NeuroImage 106 (2015) 238–251

Contents lists available at ScienceDirect

NeuroImage

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

Avoiding symmetry-breaking spatial non-uniformity in deformable image registration via a quasi-volume-preserving constraint

Iman Aganj [a,*], Martin Reuter [a,b], Mert R. Sabuncu [a,b], Bruce Fischl [a,b,c]

[a] Athinoula A. Martinos Center for Biomedical Imaging, Radiology Department, Massachusetts General Hospital, Harvard Medical School, 149, 13th St., Room 2301, Charlestown, MA 02129, USA
[b] Computer Science and Artificial Intelligence Laboratory, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 32 Vassar St., Cambridge, MA 02139, USA
[c] Harvard-MIT Division of Health Sciences and Technology, 77 Massachusetts Ave., Room E25-519, Cambridge, MA 02139, USA

ARTICLE INFO

Article history:
Accepted 26 October 2014
Available online 30 October 2014

Keywords:
Deformable image registration
Volume-preserving constraints
Symmetry
Inverse-consistency
Integral non-uniformity

ABSTRACT

The choice of a reference image typically influences the results of deformable image registration, thereby making it asymmetric. This is a consequence of a spatially non-uniform weighting in the cost function integral that leads to general registration inaccuracy. The inhomogeneous integral measure – which is the local volume change in the transformation, thus varying through the course of the registration – causes image regions to contribute differently to the objective function. More importantly, the optimization algorithm is allowed to minimize the cost function by manipulating the volume change, instead of aligning the images. The approaches that restore symmetry to deformable registration successfully achieve inverse-consistency, but do not eliminate the regional bias that is the source of the error. In this work, we address the root of the problem: the non-uniformity of the cost function integral. We introduce a new quasi-volume-preserving constraint that allows for volume change only in areas with well-matching image intensities, and show that such a constraint puts a bound on the error arising from spatial non-uniformity. We demonstrate the advantages of adding the proposed constraint to standard (asymmetric and symmetrized) demons and diffeomorphic demons algorithms through experiments on synthetic images, and real X-ray and 2D/3D brain MRI data. Specifically, the results show that our approach leads to image alignment with more accurate matching of manually defined neuroanatomical structures, better tradeoff between image intensity matching and registration-induced distortion, improved native symmetry, and lower susceptibility to local optima. In summary, the inclusion of this space- and time-varying constraint leads to better image registration along every dimension that we have measured it.

Introduction

Image registration is a crucial step in numerous clinical and neuroscientific imaging studies involving the comparison of images, such as population investigations and longitudinal analyses. In the basic case with only two images involved, pairwise registration provides dense point-wise *correspondences* between voxels of the two input images. The set of such correspondences is often thought of as a *transformation* that takes each point in one image to the corresponding point in the other image. In that case, registration is said to *align* the two images; i.e., morph and overlay a *moving* image on a *fixed* image (or alternatively move both images) so that they appear similar or identical to each other. Interpreting registration as alignment requires the definition of a *reference space*, in which the images are aligned and compared. The choice of the reference influences the results, particularly when the transformation is assumed non-rigid (beyond merely translation and rotation) — as is necessary in most cross-subject registration applications. It is common to arbitrarily select the native space of one of the images (the space where the image is undistorted) as the reference space, hence the dependence of the results on the choice of the so-called *reference image*.[1] It must be noted that interpreting registration as alignment – and consequently selection of a reference image/space – is intrinsically unnecessary for many applications (e.g., comparison of local cortical thickness of two brain images, where correspondences established between the two images do not necessarily represent alignment). With no image designated as the reference, swapping the two images should not affect the resulting point-wise correspondences, making pairwise registration inherently *symmetric* with respect to the input

* Corresponding author at: 149, 13th St., Room 2301, Charlestown, MA 02129, USA.
E-mail addresses: iman@nmr.mgh.harvard.edu (I. Aganj), mreuter@nmr.mgh.harvard.edu (M. Reuter), msabuncu@nmr.mgh.harvard.edu (M.R. Sabuncu), fischl@nmr.mgh.harvard.edu (B. Fischl).
URL: http://nmr.mgh.harvard.edu/~iman (I. Aganj).

[1] Since the reference image is not interpolated during registration, we call the other image the *interpolated image*.

http://dx.doi.org/10.1016/j.neuroimage.2014.10.059
1053-8119/

images. Obtaining the same results after reversing the direction of registration – known as *inverse-consistency*[2] – is therefore necessary for a pairwise registration method to be considered reliable and unbiased.[3]

Pairwise deformable registration is performed by maximizing some measure of similarity between the corresponding regions of the two images (for a survey, see Sotiras et al., 2013). Since a perfect match cannot generally be achieved due to noise and anatomical variability, typically a local image *mismatch measure* aggregated over the entire space is minimized. Such *cost functions* (CFs), with the most common example being the sum of squared difference (SSD) of image intensities, require the images to be aligned in a reference space in which the mismatch measure is integrated, thus raising the question of how to select a suitable reference space. Registration results depend on the choice of the reference as a consequence of the non-rigid (specifically, volume-changing) nature of the transformation, since, as we will see in the Intrinsic non-uniformity in standard deformable image registration section, spatially uniform integration in the reference space is generally equivalent to non-uniform integration in (one or both) native spaces. In other words, for an arbitrary deformation, the integral of the mismatch measure has space-varying weightings in at least one of the native spaces of the images, with a weighting that varies depending on the deformation and the choice of the reference space. As we will see, the weighting in the CF integral is the Jacobian determinant of the transformation, which represents the local volume change in the deformation field. Assigning weighting to image regions introduces a regional bias. Furthermore, given that the weighting depends on the transformation, the optimization algorithm will drive the deformation also towards 'lowering the weighting of the mismatched areas', instead of only improving the alignment (e.g. by shrinking a region of one image so much as to make it almost vanish). Figure 1 presents an example of such a phenomenon, which results in a counterproductive increase in SSD in the native space of the interpolated image during registration. The dependence of the degree of expansion and shrinkage – representing the Jacobian determinant – on the choice of the reference image has also been observed by Cachier and Rey (2000), and has been reported to bias the quantification of the evolution of lesions in multiple sclerosis studies by hampering the equal retrieval of expanding and shrinking areas (Rey et al., 2002). These sources of error exist regardless of whether the reference space is chosen to be the native space of an input image or some sort of *mid-space*, as one or both of the native spaces of the images – which are the only physically meaningful spaces – will be integrated with regional bias.

In standard implementation of deformable registration the reference space is commonly chosen as the native space of one of the images (say, the first image), and the results are consequently influenced by the ordering of the input images, thereby breaking the symmetry of registration. The spurious dependence of the point-wise correspondences on the choice of the reference image has been shown to be related to a bias introduced into the estimation of Alzheimer's disease effects (Fox et al., 2011; Hua et al., 2011; Thompson and Holland, 2011; Yushkevich et al., 2010). In longitudinal studies in particular, favoring one time point over another may result in errors dominating the subtle changes one seeks to measure (Reuter et al., 2012). In addition, in radiation therapy, the implication of registration asymmetry has been discussed for daily dose computation (Yang et al., 2008) and auto re-contouring (Ye and Chen, 2009). To address this issue, existing approaches primarily aim to restore inverse-consistency to registration by computing the integral in both image spaces and taking the average (Alvarez et al., 2007; Bondar et al., 2010; Cachier and Rey, 2000; Christensen and Johnson, 2001; Chui, 2001; Feng et al., 2009; Geng, 2007; Gholipour et al., 2010; Leow et al., 2007; Modat et al., 2012; Mohagheghian et al., 2010; Sabuncu et al., 2009; Tagare et al., 2009; Tao et al., 2009; Trouvé and Younes, 2000; Vercauteren et al., 2008b; Zeng and Chen, 2008; Zhang et al., 2006) or computing the integral in an abstract mid-space chosen to be "in between" the native spaces of the images (Beg and Khan, 2007; Chen and Ye, 2010; Joshi et al., 2004; Lorenzen et al., 2004; Lorenzi et al., 2013; Noblet et al., 2008; Škrinjar et al., 2008; Yang et al., 2008; Ye and Chen, 2009). Other approaches based on similar ideas have been proposed in the literature, including (Ashburner et al., 1999, 2000; Avants et al., 2008; Basri et al., 1998; Christensen and Johnson, 2003; Dedeoglu and Kanade, 2005; He and Christensen, 2003; Rogelj and Kovačič, 2006; Yanovsky et al., 2008b; Yeung et al., 2008). Although these methods are effective in making the registration invariant to the ordering of the input images, they do not alleviate the time-varying regional bias, which is the source of the problem. Indeed, inverse-inconsistency is merely a *symptom* of non-uniform integration of the mismatch measure on the images, and symmetrization of registration does not necessarily eliminate the underlying cause (non-uniformity of the integral) and the consequent inaccuracies introduced into the registration. Our underlying hypothesis is that the freedom for the algorithm to minimize the CF by altering the Jacobian determinant instead of improving image matching leads to suboptimal registration solutions in terms of 1) reduced image intensity matching, 2) increased (and unnecessary) distortion in the warp field, and 3) the creation of local minima that reduce the accuracy of the resulting correspondences.

In this work, instead of symmetrizing the CF, we address the root of the problem: non-uniform integrals of CFs defined on the native spaces of images. We propose to restrict the deformation such that the integrals in the native spaces of the images are (almost) unweighted, except in regions where weighting contributes (almost) no error to the CF. Our adaptive constraint – which, as we will see is quasi-volume-preserving (QVP) – keeps the deformation field away from zones that would lead to non-uniformity-induced error, by limiting the local volume changes except for regions where image intensities match well. As a result, the proposed method yields overall improvement in the alignment (Results and discussion section) when incorporated in an SSD-like deformable registration algorithm.[4] Furthermore, a natural consequence of the QVP constraint is that the values of the *native CFs*, i.e. those with uniform integral on the native space of an input image, remain arbitrarily close to each other throughout the registration. This property, which we shall name *native symmetry*, is a stronger form of symmetry and is improved by our method. Native symmetry implies that *both* native CFs agree on the progress of registration, as opposed to only one of them (as in asymmetric registration, e.g. Fig. 1) or only the average of them (as in symmetrization). An additional advantage of restricting the deformation in dissimilar regions is helping to avoid entrapment of the iterative algorithm in local minima due to too much flexibility, thereby guiding it towards a good overall QVP fit before relaxing the constraints and achieving an optimum warp. This is particularly important in registration of medical images with possibly large anatomical variation. We will show improvement in registration in terms of better label alignment, better tradeoff between intensity matching and geometric distortion, native symmetry, and lower susceptibility to local optima, using two-dimensional (2D) non-diffeomorphic and three-dimensional (3D) diffeomorphic registration on several datasets.

This article extends our previous conference version (Aganj et al., 2013a). In particular, we provide a thorough theoretical justification (Spatial non-uniformity problem description section), more detailed description of the method (Proposed methods section), more comprehensive experimental validation (Results and discussion section), and further implementation details (appendices).

[2] The terms "symmetry" and "inverse-consistency" have been used interchangeably in most of the literature.

[3] We do not denote the images as "source" and "target", and therefore register Image A and Image B, rather than Image A to Image B. However, we do not consider atlas-to-subject mapping as "pairwise registration" here, since such a registration is intrinsically asymmetric (Sabuncu et al., 2009). By pairwise registration, we mean that the input images are real and in physically existent spaces, such as registering brain images of two subjects, or of one subject at different time points.

[4] Information theoretical objective functions also suffer from the problems described here. However, addressing them is beyond the scope of this paper.