Image Processing. vol. 6144 (2006)
11. Schmah, T., Risser, L., Vialard, F.X.: Left-invariant metrics for diffeomorphic image registration with spatially-varying regularisation. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013, Springer (2013)
12. Simpson, I.J.A., Woolrich, M.W., Cardoso, M.J., Cash, D.M., Modat, M., Schnabel, J.A., Ourselin, S.: A Bayesian Approach for Spatially Adaptive Regularisation in Non-rigid Registration (2013)
13. Wahba, G.: Spline models for observational data. Society for Industrial Mathematics (1990)