Medical Image Analysis 27 (2016) 57–71

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Deformable image registration by combining uncertainty estimates from supervoxel belief propagation

Mattias P. Heinrich[a,*], Ivor J.A. Simpson[b], BartŁomiej W. Papież[c], Sir Michael Brady[d], Julia A. Schnabel[c]

[a] *Institute of Medical Informatics, Universität zu Lübeck, Germany*
[b] *Centre for Medical Image Computing, University College London, UK*
[c] *Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, UK*
[d] *Department of Oncology, University of Oxford, UK*

ARTICLE INFO

Article history:
Received 15 May 2014
Revised 20 September 2015
Accepted 22 September 2015
Available online 19 October 2015

Keywords:
Supervoxel layers
Segmentation propagation
Motion estimation
Mean-shift
Probabilistic registration

ABSTRACT

Discrete optimisation strategies have a number of advantages over their continuous counterparts for deformable registration of medical images. For example: it is not necessary to compute derivatives of the similarity term; dense sampling of the search space reduces the risk of becoming trapped in local optima; and (in principle) an optimum can be found without resorting to iterative coarse-to-fine warping strategies. However, the large complexity of high-dimensional medical data renders a direct voxel-wise estimation of deformation vectors impractical. For this reason, previous work on medical image registration using graphical models has largely relied on using a parameterised deformation model and on the use of iterative coarse-to-fine optimisation schemes. In this paper, we propose an approach that enables accurate voxel-wise deformable registration of high-resolution 3D images without the need for intermediate image warping or a multi-resolution scheme. This is achieved by representing the image domain as multiple comprehensive supervoxel layers and making use of the full marginal distribution of all probable displacement vectors after inferring regularity of the deformations using belief propagation. The optimisation acts on the coarse scale representation of supervoxels, which provides sufficient spatial context and is robust to noise in low contrast areas. Minimum spanning trees, which connect neighbouring supervoxels, are employed to model pair-wise deformation dependencies. The optimal displacement for each voxel is calculated by considering the probabilities for all displacements over all overlapping supervoxel graphs and subsequently seeking the mode of this distribution. We demonstrate the applicability of this concept for two challenging applications: first, for intra-patient motion estimation in lung CT scans; and second, for atlas-based segmentation propagation of MRI brain scans. For lung registration, the voxel-wise mode of displacements is found using the mean-shift algorithm, which enables us to determine continuous valued sub-voxel motion vectors. Finding the mode of brain segmentation labels is performed using a voxel-wise majority voting weighted by the displacement uncertainty estimates. Our experimental results show significant improvements in registration accuracy when using the additional information provided by the registration uncertainty estimates. The multi-layer approach enables fusion of multiple complementary proposals, extending the popular fusion approaches from multi-image registration to probabilistic one-to-one image registration.

1. Introduction

Medical image registration aims to find spatial correspondences between scans of different patients, modalities, the time course of a disease, or response to therapy. It also forms an integral part of many medical image analysis applications. For example, intra-patient deformable registration can be used to relate two scans, e.g. a pre-treatment planning scan to an intra-operative image for guiding an intervention. Longitudinal scans can be employed to monitor treatment or disease progression. Computed tomography (CT) scans are now used widely for motion estimation in radiotherapy planning in order to increase the accuracy of dose delivery (Weiss et al., 2007). To study the functionality or anatomical variability of human brains, registration-based segmentation propagation is widely used to

* Corresponding author. Institute of Medical Informatics, Ratzeburger Allee 160, 23562 Lübeck, Germany.
E-mail address: heinrich@imi.uni-luebeck.de, mattias@gmx.at (M.P. Heinrich).
URL: http://www.mpheinrich.de (M.P. Heinrich)

http://dx.doi.org/10.1016/j.media.2015.09.005
1361-8415/

automatically label structures in magnetic resonance images (MRI) (Klein et al., 2009). It can be used to measure the volume and shape of anatomical structures in the human brain.

Image registration algorithms are generally based on three components: a similarity metric, a transformation model, and an optimisation strategy. A large variety of approaches have been proposed for the medical image domain over the past few years (see (Sotiras et al., 2013) for an overview). Different assumptions have been made to model these three components to obtain robust, accurate and also computationally efficient algorithms. In the following, we discuss four challenges, which are, in our opinion, prevalent in current approaches.

First, the similarity metric measures the data affinity and can assume many different forms depending on the medical application and modality. Ideally, the choice of the similarity measure is affected by neither the chosen transformation model nor the employed optimisation technique. In practice, however, this is often not the case, since gradient-based optimisation techniques require the similarity term to be (at least first-order) differentiable.

Second, the transformation model may restrict the deformation between two images either to obey a certain physically motivated model (for example finite element models (FEM) (Ferrant et al., 2000)) or to be well approximated by a particular mathematical model (such as free-form transformations (Rueckert et al., 1999)). Due to its simplicity, most common registration algorithms define a transformation model on a regular, equally spaced grid. However, sparsely or irregularly spaced models, which have been e.g. presented by Schnabel et al. (2001), Rohde et al. (2003) (with spline basis functions), Glocker et al. (2010) (triangular mesh) and Popuri et al. (2013) (FEM), might in many cases provide a more realistic model of real physical motion, in particular when the smoothness of deformations varies across the image domain.

Third, the chosen optimisation strategy impacts the space of obtainable deformation vectors. Continuous optimisation approaches yield excellent results for subtle sub-voxel changes across scans, which is important for the analysis of longitudinal brain development (Ashburner and Ridgway, 2012). A disadvantage is that they are susceptible to local minima, especially in the presence of large, complex motions.

Fourth, a further limitation of the majority of current methods is that they only estimate the most probable transformation (the *maximum a posteriori solution*). However, quantifying the uncertainty distribution of a registration can provide improvements with respect to the immediate goals, such as segmentation or motion vector estimation, as well as give a confidence measure of the generated results.

Recent work (Shekhovtsov et al., 2008; Glocker et al., 2008a; Heinrich et al., 2013a; Cobzas and Abhishek, 2011), has demonstrated a number of advantages of discrete optimisation techniques over the more commonly used continuous counterparts: they do not necessitate the computation of derivatives of the image similarity metric; they are computationally efficient; the space of displacements can be defined to capture a large range of deformations; and, under certain circumstances, local minima can be avoided. For a more in-depth discussion on optimality guarantees see Komodakis and Tziritas (2007) for linear programming and Felzenszwalb and Zabih (2011) for belief propagation techniques. Discrete approaches are, however, restricted to a quantisation of displacement vectors, which causes limits on the achievable accuracy. An extended review of deformable medical image registration using Markov random field formulations can be found in Glocker et al. (2011).

In this paper, we further contribute to medical image registration based on graphical models by introducing three new concepts, which have to date not been deeply explored. First, we make use of a more flexible image representation using supervoxel graphs. This is, to the best of our knowledge, the first time that multiple complementary layers have been used, which enable us to represent the complex nature of 3D deformations with spatially varying smoothness. Second, the probability (inversely, uncertainty) of a large set of potential displacements is calculated for our graphical model using belief propagation (with the min-sum algorithm). Supervoxels are inter-connected using a tree, which enables us to obtain marginal energies for every displacement and to regularise the displacement field using pair-wise interactions. Third, the complementary information from multiple supervoxel layers is combined on a voxel-wise level. A mode seeking algorithm over all potential (and probable) displacements for every voxel is used to find subvoxel accurate motion vectors (or the most likely fusion of many potential segmentation labels).

This paper builds upon previous work by the authors including the use of a graphical model that represents the image domain by a number of overlapping layers of supervoxels, which are connected by a minimum spanning tree (MST) (Heinrich et al., 2013b); the optimisation of the MST model using belief propagation (Heinrich et al., 2013a) and (Heinrich et al., 2012) and a similar calculation of min-marginal probabilities over all potential displacements (Heinrich et al., 2013d). Here, we unify those approaches, extend the motion estimation by the mean-shift mode seeking algorithm, and report substantial additional experimental validation.

2. Background

Supervoxel: Superpixel clustering describes the parcellation of the image domain into perceptually meaningful regions, which group pixels based on their appearance and spatial closeness. Ren and Malik, 2003 introduced the term superpixel, but previous work, such as watershed segmentation, have followed the same principle. Because they significantly reduce computational complexity, they have attracted a lot of attention in a range of image analysis tasks including optical flow estimation (Lei and Yang, 2009; Zitnick and Kang, 2007); and image segmentation (Shi and Malik, 2000). In medical imaging, 3D supervoxels have been introduced relatively recently for cell segmentation by Lucchi et al. (2012). Brain tumour segmentation has been addressed by Wang and Yushkevich (2013) using supervoxel matching without regularisation, following the superparsing framework of Tighe and Lazebnik (2013). Another recent approach by Tang and Hamarneh (2014) used supervoxels in an aggregation step for random walk registration. In Felzenszwalb and Huttenlocher (2006) a uniform hierarchical grouping of nodes was used to improve the convergence of belief propagation, while Willsky (2002) proposed organised trees of multiple scales in a pyramid form to solve large scale MRF problems.

Uncertainty estimates: Probabilistic registration methods based on continuous optimisation have been used to estimate the spatial variation of the displacements close to a local optimum to improve deformable registration based on locally adaptive smoothing (Simpson et al., 2011) and in order to boost classification accuracy (Simpson et al., 2013). This Bayesian framework was extended by Wassermann et al. (2014) for large diffeomorphic mappings. Iglesias et al. (2013) use uncertainty of registration parameters to improve segmentation propagation by using multiple probable warps from atlas to target volume. Registration uncertainties are also used in Risholm et al. (2013) to estimate the cumulated dose delivery in radiotherapy and in Kybic (2010) to estimate registration accuracy. The limitation of these approaches, based either on bootstrapping (Kybic, 2010), variational Bayes (Simpson et al., 2012), or Monte Carlo sampling (Iglesias et al., 2013; Risholm et al., 2013), is that a dense sampling of the uncertainty of the displacement space is impossible (or at least extremely computationally expensive) and distributions with multiple local optima cannot be easily dealt with. The AQUIRC method (Datteri and Dawant, 2012) has been used to estimate registration uncertainty (and thus accuracy) using transitivity errors in registration circuits (at least four pair-wise registrations) and