Bayesian Characterization of Uncertainty in Multi-modal Image Registration

Firdaus Janoos, Petter Risholm, and William Wells (III)

Harvard Medical School, Boston, USA
{fjanoos,pettri,sw}@bwh.harvard.edu

Abstract. Understanding and quantifying the uncertainty involved when registering images is an important problem in medical imaging, where clinical decisions are made based on the registered solution. This is especially important in non-rigid registration where the higher degrees of freedom may provide unwarranted confidence in the results, through over-fitting. The Bayesian approach, which defines uncertainty as the posterior distribution on deformations, requires a generative model of the image formation process where the fixed image is modeled as a deformed version of the moving image plus a noise term. As per this model, the likelihood term is equivalent to the sum-of-squared differences image matching metric and is therefore valid only for same-mode image registration. In this paper, we propose a general formalism to quantify Bayesian uncertainty in the registration of multi-modal images through an extended probability model that introduces and then marginalizes out a stochastic transfer function between moving and fixed image intensities.

1 Introduction

1.1 Motivation

Registration is a fundamental tool for many bio-medical image analysis tasks such as longitudinal and population studies, and image guided surgery. However, assuming the physical validity of the deformation mechanism used in the registration procedure, imaging noise and artifacts, such as distortion or bias-field, along with the highly variable presentation of pathology affect the confidence in the optimal solution. This problem is compounded by the high degrees of freedom afforded by non-rigid registration models which introduces the possibility of over-fitting, for example, by through complex warps of regions which have insufficient contrast to guide the registration. This can happen primarily through the use of regularization, which is needed to condition the ill-posed model inversion, but which introduces long range dependencies in the solution. And finally, there is the uncertainty in the specification of model hyper-parameters, such as the mechanical properties of the underlying tissue or the statistics of imaging noise, all of which degrade the validity, sufficiency and accuracy of the deformation obtained through optimization. Therefore, quantifying and conveying the uncertainty in registration is extremely important, especially when clinical decisions are based on registration results.

In the Bayesian approach to non-rigid registration, the registration parameters are random variables and optimization may be used to obtain their *maximum a posteriori*

B.M. Dawant et al. (Eds.): WBIR 2012, LNCS 7359, pp. 50–59, 2012.

(MAP) estimates. More importantly, however, the Bayesian approach enables quantification of uncertainty as the posterior distribution over deformations, via measures such as variance, inter-quartile ranges, credibility intervals and entropy [2,9,5]. Here, the likelihood function, corresponding to the data fidelity term, measures the alignment of the two images under a transformation, while the prior corresponds to the regularization term penalizing implausible deformations. While different deformation and regularization combinations – such as freeform deformations with b-splines [9] or finite element (FE) meshes with elastic deformation penalty [2,5] – have been used, the likelihood terms have allowed for only same-modality image registration while assuming additive normal noise.

Specifically, denote the moving image as $\mathrm{M} : \Omega_{\mathrm{M}} \rightarrow \mathbb{I}_{\mathrm{M}}$ and $\mathrm{F} : \Omega_{\mathrm{F}} \rightarrow \mathbb{I}_{\mathrm{F}}$, where $\Omega_{\mathrm{M}} \subseteq \mathbb{R}^3$ and $\Omega_{\mathrm{F}} \subseteq \mathbb{R}^3$ are the spatial domain of the fixed and moving images respectively, and $\mathbb{I}_{\mathrm{M}} = \mathbb{I}_{\mathrm{F}} \subseteq \mathbb{R}$ are their equal intensity ranges. The images F and M are treated as random fields and the fixed image is assumed to be generated by applying a transformation $\mathbf{u}$ to the moving image domain as per: $\mathrm{F} = \mathrm{M} \circ \mathbf{u} + \epsilon$ which also introduces normal noise $\epsilon \sim \mathcal{N}(0, \tau_\epsilon)$. Here, the transform $\mathbf{u}$ too is treated as a random variable with a probability measure, given by the prior distribution. Additionally, the likelihood is assumed to spatially iid.

Therefore, the log-posterior distribution of the transform is:

$$\ln p\left(\mathbf{u} \mid \mathrm{M}, \mathrm{F}\right) = -\int_{\Omega_{\mathrm{F}}} \frac{\left|\mathrm{F}_{\mathbf{x}} - \mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right|^2}{2\tau_\epsilon} d\mathbf{x} - \frac{E_{\mathrm{reg}}(\mathbf{u})}{2\tau_{\mathrm{reg}}} + \text{const}, \tag{1}$$

where $E_{\mathrm{reg}}(\mathbf{u})$ is the regularization energy of the transformation. Here, the first term is proportional to the log-likelihood term $\ln p(\mathrm{F}, \mathrm{M}|\mathbf{u})$ while the second is proportional to the log-prior on the transformations $\ln p(\mathbf{u})$. The temperatures τ_ϵ and τ_{reg} are model hyper-parameters, where τ_ϵ is related to the variance of the image noise, while τ_{reg} controls the variance of the prior on transformations.

This model provides a principled basis for the interpretation of the posterior density as uncertainty in parameter estimates, for setting priors on the model hyper-parameters and as shown in [4], for eliminating the uncertainty due to HPs by marginalizing them out.

1.2 Contribution

One of the main drawbacks, however, of this framework is that it restricts the image similarity term to sum-of-squared differences (SSD) metric and is applicable to only same-mode images. Here, we present an extension for multi-modal image registration through a generalization of the SSD metric that accounts for arbitrary intensity transformations. Specifically, we introduce a latent random process $\eta_{\mathbf{u}[\mathbf{x}]}(m) \in \mathbb{I}_{\mathrm{F}}$ defined on moving image intensities $m \in \mathbb{I}_{\mathrm{M}}$ which serves as a *link function* between the moving image intensity range to that of the fixed image. The posterior density of the link process is directly estimated from the data, and is marginalized out by means of the free-energy equivalence. In this paper, non-parametric kernel density estimation is used, although this framework supports any alternative parametric or non-parametric density estimation method. The new registration model is evaluated on a synthetic T1-to-T2 MR image registration problem with ground-truth. A clinical application (cryoablation) to register pre-operative abdominal MR with an intra-procedural CT image is also demonstrated.

1.3 Related Work

In addition to fixed (known) transfer functions, parametric and non-parametric methods to estimate the intensity transformation from the data are commonly used in multi-modal registration [7]. For example, Guimond *et al.* [3] learn a polynomial mapping between the intensities. Roche *et al.* [6] use the conditional expectation $\mathbb{E}\{m \in \mathbb{I}_{\mathrm{M}} | f \in \mathbb{I}_{\mathrm{F}}\}$ as the intensity transformation as which leads to the correlation ratio as the image-matching metric. These metrics use point-estimates of the transfer function, in contrast to the full posterior as done here, thereby limiting their ability to deal with non-stationary and noisy intensity mappings.

While image matching based on the joint-histogram of intensities, such as mutual information, generalize to a very wide class of intensity transformations, they do not have an associated probabilistic model and therefore cannot be used to compute a posterior. Zöllei *et al.* [10] present an alternative probabilistic model for image registration using Dirichlet priors on the latent parameters of joint multinomial models on discrete intensities. Marginalization led to objective functions that approximate entropy or likelihood formulations and only MAP estimates were sought.

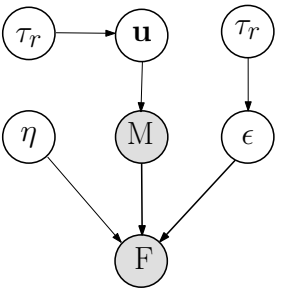

Fig. 1. Proposed probability model for multi-modal image registration. Here, F and M are the fixed and moving images respectively, ϵ is white normal noise with variance τ_ϵ, $\mathbf{u}$ are the transformation parameters with prior distribution variance controlled by τ_{reg}, and η is the link process that maps moving image intensities to fixed image intensities.

2 Method

2.1 Multi-modal Registration Model

The proposed Bayesian model for the multi-modal registration problem is as follows:

$$\mathrm{F}_{\mathbf{x}} = \eta_{\mathbf{u}[\mathbf{x}]}\left(\mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right) + \epsilon_{\mathbf{x}}. \tag{2}$$

In this model, the moving and fixed image M and F and additive normal noise ϵ are all spatially iid random processes, while the transformation parameters $\mathbf{u}$ have a prior distribution specified by the regularization term. The link function η is a stochastic process defined on $\Omega_{\mathrm{F}} \times \mathbb{I}_{\mathrm{M}}$, with $\eta_{\mathbf{u}[\mathbf{x}]}$ defined on the moving image intensity range $\mathbb{I}_{\mathrm{M}}$, such that $\eta_{\mathbf{u}[\mathbf{x}]}\left(\mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right) \in \mathbb{I}_{\mathrm{F}}$ maps moving image intensity $\mathrm{M}_{\mathbf{u}[\mathbf{x}]} \in \mathbb{I}_{\mathrm{M}}$ to a fixed image intensity. We assume η to be iid in the space dimension (Ω_{F}) and $\eta_{\mathbf{u}[\mathbf{x}]}$ to be independent in the moving intensity dimension ($\mathbb{I}_{\mathrm{M}}$). In the following discussion, define $f_{\mathbf{u}[\mathbf{x}]} \triangleq \eta_{\mathbf{u}[\mathbf{x}]}(\mathrm{M}_{\mathbf{u}[\mathbf{x}]})$. Also, we will drop explicit conditioning on $\mathbf{u}$, F and M except when there is ambiguity.

Under the spatial iid assumptions of F and M, and approximating $\epsilon_{\mathbf{x}}$ by its conditional expectation $\mathbb{E}\left\{\epsilon_{\mathbf{x}} \mid \mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right\} = 0$, the strong law of large numbers yields that the marginal posterior density of the link process

$$p\left(\eta_{\mathbf{u}[\mathbf{x}]}(\mathrm{M}_{\mathbf{u}[\mathbf{x}]}) \mid \mathbf{u}, \mathrm{F}, \mathrm{M}\right) = p\left(f_{\mathbf{u}[\mathbf{x}]} \mid \mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right) \approx \frac{\mathrm{vol}\left\{[\mathrm{F} = f_{\mathbf{u}[\mathbf{x}]}] \cap [\mathrm{M} \circ \mathbf{u} = \mathrm{M}_{\mathbf{u}[\mathbf{x}]}]\right\}}{\mathrm{vol}\left\{[\mathrm{M} \circ \mathbf{u} = \mathrm{M}_{\mathbf{u}[\mathbf{x}]}]\right\}} \tag{3}$$

We can marginalize out the latent process $\eta_{\mathbf{u}[\mathbf{x}]}$ under its posterior using the free-energy equivalence:

$$\begin{aligned}\ln p(\mathbf{u} \mid \mathrm{F}, \mathrm{M}) &= \int p(\eta|\mathbf{u}, \mathrm{F}, \mathrm{M}) \ln p(\mathbf{u}, \eta \mid \mathrm{F}, \mathrm{M}) d\eta - \int p(\eta|\mathbf{u}, \mathrm{F}, \mathrm{M}) \ln p(\eta|\mathbf{u}, \mathrm{F}, \mathrm{M}) d\eta \\ &= \int p(\eta|\mathbf{u}, \mathrm{F}, \mathrm{M}) \ln p(\mathbf{u}, \eta \mid \mathrm{F}, \mathrm{M}) d\eta + \mathbb{H}\left\{\eta|\mathbf{u}, \mathrm{F}, \mathrm{M}\right\}, \end{aligned} \tag{4}$$

where $\mathbb{H}\left\{\eta_{\mathbf{u}[\mathbf{x}]}|\mathbf{u}, \mathrm{F}, \mathrm{M}\right\}$ is the differential entropy of $p(\eta_{\mathbf{u}[\mathbf{x}]}|\mathbf{u}, \mathrm{F}, \mathrm{M})$. In the case of the model given in eqn. (2), the log-posterior of the deformations is:

$$\begin{aligned}\ln p\left(\mathbf{u} \mid \mathrm{F}, \mathrm{M}\right) = &- \int_{\Omega_\mathrm{F}} \int_{\mathbb{I}_\mathrm{F}} p\left(\eta_{\mathbf{u}[\mathbf{x}]}(\mathrm{M}_{\mathbf{u}[\mathbf{x}]})\right) \frac{\left|\mathrm{F}_\mathbf{x} - \eta_{\mathbf{u}[\mathbf{x}]}(\mathrm{M}_{\mathbf{u}[\mathbf{x}]})\right|^2}{2\tau_\epsilon} d\eta_{\mathbf{u}[\mathbf{x}]}(\mathrm{M}_{\mathbf{u}[\mathbf{x}]}) d\mathbf{x} \\ &+ \int_{\Omega_\mathrm{F}} \mathbb{H}\left\{\eta_{\mathbf{u}[\mathbf{x}]}(\mathrm{M}_{\mathbf{u}[\mathbf{x}]})\right\} d\mathbf{x} - \frac{E_\mathrm{reg}(\mathbf{u})}{2\tau_\mathrm{reg}} + \text{const.}\end{aligned}$$

Now, under the spatial iid assumptions of the model, the integral

$$\int_{\Omega_\mathrm{F}} \mathbb{H}\left\{\eta_{\mathbf{u}[\mathbf{x}]}(\mathrm{M}_{\mathbf{u}[\mathbf{x}]})\right\} d\mathbf{x} = \int_{\Omega_\mathrm{F}} \int_{\mathbb{I}_\mathrm{F}} p\left(f_{\mathbf{u}[\mathbf{x}]} \mid \mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right) \ln p\left(f_{\mathbf{u}[\mathbf{x}]} \mid \mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right) df_{\mathbf{u}[\mathbf{x}]} d\mathbf{x}$$

is equal to $\mathrm{vol}\{\Omega_\mathrm{M}\}\mathbb{H}\left\{f_{\mathbf{u}[\mathbf{x}]} \mid m_\mathbf{x}\right\}$, where

$$\mathbb{H}\left\{f_{\mathbf{u}[\mathbf{x}]} \mid m_\mathbf{x}\right\} = \int_{\mathbb{I}_\mathrm{M}} p\left(\mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right) \int_{\mathbb{I}_\mathrm{F}} p\left(f_{\mathbf{u}[\mathbf{x}]} \mid \mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right) \ln p\left(f_{\mathbf{u}[\mathbf{x}]} \mid \mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right) df_{\mathbf{u}[\mathbf{x}]} d\mathrm{M}_{\mathbf{u}[\mathbf{x}]},$$

is the conditional entropy of the link-process posterior.

Putting it all together, the log-posterior of the deformation model becomes:

$$\begin{aligned}\ln p\left(\mathbf{u} \mid \mathrm{F}, \mathrm{M}\right) = &- \frac{1}{2\tau_\epsilon} \int_{\Omega_\mathrm{F}} \int_{\mathbb{I}_\mathrm{F}} p\left(f_{\mathbf{u}[\mathbf{x}]} \mid \mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right) \left|\mathrm{F}_\mathbf{x} - f_{\mathbf{u}[\mathbf{x}]}\right|^2 df_{\mathbf{u}[\mathbf{x}]} d\mathbf{x} \\ &+ \mathrm{vol}\{\Omega_\mathrm{F}\}\mathbb{H}\left\{f_{\mathbf{u}[\mathbf{x}]} \mid \mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right\} - \frac{E_\mathrm{reg}(\mathbf{u})}{2\tau_\mathrm{reg}} + \text{const.}\end{aligned} \tag{5}$$

It can be easily seen that this distribution satisfies an intensity transformation invariance, *i.e.* $p(\mathbf{u}|\mathrm{F}, \mathrm{M}) = p(\mathbf{u}|\mathrm{F}, \alpha(\mathrm{M}))$ where $\alpha : \mathbb{I}_\mathrm{M} \to \mathbb{R}$ is a monotonic transfer function on the moving image intensity range such that $\alpha'(m) \neq 0$, at all $m \in \mathbb{I}_\mathrm{M}$.

2.2 Estimating the Link Process Posterior

The marginal posterior $p\left(f_{\mathbf{u}[\mathbf{x}]} \mid \mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right)$ of the link process $\eta_{\mathbf{u}[\mathbf{x}]}(\mathrm{M}_{\mathbf{u}[\mathbf{x}]})$ (§ $eqn.$ (3)), are obtained in non-parametric form using kernel density estimation (KDE) [8]:

$$p\left(f_{\mathbf{u}[\mathbf{x}]} | \mathrm{M}_{\mathbf{u}[\mathbf{x}]}\right) = \frac{\sum_{i=1}^N k_1(f_{\mathbf{u}[\mathbf{x}]} - f_i) k_2(\mathrm{M}_{\mathbf{u}[\mathbf{x}]} - m_i)}{\sum_{j=1}^N k_2(\mathrm{M}_{\mathbf{u}[\mathbf{x}]} - m_j)}, \tag{6}$$

where k_1 and k_2 are two non-negative symmetric kernel functions that integrate to unity, with scales h_1 and h_2 respectively. And, $f_i \triangleq \mathrm{F}[\mathbf{x}_i]$ and $m_i \triangleq \mathrm{M}_{\mathbf{u}[\mathbf{x}_i]}$ are fixed and moving image values sampled at locations $\mathbf{x}_i,\ i = 1 \dots N$.

Therefore,

$$\int_{\Omega_\mathrm{F}} \int_{\mathbb{I}_\mathrm{F}} p\left(f_{\mathbf{u}[\mathbf{x}]} \middle| \mathrm{M}_{\mathbf{u}[\mathbf{x}]}, \mathbf{u}, \mathrm{F}, \mathrm{M}\right) \left|\mathrm{F}_\mathbf{x} - f_{\mathbf{u}[\mathbf{x}]}\right|^2 df_{\mathbf{u}[\mathbf{x}]} d\mathbf{x} = \int_{\Omega_\mathrm{F}} \frac{\sum_{i=1}^N k_2(\mathrm{M}_{\mathbf{u}[\mathbf{x}]} - m_i)\left|\mathrm{F}_\mathbf{x} - f_i\right|^2}{\sum_{j=1}^N k_2(\mathrm{M}_{\mathbf{u}[\mathbf{x}]} - m_j)} d\mathbf{x} + h_1^2 \mathrm{vol}\{\Omega_\mathrm{F}\},$$

and

$$\int_{\mathbb{I}_\mathrm{M}} \mathbb{H}\left\{f_{\mathbf{u}[\mathbf{x}]} \middle| m\right\} p\left(m \mid \mathbf{u}, \mathrm{F}, \mathrm{M}\right), d\mathrm{M}_{\mathbf{u}[\mathbf{x}]} = -\frac{1}{N} \sum_{i=1}^N \ln \frac{\sum_{j=1}^N k_1(f_i - f_j) k_2(m_i - m_j)}{\sum_{j=1}^N k_2(m_i - m_j)}.$$

As a result, the KDE version of log-posterior of the transformation model is:

$$\begin{aligned} \ln p\left(\mathbf{u} \mid \mathrm{F}, \mathrm{M}\right) = &- \frac{1}{2\tau_\epsilon} \int_{\Omega_\mathrm{F}} \frac{\sum_{i=1}^N k_2(\mathrm{M}_{\mathbf{u}[\mathbf{x}]} - m_i)\left|\mathrm{F}_\mathbf{x} - f_i\right|^2}{\sum_{j=1}^N k_2(\mathrm{M}_{\mathbf{u}[\mathbf{x}]} - m_j)} d\mathbf{x} \\ &- \frac{\mathrm{vol}\{\Omega_\mathrm{F}\}}{N} \sum_{i=1}^N \ln \frac{\sum_{j=1}^N k_1(f_i - f_j) k_2(m_i - m_j)}{\sum_{j=1}^N k_2(m_i - m_j)} - \frac{E_\mathrm{reg}(\mathbf{u})}{2\tau_\mathrm{reg}} + \text{const.} \end{aligned} \tag{7}$$

3 Results

In this section, we show results registering both synthetic and clinical multi-modal images. A tetrahedral finite-element (FE) model together with a bio-mechanically plausible elastic energy penalty on mesh deformations was used in the experiments. The posterior distribution on deformations was characterized by the Metropolis-Hastings (MH) Markov Chain Monte Carlo (MCMC) method described in [5].

3.1 Synthetic Data

From the BrainWeb[1] database, we acquired simulated T1 and T2 weighted MR images of the brain that are in perfect alignment. The images were resampled to a resolution of $2 \times 2 \times 4$ mm and size $90 \times 108 \times 45$ voxels, and were Gaussian smoothed with 1mm variance. Normally distributed white noise of standard deviation of 0.02 was added to the T1 weighted image which was treated as the fixed image in all the synthetic experiments. Two synthetic moving images were created by applying the same b-spline deformation field (maximum and average displacement of 10.1mm and 4.3mm respectively) to (a) the T1 image; and (b) the T2 image, as shown in Fig. 2.

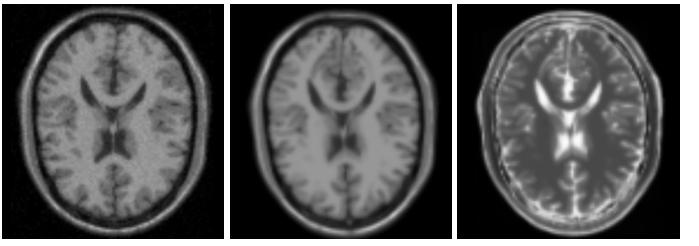

(a) Fixed T1 (b) Deformed T1 (c) Deformed T2

Fig. 2. The three images used in the synthetic experiments. **(a)**: The T1-weighted MR image. **(b)**: The deformed T1-weighted image. **(c)**: The deformed T2-weighted image.

For registration, an FE-based deformation model was employed where the moving image domain containing brain tissue was discretized with 104 tetrahedral elements and 44 FE vertices (giving 3×44 deformation parameters in 3D). The posterior distribution on deformations was characterized by MH-MCMC sampling with a noise temperature of $\tau_\epsilon = 0.04$ and prior temperature of $\tau_{\text{reg}} = 200$. The bandwidths h_1 and h_2 of the kernels k_1 and k_2 were selected using cross-validation. A total of $800{\times}10^3$ samples were generated for each MCMC chain, and with burn-in of $300{\times}10^3$ samples and a thinning factor of 10, effectively giving $50{\times}10^3$ samples from the posterior distribution in each chain.

The posterior mode (MAP), serving as a point estimate, and inter-quartile ranges (IQR), serving as a measure of uncertainty, were computed from the MCMC samples using kernel density estimation for each of the 3×44 components of the deformation field. In Fig. 3(a), the error in the T1-T1 registration versus that in the T1-T2 registration (determined with respect to the ground truth deformations) of the MAP estimate shown. For the T1-T1 experiment, the maximum and median absolute error was 1.6mm and 0.37mm respectively, while for the T1-T2 experiment it was 1.6mm and 0.45mm. Here, we can observe a strong linear relationship between the errors in the estimating same deformation from two different modalities ($r^2 = 0.80$). The T1-T1 registration IQRs are plotted against the T1-T2 registration IQRs in Fig. 3(b). The IQRs are highly correlated across modalities ($r^2 = 0.89$), but the correlation diminishes at higher uncertainties. For the T1-T1 case, the maximum and median IQRs were 0.67mm and 0.34mm respectively, while for the T1-T2 case they were 0.76mm and 0.37mm. These results imply that there is a slight but statistically insignificant decrease in registration accuracy and precision for multi-modal data (one-sided two-sample t-test, no effect for any $p < 0.26$).

For all $3{\times}44$ deformation components, the KL-divergence between their posterior distributions from the T1-T1 and T1-T2 cases is graphed in Fig. 4. It can be seen that the posterior distributions over most components are very similar but with a few outliers. Estimating the null distribution of KL-divergences by bootstrapping from the MCMC chain of the T1-T1 registration case, the difference between the posterior distributions of the T1-T1 and T1-T2 cases were not significant at any $p < 0.31$ (false discovery rate corrected).

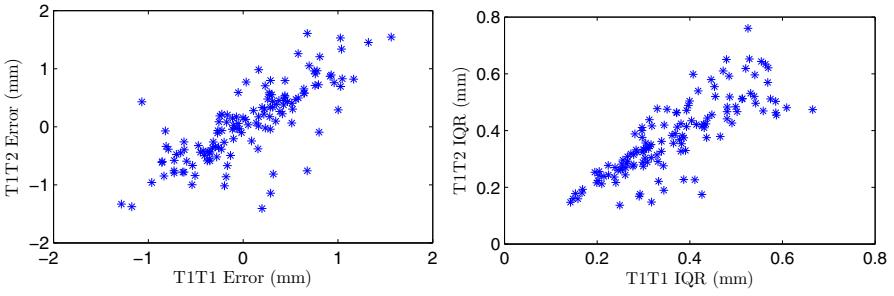

(a) Error (b) IQR

Fig. 3. (a): Error in MAP estimates of the T1-T1 versus that from T1-T2 registration cases (with respect to ground truth) plotted for each of the 3×44 vertex deformation components. **(b)**: The IQRs posterior distributions of each displacement component for the T1-T1 case versus the T1-T2 case.

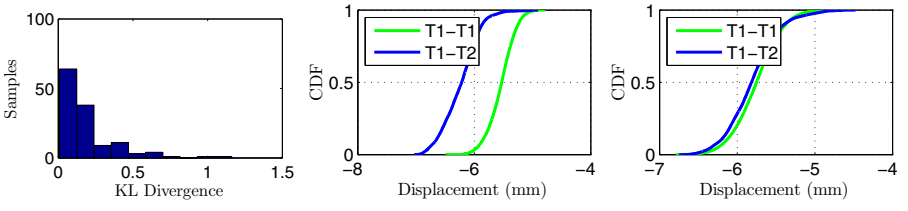

(a) Histogram of $\mathbb{KL}$ divergences. (b) Cdf for maximum $\mathbb{KL}$-divergence. (c) Cdf for median $\mathbb{KL}$-divergence.

Fig. 4. (a): Histogram of the KL-divergences between the posterior distributions (per displacement component) of the T1-T1 and T1-T2 registration cases. The median and maximum KL-divergences were 0.2 and 1.1 respectively. (b) Cumulative distribution function (cdf) of the posterior distribution corresponding to the displacement component with the maximum KL-divergence. (c) Cdf of the posterior distribution corresponding to the displacement component with the median KL-divergence.

3.2 Clinical Data

In ablation therapy of liver tumors, the procedure is often planned on pre-operative MR images which provide superior soft-tissue contrast while CT is used for intra-operative guidance. Although registering the pre-operative MR with the intra-operative CT enables real-time guidance of the ablation probe using the enhanced contrast provided by MR, the uncertainty in the results can provide equally important information for the decisions of the surgeon. Next, we demonstrate the quantification of this uncertainty using a data-set obtained during such a procedure (§ Fig. 5). A T1-weighted MR image (size: $512\times512\times96$, spacing: $0.8 \times 0.8 \times 2.5$ mm) was acquired pre-procedurally, while a CT image (size: $512\times512\times41$, spacing: $1.0 \times 1.0 \times 5.0$ mm) was acquired intra-procedurally. The MR image was sub-sampled by a factor of 2 in all dimensions, while the CT image was subsampled by a factor of 2 in the x- and y-directions. Both

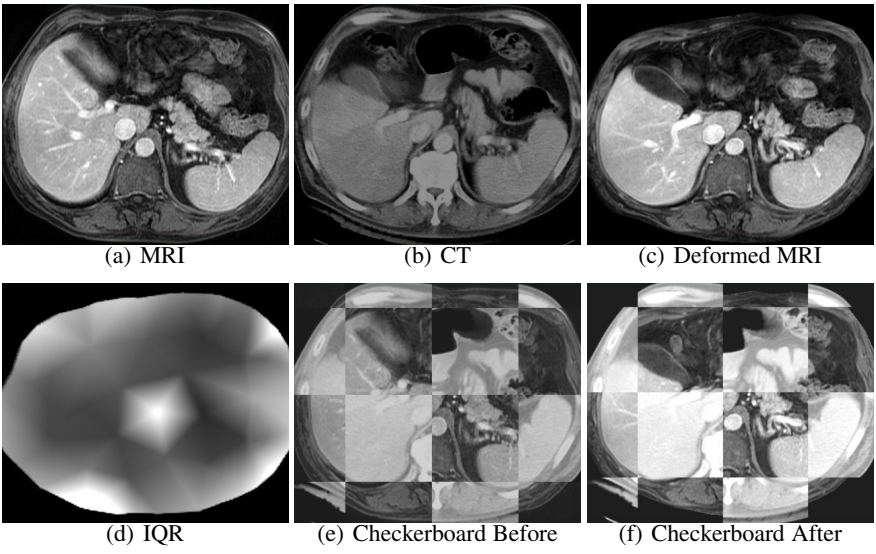

(a) MRI (b) CT (c) Deformed MRI

(d) IQR (e) Checkerboard Before (f) Checkerboard After

Fig. 5. Registration of pre-operative abdominal MRI in **(a)** with the CT in **(b)** acquired prior to insertion of cryo probe. **(c)**: The registered MRI. **(d)**: Spatial IQR along the z-direction. The brightest spot in the image has IQR of 3.8mm, while darkest has IQR of 1.0mm. **(e)**: Original MR image checker-boarded with CT image. **(f)**: Registered MR image checker-boarded with CT image. Notice that the boundaries liver and spleen are well aligned after registration.

images were smoothed with a Gaussian filter of 2.0mm variance and intensities were normalized between the $[0, 1]$ interval. The anatomy in the MR image was fitted with an FE-mesh consisting of 155 vertices and 512 tetrahedra. Starting from a manually determined rigid alignment of the images, 10^6 deformation samples were generated through MH-MCMC, with a burn-in factor of 50% and a thinning factor of 10. The remaining 50×10^3 samples were used to compute the posterior statistics.

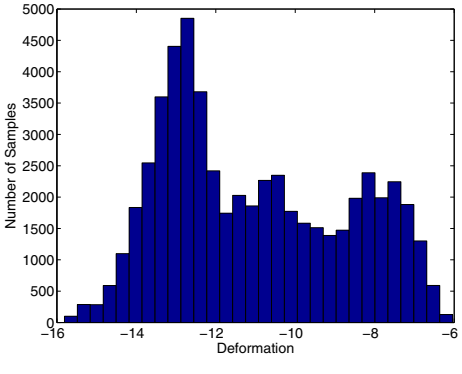

Fig. 6. Marginal probability distribution on displacements (in mm) along the z-direction for the FE-node located in the zone with high uncertainty (IQR)

Fig. 5 shows qualitative results from the alignment as well as corresponding uncertainty estimates. It can be seen that that there is relatively high uncertainty in the center of the slice in the abdominal aortal region. The marginal distribution over displacements in the z-direction of one FE-node in this location of high uncertainty is shown in Fig. 6. It can be observed that it has one distinct mode, but in addition has two two other, smaller, modes.

4 Conclusion

In this paper, we have presented a principled approach to quantifying the uncertainty associated with multi-modal image registration, based on a forward model of the image generation process. This approach augments the standard Bayesian framework for same-mode registration by introducing a stochastic link process that maps moving a image intensity to the fixed image intensity range, and is associated with a probability measure. This can capture a wider range of complex relationships than possible by using a parametric or specific functional representation of the map, similar to the MI metric. The framework furthermore marginalizes out the link-process using the free-energy equivalence. Therefore, in contrast to MI, the fully specified Bayesian model enables measuring the posterior over transformations, without dependencies on the intensity transfer function.

The formulation presented here and by Roche *et al.* [6] are examples of kernel regression. The main difference is that they use kernel regression to estimate the expected value of the fixed image intensity given moving image intensity, while we use kernel regression to estimate the expected difference between the observed and predicted fixed image values. In a regression framework, this is equivalent to the difference between using only the squared bias of the estimator (*i.e.* of the intensity transformation) as in [6] versus using the full mean squared error (MSE) as here, in the image similarity function. Moreover, marginalizing out the link process requires including the entropy of the conditional distribution in the cost function.

The differentiability of the kernel used in eqn. (7) permits computation of gradients of $p(\mathbf{u} \mid \mathrm{F}, \mathrm{M})$. Therefore, we can perform direct MAP estimation of the registration parameters, without MCMC sampling in an expectation maximization framework. In the n-th iteration of EM, the E-step computes the link-process posterior $p\left(f_{\mathbf{u}^{(n)}[\mathbf{x}]} \mid \mathrm{M}_{\mathbf{u}^{(n)}[\mathbf{x}]}\right)$, while the M-step optimizes $p(\mathbf{u} \mid \mathrm{F}, \mathrm{M})$ with respect to $\mathbf{u}$.

Acknowledgements. This work was partially funded by the NIH grants P41EB015898, P41RR019703, P41-RR-013218 and P41-EB-015902.

References

1. Brainweb database, `http://www.bic.mni.mcgill.ca/brainweb/` 5
2. Gee, J.C., Bajcsy, R.K.: Elastic matching: Continuum mechanical and probabilistic analysis. In: Brain Warping, p. 183. Academic Press 2

3. Guimond, A., Roche, A., Ayache, N., Meunier, J.: Three-dimensional multimodal brain warping using the demons algorithm and adaptive intensity corrections. IEEE Trans. Med. Imaging 20(1), 58–69 (2001) 3
4. Janoos, F., Risholm, P., Wells, W.M.: Robust non-rigid registration and characterization of uncertainty. In: Zhou, K., Duncan, J.S., Ourselin, S. (eds.) Methods in Biomedical Image Analysis (MMBIA), vol. 1 (2012) 2
5. Risholm, P., Pieper, S., Samset, E., Wells III, W.M.: Summarizing and Visualizing Uncertainty in Non-rigid Registration. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010, Part II. LNCS, vol. 6362, pp. 554–561. Springer, Heidelberg (2010) 2, 5
6. Roche, A., Malandain, G., Pennec, X., Ayache, N.: The Correlation Ratio as a New Similarity Measure for Multimodal Image Registration. In: Wells, W.M., Colchester, A.C.F., Delp, S.L. (eds.) MICCAI 1998. LNCS, vol. 1496, pp. 1115–1124. Springer, Heidelberg (1998) 3, 9
7. Rogelj, P., Kovačič, S., Gee, J.C.: Point similarity measures for non-rigid registration of multi-modal data. Comput. Vis. Image Underst. 92, 112–140 (2003) 3
8. Silverman, B.: Density Estimation for Statistics and Data Analysis. Chapman & Hall, London (1998) 5
9. Simpson, I.J., Schnabel, J.A., Groves, A.R., Andersson, J.L., Woolrich, M.W.: Probabilistic inference of regularisation in non-rigid registration. NeuroImage (2011) 2
10. Zöllei, L., Jenkinson, M., Timoner, S., Wells, W.: A Marginalized MAP Approach and EM Optimization for Pair-Wise Registration. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 662–674. Springer, Heidelberg (2007) 3