Probabilistic non-linear registration with spatially adaptive regularisation☆

I.J.A. Simpson [a,b,*], M.J. Cardoso [a,b], M. Modat [a,b], D.M. Cash [b], M.W. Woolrich [d,e], J.L.R. Andersson [e], J.A. Schnabel [c], S. Ourselin [a,b], for the Alzheimer's Disease Neuroimaging Initiative

[a] Centre for Medical Image Computing, University College London, United Kingdom
[b] Dementia Research Centre, University College London, United Kingdom
[c] Institute of Biomedical Engineering, University of Oxford, United Kingdom
[d] Oxford Centre for Human Brain Activity, University of Oxford, United Kingdom
[e] Centre for Functional Magnetic Resonance Imaging of the Brain, University of Oxford, United Kingdom

ARTICLE INFO

ABSTRACT

This paper introduces a novel method for inferring spatially varying regularisation in non-linear registration. This is achieved through full Bayesian inference on a probabilistic registration model, where the prior on the transformation parameters is parameterised as a weighted mixture of spatially localised components. Such an approach has the advantage of allowing the registration to be more flexibly driven by the data than a traditional globally defined regularisation penalty, such as bending energy. The proposed method adaptively determines the influence of the prior in a local region. The strength of the prior may be reduced in areas where the data better support deformations, or can enforce a stronger constraint in less informative areas. Consequently, the use of such a spatially adaptive prior may reduce unwanted impacts of regularisation on the inferred transformation. This is especially important for applications where the deformation field itself is of interest, such as tensor based morphometry. The proposed approach is demonstrated using synthetic images, and with application to tensor based morphometry analysis of subjects with Alzheimer's disease and healthy controls. The results indicate that using the proposed spatially adaptive prior leads to sparser deformations, which provide better localisation of regional volume change. Additionally, the proposed regularisation model leads to more data driven and localised maps of registration uncertainty. This paper also demonstrates for the first time the use of Bayesian model comparison for selecting different types of regularisation.

1. Introduction

Non-linear image registration is a fundamental tool in medical image analysis with a great many applications (Sotiras et al., 2013). One widely explored application of non-linear registration is the analysis of human brain morphology from structural magnetic resonance (MR) images. In this context, non-linear image registration has been used to accurately quantify localised cross-sectional differences between populations, such as subjects with Alzheimer's disease (AD) compared to normal ageing. It has also been used to measure longitudinal changes within individuals. Differences in morphology between populations can be identified using approaches such as tensor based morphometry (TBM) (Ashburner and Friston, 2000; Chung et al., 2001), where statistical analysis is performed on the Jacobian tensor of deformation fields calculated from registering individual subjects to a common space. TBM offers a whole brain approach to statistical analysis, and has the potential to extract rich features that accurately summarise anatomical differences.

TBM features are wholly defined by the registration process, which is complicated by the fact that non-linear registration is an ill-posed problem. In a typical structural MR image there are more than one million voxels in the human brain, where the intensity of a voxel is a noisy surrogate of tissue type. As such, there is a great deal of ambiguity in matching intensities, making it implausible for a unique voxelwise mapping to be determined purely from the image data.

1.1. Regularisation

As no unique mapping can be determined purely from the data, a "reasonable" mapping between images is sought. This is achieved through the use of a data matching term and regularisation, which maximises the similarity of image appearance whilst maintaining a plausible deformation, i.e. with an appropriate magnitude of displacement and spatial smoothness. Regularisation can be considered as a prior on the set of expected deformations, which reduces the space of potential solutions and hence limits the variance of any estimated solution. The form of bias induced by the prior is generally selected based on some physical model of deformation, such as linear elasticity (Miller et al., 1993) or thin-plate spline bending energy (Bookstein, 1997).

Regularisation models are commonly described as having the same effect across the image. However, such models may well be unreasonable in brain registration for two reasons: Firstly, different regions of the image contain different amounts of information. Uninformative image areas should be strongly influenced by the priors as they contain little information, whereas feature-rich regions should be given more freedom. Furthermore, the magnitude of anatomical mis-correspondence is likely to be variable across space, and some regions will require more complex deformations than others to allow an adequate mapping. Therefore, the use of a global spatial regularisation prior may introduce either an unwanted or insufficient bias on the deformation in certain image regions. This could have substantial adverse effects on an application, such as TBM, which directly relies on the interpretability of the deformation field.

1.1.1. Previous approaches to spatially varying regularisation in registration

There have been several previous works on the use of spatially varying regularisation in non-linear registration. These include approaches that vary based on tissues or structures derived from segmentations (Lester et al., 1999; Davatzikos, 1997; Staring et al., 2007; Schmah et al., 2013). These approaches are ideal in cases when an informative deformation prior is known for a specific region or tissue type, which can be robustly defined. However, in the majority of registration applications, this is unlikely to be the case.

More data driven approaches have been proposed, which include anisotropic smoothing of image similarity gradients according to image information (Hermosillo et al., 2002; Papież et al., 2013). Alternative approaches include weighting similarity gradients based on measures of local image reliability (Tang et al., 2010). These approaches allow the image information to affect the local regularisation strength, although are still somewhat ad-hoc, being dependent on the definition of a heuristic weighting between regularisation and data fidelity.

Inference of geometric deviation from an estimated atlas for use as a spatial prior is an alternative approach to define regularisation priors, Allassonniére et al. (2007) proposed a small deformation Bayesian framework for atlas estimation and registration. Gori et al. (2013) proposed a Bayesian approach for estimating an atlas and structure specific regularisation terms for a registration model based on the metric of currents. A recently published approach by Xu et al. (2014) propose a method for deriving an average atlas and a spatial distance metric based on the geometric variability of the atlas. Zhang et al. (2013) proposed a generative registration model using Geodesic shooting for atlas and regularisation estimation, this work was extended to sparsely estimate the principal geodesic modes of variation (Zhang and Fletcher, 2014). Durrleman et al. (2013) also estimate sparse parametrisations of variability from an estimated atlas.

Most similarly to this work, Risholm et al. (2010b, 2013) presented a Bayesian inference scheme that allows linear elastic parameters to be inferred from the data. These parameters can also vary spatially, as demonstrated by Risholm et al. (2011b). This approach does not require the definition of strong heuristics, although informative priors are required for the elastic model parameters. The limitations of the framework lie in the numerical integration inference strategy, which comes with vast computational complexity. Modern sampling techniques may help alleviate this burden (Zhang et al., 2013).

1.2. Contribution of this paper

This paper proposes a novel non-linear registration model and Bayesian inference scheme that allows for data-driven spatially varying regularisation. This approach alleviates the difficulties associated with previous attempts at spatially varying regularisation. Firstly, it is fully data driven, requiring no segmentations or informative priors. Secondly, the trade-off of data fidelity and regularisation is inferred directly from the data and finally, inference is tractable.

This work follows from our previous conference paper (Simpson et al., 2013b), with a second-order inference scheme for the regularisation parameters, a full mathematical derivation and broader validation. Additionally, this paper investigates objective Bayesian model comparison and the effects of the spatially varying prior on registration uncertainty. The proposed framework describes registration using a hierarchical probabilistic model, with a transformation prior that is parameterised by a set of hyper-parameters. Each hyper-parameter influences a spatially localised region of the prior. Through the use of full Bayesian inference, posterior distributions of hyper-parameter weights can be inferred alongside the transformation. This allows the effects of the prior to be locally determined during the registration.

This approach is demonstrated through an application of TBM on synthetic images, as well as comparing subjects with AD to healthy controls. Our results demonstrate the strength of our approach in terms of reducing false positive results, which may improve interpretability. We also highlight additional benefits of the proposed framework including: objective comparison of regularisation models, and more reasonable uncertainty estimates of the deformation fields.

2. Method

2.1. Model

Image registration can be described in a probabilistic manner using a generative model of the target image, $\mathbf{y}$, which is predicted by the deformed source image, $\mathbf{t}(\mathbf{x}, \mathbf{w})$. Here, $\mathbf{t}$ is a transformation model, $\mathbf{x}$ is the source image and $\mathbf{w}$ parametrises the transformation. In this paper, a cubic B-spline free form deformation model (Rueckert et al., 1999; Andersson et al., 2007) is used for $\mathbf{t}$, with $\mathbf{w}$ corresponding to the control point displacement. However, in principle any deformation model could be used.

The generative model also contains an additive noise term, $\mathbf{e}$, which describes the error in model fit. In this work, $\mathbf{e}$, is modelled as independently and identically distributed across voxels and follows a normal distribution:

$$\mathbf{e} \approx \mathcal{N}(0, \mathbf{I}\phi^{-1}\alpha), \tag{1}$$

where $\mathbf{I}$ is an identity matrix the size of the number of voxels, N_v. ϕ corresponds to the noise precision (inverse variance) of the additive Gaussian noise under the assumption of being independently distributed. α corresponds to the virtual decimation factor (Groves et al., 2011), which is a data driven term used to compensate for spatial covariance in the residual, weakening the assumption of independent noise. The assumption of identically distributed noise could also be relaxed in this approach as in Simpson et al. (2012a). The full generative model for registration is therefore given as:

$$\mathbf{y} = \mathbf{t}(\mathbf{x}, \mathbf{w}) + \mathbf{e}. \tag{2}$$