Published in final edited form as:

Biomed Image Regist Proc. 2012 ; 7359: 120–130. doi:10.1007/978-3-642-31340-0_13.

Spatial Confidence Regions for Quantifying and Visualizing Registration Uncertainty

Takanori Watanabe and Clayton Scott

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA

Takanori Watanabe: takanori@umich.edu; Clayton Scott: clayscot@umich.edu

Abstract

For image registration to be applicable in a clinical setting, it is important to know the degree of uncertainty in the returned point-correspondences. In this paper, we propose a data-driven method that allows one to visualize and quantify the registration uncertainty through spatially adaptive confidence regions. The method applies to various parametric deformation models and to any choice of the similarity criterion. We adopt the B-spline model and the negative sum of squared differences for concreteness. At the heart of the proposed method is a novel shrinkage-based estimate of the distribution on deformation parameters. We present some empirical evaluations of the method in 2-D using images of the lung and liver, and the method generalizes to 3-D.

1 Introduction

Image registration is the process of finding the spatial transformation that best aligns the coordinates of an image pair. Its ability to combine physiological and anatomical information has led to its adoption in a variety of clinical settings. However, the registration process is complicated by several factors, such as the variation in the appearance of the anatomy, measurement noises, deformation model mismatch, local minima, etc. Thus, registration accuracy is limited in practice, and the degree of uncertainty varies at different image regions. For image registration to be used in clinical practice, it is important to understand its associated uncertainty.

Unfortunately, evaluating the accuracy of a registration result is non-trivial, mainly due to the scarcity of ground-truth data. For rigid-registration, there have been studies where physical landmarks are used to perform error analysis [3]. Statistical performance bounds for simple transformation models have been presented under a Gaussian noise condition [11,13]. However, it is generally difficult or impractical to extend these methods to nonrigid registration, which limits their applicability since many part of the human anatomy cannot be described by a rigid model.

While characterizing the accuracy of a nonrigid registration algorithm is even more challenging, there have been recent works addressing this issue. Christensen *et al.* initiated a project which aims to allow researchers to perform comparative evaluation of nonrigid registration algorithms on brain images [1]. Kybic used bootstrap resampling to perform multiple registrations on each bootstrap sample, and used the results to compute the statistics

of the deformation parameter [8]. In [6], Hub *et al.* proposed an algorithm and a heuristic measure of local uncertainty to evaluate the fidelity of the registration result. Risholm *et al.* adopted a Bayesian framework in [10], where they proposed a registration uncertainty map based on the inter-quartile range (IQR) of the posterior distribution of the deformation field. Simpson *et al.* also adopted the Bayesian paradigm in [12], where they introduced a probabilistic model that allows inference to take place on both the regularization level and the posterior of the deformation parameters. The mean-field variational Bayesian method was used to approximate the posterior of the deformation parameters, providing an efficient inference scheme.

We view the deformation as a random variable and propose a method that estimates the distribution of the deformation parameters given an image pair and registration algorithm. For illustration purpose, we use the cubic B-spline deformation model and the negative sum of squared differences as the similarity criterion, but the idea is applicable for other forms of parametric model (see [5] for other possible choices) and intensity-based registration algorithms. The estimated distribution will allow us to simulate realizations of registration errors, which can be used to learn spatial confidence regions. To the best of our knowledge, none of the existing methods view the registration uncertainty through spatial confidence regions represented in the pixel-domain. The confidence regions can be used to create an interactive visual interface, which can be used to assess the accuracy of the original registration result. A conceptual depiction of this visual interface is shown in Fig. 1. When a user, such as a radiologist, selects a pixel in the reference image, a confidence region appears around the estimated corresponding pixel in the homologous image. If the prespecified confidence level is, say $\gamma = 0.95$, then the actual corresponding point is located within the confidence region with at least 95% probability. The magnitude and the orientation of the confidence region offers an understanding of the geometrical fidelity of the registration result at different spatial locations.

2 Method

For clarity, the idea is presented in a 2-D setting, but the method generalizes directly to 3-D.

Nonrigid Registration and Deformation Model

When adopting a parametric deformation model, it is common to cast image registration as an optimization problem over a real valued function Ψ , a similarity measure quantifying the quality of the overall registration. Formally, this is written

$$\arg \max_{\theta} \Psi(f_r(\cdot), f_h \circ T(\cdot; \theta)), \quad (1)$$

where $f_r, f_h: \mathbb{R}^2 \rightarrow \mathbb{R}$ are the reference and the homologous images respectively, and $T(\cdot; \theta): \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a transformation parametrized by θ . Letting $\mathbf{x} = (x, y)$ denote a pixel location, a nonrigid transformation can be written $T(\mathbf{x}; \theta) = \mathbf{x} + \mathbf{d}(\mathbf{x}; \theta)$, where $\mathbf{d}(\cdot; \theta)$ is the deformation. To model the deformation, we adopt the commonly used tensor product of the cubic B-spline basis function β [7], where the deformation for each direction $q \in \{x, y\}$ is described independently by parameter coefficients $\{\theta_q\}$ as follows:

$$d_q(\mathbf{x}; \boldsymbol{\theta}_q) = \sum_{i,j} \theta_q^{(i,j)} \beta \left(\frac{x}{m_x} - i \right) \beta \left(\frac{y}{m_y} - j \right). \quad (2)$$

The scale of the deformation is controlled by m_q , which is the knot spacing in the q direction. If K knots are placed on the image, the total dimension of the parameter $\boldsymbol{\theta} = \{\boldsymbol{\theta}_x, \boldsymbol{\theta}_y\}$ is $2K$ since $\boldsymbol{\theta}_x, \boldsymbol{\theta}_y \in \mathbb{R}^K$.

Spatial Confidence Regions

Given the image pair f_r and f_h , let $\Omega_r \subset \mathbb{R}^2$ and $\Omega_h \subset \mathbb{R}^2$ denote the regions of interest in the reference and homologous image respectively. Also, let $\hat{\boldsymbol{\theta}}$ be the deformation coefficients estimated from registration (1). We will assume that the underlying ground-truth deformation belongs to the adopted deformation class, with deformation parameter $\boldsymbol{\theta}$. Then, the registration error $\mathbf{e}$ for pixel $\mathbf{x} \in \Omega_r$ is expressed as

$$\mathbf{e}(\mathbf{x}) = (e_x(\mathbf{x}), e_y(\mathbf{x})) = \mathbf{T}(\mathbf{x}; \hat{\boldsymbol{\theta}}) - \mathbf{T}(\mathbf{x}; \boldsymbol{\theta}). \quad (3)$$

We will view the true deformation $\boldsymbol{\theta}$ as a random variable, which introduces a distribution on $\mathbf{e}(\mathbf{x})$ for each $\mathbf{x}$. The confidence region $\Phi(\mathbf{x}) \subseteq \Omega_h$ is a set such that $\Pr(\mathbf{e}(\mathbf{x}) \in \Phi(\mathbf{x})) = \gamma$, where $\gamma \in [0, 1]$ is a prespecified confidence level. To estimate the spatial confidence regions, we adopt the following two-step process.

First, we estimate the distribution of $\boldsymbol{\theta}$. We assume $\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta)$, so the problem reduces to estimating $\boldsymbol{\mu}_\theta$ and $\boldsymbol{\Sigma}_\theta$. This is a challenging task because there is only a single realization of $\boldsymbol{\theta}$, corresponding to the given reference and homologous images, and this realization is not observed.

Second, given the estimates of $\boldsymbol{\mu}_\theta$ and $\boldsymbol{\Sigma}_\theta$ we can then simulate approximate realizations of $\boldsymbol{\theta}$ and thereby simulate spatial errors $\mathbf{e}(\mathbf{x})$. From this it is straight-forward to estimate $\Phi(\mathbf{x})$. However, sampling from $\mathcal{N}(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta)$ is potentially computationally intensive. The total dimension of $\boldsymbol{\theta}$ for the B-spline model is $2K$ in 2-D and $3K$ in 3-D. For a high resolution CT data-set of image size $512 \times 512 \times 480$ with voxel dimensions $1 \times 1 \times 1 \text{ mm}^3$, B-spline knots placed every 5 mm leads to a dimension on the order of millions. Sampling from a multivariate normal distribution requires a matrix square root of $\boldsymbol{\Sigma}_\theta$ but this is clearly prohibitive in both computational cost and memory storage. Therefore it is essential that the estimate $\hat{\boldsymbol{\Sigma}}_\theta$ have some structure that facilitates efficient sampling.

Estimation of Deformation Distribution

We use the registration result $\hat{\boldsymbol{\theta}}$ as the estimate for $\boldsymbol{\mu}_\theta$ and propose the following convex combination for $\boldsymbol{\Sigma}_\theta$

$$\hat{\boldsymbol{\Sigma}}_\theta = (1-\rho) \boldsymbol{\Sigma}_o + \rho \hat{\boldsymbol{\theta}} \hat{\boldsymbol{\theta}}^T. \quad (4)$$

The first term Σ_o is a positive-definite matrix which is an *a priori* baseline we impose on the covariance structure, and the second term is a rank-1 outer product that serves as the data-driven component. The weighting between the two terms is controlled by $\rho \in [0, 1)$. Note that (4) has a form of a shrinkage estimator reminiscent of the Ledoit-Wolf type covariance estimate [9], but only using the registration result $\hat{\theta}$.

For the baseline covariance Σ_o , we propose to use a covariance matrix which is motivated from the autoregressive model. Let $\Sigma_{AR} \in \mathbb{R}_{++}^{K \times K}$ denote the covariance of a first order 2-D autoregressive model, whose entries are given as

$$\Sigma_{AR}(i, j) = r_x^{|x(i) - x(j)|} r_y^{|y(i) - y(j)|}, \quad 1 \leq i, j \leq K. \quad (5)$$

Here, $|r_x| < 1$ and $|r_y| < 1$ are parameters that control the smoothness between neighboring knots, and $x(i) = \text{mod}(i - 1, n_x)$, $y(i) = \lfloor (i - 1)/n_x \rfloor$ are the mappings from the lexicographic index i to its corresponding (x, y) coordinate, assuming an $(n_x \times n_y)$ grid of knots. A key property of this dense matrix is that its inverse, or the precision matrix $\Theta_{AR} = \Sigma_{AR}^{-1}$, is block-tridiagonal with tridiagonal blocks. Specifically, Θ_{AR} has an n_y -by- n_y block matrix structure with each blocks of size $(n_x \times n_x)$, and only the main diagonal and the subdiagonal blocks are non-zero. Furthermore, these non-zero blocks are tridiagonal with the values of the non-zero entries known as a function of r_x and r_y .

Based on Σ_{AR} , we propose to use the following baseline covariance $\Sigma_o \in \mathbb{R}_{++}^{2K \times 2K}$ having a 2-by-2 block matrix structure expressed by the Kronecker product:

$$\Sigma_o = \begin{bmatrix} c_x \Sigma_{AR} & c_{xy} \Sigma_{AR} \\ c_{xy} \Sigma_{AR} & c_y \Sigma_{AR} \end{bmatrix} = \begin{bmatrix} c_x & c_{xy} \\ c_{xy} & c_y \end{bmatrix} \otimes \Sigma_{AR}. \quad (6)$$

The coefficients c_x and c_y assign the prior variance level on θ_x and θ_y , whereas c_{xy} assigns the prior cross-covariance level between θ_x and θ_y . The only restriction on these values is $(c_x c_y) > c_{xy}^2$, which ensures Σ_o is positive-definite. It is important to note that the precision matrix Θ_o of this baseline covariance is sparse, also having a 2-by-2 block matrix structure

$$\Theta_o = \Sigma_o^{-1} = \begin{bmatrix} c_x & c_{xy} \\ c_{xy} & c_y \end{bmatrix}^{-1} \otimes \Sigma_{AR}^{-1} = \begin{bmatrix} p_x & p_{xy} \\ p_{xy} & p_y \end{bmatrix} \otimes \Theta_{AR}, \quad (7)$$

where $\{p_x, p_y, p_{xy}\}$ are obtained by inverting the 2×2 coefficient matrix. The sparsity structure of Θ_o can be interpreted intuitively under a Gaussian graphical model framework. The conditional dependencies between knots are described by the non-zero entries in the matrix, which are represented as edges in an undirected graph. For our model, a knot $\theta_x(i, j)$ has 17 edges, 8 connected to its 8-nearest neighbors and the other 9 connected to the corresponding $\theta_y(i, j)$ knot and its 8-nearest neighbors. Fig. 2 provides an illustration of Σ_o

and the sparsity structure of its inverse Θ_o , along with an example realization of B-spline coefficients $\theta = (\theta_x, \theta_y)$.

Error Simulations and Spatial Confidence Regions

Since the estimate $\hat{\Sigma}_{\theta}(4)$ is a rank-1 updated form of the baseline Σ_o , we can exploit the sparsity structure of Θ_o to efficiently draw realizations from $\mathcal{N}(\hat{\theta}, \hat{\Sigma}_{\theta})$ without explicitly storing or computing a matrix square root for the dense matrix $\hat{\Sigma}_{\theta}$. We only need to store the sparse precision matrix Θ_o and compute its cholesky factor L_o , which can be done in $\mathcal{O}(K)$ operations [4]. This allows the sampling procedure to scale gracefully to 3-D.

Using such sampling procedure, we can now generate realizations of registration error $e(x)$ as follows:

1. Sample $\theta_i \sim \mathcal{N}(\hat{\mu}_{\theta}, \hat{\Sigma}_{\theta})$.
2. Synthesize reference image $f_r^{(i)}(x) \leftarrow f_h \circ T(x; \theta_i)$.
3. Register f_h on to $f_r^{(i)}$ to get estimate $\hat{\theta}_i$.
4. Compute error $e_i(x) = T(x; \hat{\theta}_i) - T(x; \theta_i)$.

We assume that $e(x) \sim \mathcal{N}(\mu_e(x), \Sigma_e(x))$ for all x . Then the spatial confidence region associated with pixel $x \in \Omega_r$ is defined by the ellipsoid

$$\Phi(x) = \{x' : (x' - \mu_e(x))^T \sum_e^{-1} (x' - \mu_e(x)) < \chi_2^2(1 - \gamma)\}, \quad (8)$$

which is the 100% level set of the bivariate normal distribution. Under this formulation, confidence region estimation becomes the problem of estimating $\{\mu_e(x), \Sigma_e(x)\}$, the mean and covariance of the registration error at pixel location x . We estimate these with the sample mean and covariance based on the simulated errors $\{e_i(x)\}$. Algorithm 1 outlines the overall spatial confidence region estimation process.

Note that since we are using $\hat{\theta}$ as the estimate for μ_{θ} it is important for the original registration to return a sensible result, as severe inaccuracy could negatively impact the quality of the spatial confidence regions.

3 Experiments

We demonstrate an application of the method, and also present preliminary experiments performed in 2-D. For illustration purpose, we used the negative sum of squared differences as the similarity criterion, but other metrics such as mutual information are also appropriate. To encourage the estimated deformation to be topology-preserving, we included the penalty term introduced by Chun *et al.* [2] into the cost function for all experiments.

Algorithm 1**Spatial Confidence Regions Generation****Input:** f_r, f_h **Output:** $\{\mu_e(x), \Sigma_e(x)\}$ for all $x \in \Omega_r$

$$\hat{\theta} \leftarrow \arg \max_{\theta'} \Psi(f_r(\cdot), f_h \circ T(\cdot; \theta'))$$

$$\mu_{\hat{\theta}} \leftarrow \hat{\theta}$$

$$\Sigma_{\hat{\theta}} \leftarrow (1 - \rho) \Sigma_o + \rho \hat{\theta} \hat{\theta}^T$$

for $i = 1, \dots, N$ sample $\theta_i \leftarrow \mathcal{N}(\mu_{\hat{\theta}}, \Sigma_{\hat{\theta}})$ generate $f_r^{(i)}(x) \leftarrow f_h \circ T(x; \theta_i)$

$$\hat{\theta}_i \leftarrow \arg \max_{\theta'} \Psi(f_r^{(i)}(\cdot), f_h \circ T(\cdot; \theta'))$$

register

compute $e_i(x) = T(x; \hat{\theta}_i) - T(x; \hat{\theta})$ **end**

$$\hat{\mu}_e(x) \leftarrow \frac{1}{N} \sum_{i=1}^N e_i(x)$$

$$\hat{\Sigma}_e(x) \leftarrow \frac{1}{N} \sum_{i=1}^N (e_i(x) - \hat{\mu}_e(x)) (e_i(x) - \hat{\mu}_e(x))^T$$

Application

We first applied the proposed method to two coronal CT slices in the lung region, shown in Fig. 3. Both images are size 256×360 , and the exhale-frame served as the homologous image while the inhale-frame served as reference. The notable motion in this data-set is the sliding of the diaphragm with respect to the chest wall. Due to the opposing motion fields at this interface, registration uncertainty is expected to be higher around this region. To model the deformation, we used a knot spacing of $(m_x, m_y) = (3, 8)$, resulting in a parameter dimension of $\theta \in \mathbb{R}^{7650}$. A tighter knot spacing was used for m_x since a finer scale of deformation was needed in the x -direction to model the sliding motion at the chest wall. Since the degree of this slide is relatively small for this data-set, the registration result shown in Fig. 3 looks reasonably accurate based on visual inspection.

Using $\hat{\theta}$ obtained from registering these images, we used the single-shot mean and covariance estimate and the efficient sampling scheme to obtain 100 new realizations of deformations. For the baseline covariance Σ_o , we used values of $(r_x, r_y) = (0.9, 0.9)$ and $\{c_x, c_y, c_{xy}\} = \{2, 4, 0.5\}$. A relatively high value for c_y was used since the magnitude of the overall deformation was higher in the y -direction. Finally, $\rho = 0.1$ was used, as it was found to produce sensible deformation samples. One of the synthesized reference images is shown in Fig. 3. Following Algorithm 1, we obtained a set of spatial confidence regions $\{\Phi(x)\}$ for all x in the region of anatomical interest, using a confidence level of $\gamma = 0.9$. A few of these

are displayed in Fig. 3(a)–(h), along with 100 simulated errors. It is important to note how the shapes of these confidence regions reflect the local image structure. The principal major axes of the ellipses are oriented along the edge, indicating higher uncertainty for those directions. The confidence regions for (c) and (g) take on isotropic shapes due to the absence of well-defined image structures. Finally, notice how the confidence region for (e) is quite large, illustrating how difficult it is to accurately register the sliding diaphragm at the chest wall.

Experimental Result

To quantitatively evaluate our method, we manually assigned μ_θ and Σ_θ for the cubic B-spline deformation-generating process. The mean deformation μ_θ was designed to model the exhale to inhale motion in the abdominal area around the liver region, simulated by a contracting motion field. Manually assigning a sensible ground-truth value for the covariance Σ_θ is extremely difficult due to its high dimension and positive-definite constraint. Therefore, we took the shrinkage-based covariance model (4) as the ground-truth, using values of $(r_x, r_y) = (0.95, 0.95)$, $\{c_x, c_y, c_{xy}\} = \{2, 3, 0.5\}$, and $\rho = 0.1$. These values imply that the covariance is smooth with moderate level of correlation in the x and y deformations. We sampled a single instance of deformation θ from this ground-truth distribution, and used it to deform a 2D axial CT slice in the liver region, having image size 512×420 . We labeled the original image as the homologous and the deformed image as the reference. This resulting image pair and their difference image are shown in Fig. 4. A knot spacing of $(m_x, m_y) = (8, 8)$ was used to define the scale of the ground-truth deformation, resulting in a parameter dimension of $\theta \in \mathbb{R}^{6656}$.

Next, we generated three classes of spatial confidence regions for this image pair, using confidence levels of $\gamma = 0.9$ and 0.95 . The first confidence region $\Phi_1(x)$ corresponds to the case where a correct deformation model is used for registration, and the parameter values for the shrinkage-based covariance estimate $\hat{\Sigma}_\theta$ matches that of the ground truth. The second confidence region $\Phi_2(x)$ corresponds to the case where there is a mismatch in the deformation model. Here, we used a fifth-order B-spline function during registration, with a knot spacing of $(m_x, m_y) = (6, 6)$. In addition, we introduced some discrepancies in the parameter values for $\hat{\Sigma}_\theta$. Finally, the third confidence region $\Phi_3(x)$ corresponds to the ideal case, and is constructed for the purpose of comparison. Here, a correct deformation model is used for registration, and the deformations used to train the spatial confidence regions were sampled from the ground-truth $\mathcal{N}(\mu_\theta, \Sigma_\theta)$ rather than the estimated distribution. The descriptions of these confidence regions are summarized in Table 1. All confidence regions were generated using $N = 200$ simulated errors.

To assess the quality of these spatial confidence regions, we evaluated their *coverage rates* by sampling $M = 500$ additional deformations from the ground-truth distribution $\mathcal{N}(\mu_\theta, \Sigma_\theta)$. Coverage rate for a given pixel x is defined as the percentage of registration errors that are confined within the confidence region $\Phi(x)$, and is written mathematically as

$$\frac{1}{M} \sum_{i=1}^M \mathbb{1} \{ \tilde{\mathbf{e}}_i(\mathbf{x}) \in \Phi(\mathbf{x}) \}, \quad (9)$$

where $\mathbb{1}\{\cdot\}$ is the indicator function, and $\tilde{\mathbf{e}}_i(\mathbf{x})$ are registration errors generated from deformations sampled from the ground-truth distribution. We computed the coverage rate for the pixels that are located within the region of anatomy. The resulting coverage rates are rendered as heatmaps and are displayed in Fig. 5, along with their corresponding histograms. It can be observed that the coverage rates for the first two confidence regions, $\Phi_1(\mathbf{x})$ and $\Phi_2(\mathbf{x})$, generally come close to the prespecified confidence level γ , although some degree of discrepancy can be observed at some image regions. The third confidence region $\Phi_3(\mathbf{x})$ gave the best result as expected; the coverage rate for all pixels comes very close to γ .

In summary, the performance of the spatial confidence regions $\Phi_1(\mathbf{x})$ and $\Phi_2(\mathbf{x})$ turned out to be reasonably close, having results comparable to the ideal case of $\Phi_3(\mathbf{x})$. Although further validation studies are required to obtain a more conclusive finding, this is an encouraging preliminary result.

4 Discussion and Conclusion

In this paper, we presented a new method to evaluate the accuracy of a registration algorithm using spatially adaptive confidence regions. Preliminary experimental test results in 2-D suggest the confidence regions are effective based on their coverage rates. However, it is important to note that the computational cost of the proposed method is N times the original registration algorithm, since we must register each of the sampled deformations. Depending on the user's choice, this N can be in the order of hundreds to even thousands, with higher values likely to return more reliable confidence regions. We note that the process is easily parallelizable. Furthermore, in application such as surgical planning and radiation therapy, it may not be necessary to have spatial confidence regions for every voxel in the image volume. Therefore, after completing the original full 3-D registration, we suggest to run the N registrations only within a subregion where the accuracy of the initial registration must be known. This allows one to obtain spatial confidence regions for these locations at a much more reasonable computational expense.

In the future, we will perform more extensive validation studies in 3-D using various similarity criteria and deformation models, and explore a way to quantify the robustness of the method. Furthermore, other choices of *a priori* baseline for the shrinkage-based covariance estimate will be investigated. Finally, we will seek a way to incorporate more data into our model to allow a more sophisticated parameter selection to take place.

Acknowledgments

This work was supported by NIH grant P01CA087634.

The authors would like to thank C. Meyer, A. Hero, and J. Fessler for the valuable discussions and their insightful feedback.

References

1. Christensen, GE.; Geng, X.; Kuhl, JG.; Bruss, J.; Grabowski, TJ.; Pirwani, IA.; Vannier, MW.; Allen, JS.; Damasio, H. Introduction to the non-rigid image registration evaluation project (NIREP). Third International Workshop on Biomedical Image Registration; LNCS; 2006. p. 128-135.
2. Chun SY, Fessler JA. A simple regularizer for B-spline nonrigid image registration that encourages local invertibility. *IEEE J Sel Top Sig Proc.* Feb; 2009 3(1):159–69. special Issue on Digital Image Processing Techniques for Oncology.
3. Fitzpatrick JM, West JB. The distribution of target registration error in rigid-body, point-based registration. *IEEE Trans Med Imaging.* 2001; 20(9):917–927. [PubMed: 11585208]
4. Golub, GH.; Van Loan, CF. Matrix computations. 3. Johns Hopkins University Press; 1996.
5. Holden M. A review of geometric transformations for nonrigid body registration. *IEEE Trans Med Imag.* Jan; 2008 27(1):111–28.
6. Hub M, Kessler ML, Karger CP. A stochastic approach to estimate the uncertainty involved in B-spline image registration. *IEEE Trans Med Imaging.* Nov; 2009 28(11):1708–1716. [PubMed: 19447703]
7. Kybic J, Unser M. Fast parametric elastic image registration. *IEEE Transactions on Image Processing.* Nov; 2003 12(11):1427–42. [PubMed: 18244700]
8. Kybic J. Bootstrap resampling for image registration uncertainty estimation without ground truth. *IEEE Transactions on Image Processing.* Jan; 2010 19(1):64–73. [PubMed: 19709978]
9. Ledoit O, Wolf M. Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection. *Journal of Empirical Finance.* 2003; 10:603–621.
10. Risholm, P.; Pieper, S.; Samset, E., III; WMW. Medical Image Computing and Computer-Assisted Intervention. Vol. 6362. LNCS; 2010. Summarizing and visualizing uncertainty in non-rigid registration; p. 554-561.
11. Robinson MD, Milanfar P. Fundamental performance limits in image registration. *IEEE Transactions on Image Processing.* Sep; 2004 13(9):1185–1199. [PubMed: 15449581]
12. Simpson IJ, Schnabel JA, Groves AR, Andersson JL, Woolrich MW. Probabilistic inference of regularisation in non-rigid registration. *NeuroImage.* 2012; 59(3):2438–2451. [PubMed: 21939772]
13. Yetik IS, Nehorai A. Performance bounds on image registration. *IEEE Transactions on Signal Processing.* May; 2006 54(5):1737–1749.

Fig. 1.

Conceptual illustration of the proposed method. The marks in (a)–(b) are a few point-correspondences estimated by registration. The confidence regions in (c) offer an understanding of the possible registration error for these pixels. We expect the shape of the confidence regions to reflect the local image structure, as demonstrated in (c).

Fig. 2.

Illustration of the properties of the baseline covariance Σ_o . The values used are $(n_x, n_y) = (50, 50)$, $(r_x, r_y) = (0.95, 0.8)$, and $\{c_x, c_y, c_{xy}\} = \{1, 2, 0.5\}$. (a) The baseline covariance Σ_o ,

(b) the sparsity structure of $\Theta_o = \Sigma_o^{-1}$ (c)–(d) B-spline coefficients θ_x and θ_y obtained from sample $\theta = (\theta_x, \theta_y) \sim \mathcal{N}(\mathbf{0}, \Sigma_o)$.

Fig. 3.

The top row shows the 2-D data-set used in the first experiment, along with the registration result and an image synthesized using one of the sampled deformations. A few of the confidence regions from $\mathbf{x} \in \Omega_r$ are shown in (a)–(h), with the red marks representing 100 realizations of registration error. Note how the confidence regions reflect the local image structure.

Fig. 4.

The data-set used for validation: (a) the homologous image $f_h(\mathbf{x})$, (b) the reference image $f_r(\mathbf{x}) = f_h(T(\mathbf{x}; \boldsymbol{\theta}))$ generated by a deformation coefficient sampled from the ground-truth distribution $\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta)$, (c) the absolute difference image.

Fig. 5.

The coverage rates evaluated for the three classes of spatial confidence regions presented in Table 1, displayed in the form of heatmap and histogram. Note that the performances of $\Phi_1(x)$ and $\Phi_2(x)$ are fairly comparable to the ideal confidence region $\Phi_3(x)$, as the coverage rates for many of the pixels come close to the prespecified confidence level γ .

Table 1

Spatial Confidence Regions Generated for Validation

	Def. Basis	Def. Scale	Parameter values used for $\hat{\Sigma}_{\theta}$
Conf. Reg. 1 $\Phi_1(x)$	Cubic	$m_x = 8$	$\rho = 0.1, (r_x, r_y) = (0.95, 0.95)$
	B-spline	$m_y = 8$	$\{c_x, c_y, c_{xy}\} = \{2, 3, 0.5\}$
Conf. Reg. 2 $\Phi_2(x)$	Fifth order	$m_x = 6$	$\rho = 0.15, (r_x, r_y) = (0.9, 0.9)$
	B-spline	$m_y = 6$	$\{c_x, c_y, c_{xy}\} = \{2, 2, 0\}$
Conf. Reg. 3 $\Phi_3(x)$	Cubic	$m_x = 8$	$\mu_{\hat{\theta}} = \mu_{\theta} \ \Sigma_{\hat{\theta}} = \Sigma_{\theta}$
	B-spline	$m_y = 8$	(Oracle)