

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}

{iman, mreuter, msabuncu, fischl}@nmr.mgh.harvard.edu

^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

* Corresponding author:

Iman Aganj

Email: iman@nmr.mgh.harvard.edu

Telephone: +1 (617) 726-4050

Web: <http://nmr.mgh.harvard.edu/~iman>

Address: 149, 13th St., Room 2301, Charlestown, MA 02129, USA

NeuroImage, October 2014

<http://doi.org/10.1016/j.neuroimage.2014.10.059>

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.

Keywords: Deformable image registration, volume-preserving constraints, symmetry, inverse-consistency, integral non-uniformity.

1. 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 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 Section 2.3, 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

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

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.

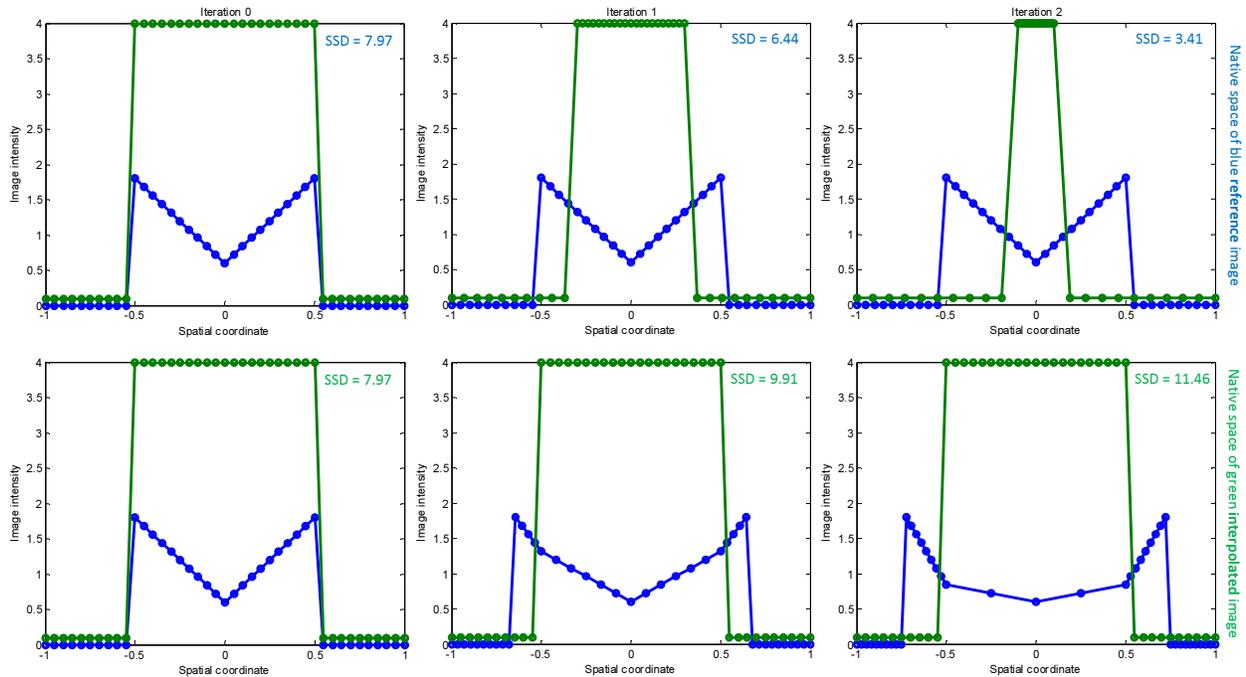


Figure 1. In this illustrative example, two one-dimensional images are registered by a given SSD-minimizing algorithm, with the blue image (the M-shaped intensity profile) chosen as the reference and the green one chosen as the interpolated (moving) image. The evolution of the registration is depicted from left to right, where the two rows visualize the process in the native spaces of the blue (top row) and the green (bottom row) images. Therefore, the deformations of the green (top row) and blue (bottom row) images are the inverse of each other. Given that the blue image matches better the background rather than the bright spot of the green image, the latter is shrunk in the reference space (top row). It can be seen that the SSD, while decreasing in the native space of the reference image (top row, where the CF integral is uniform), *increases* in the native space of the interpolated image (bottom row, where the CF integral is weighted by the volume change). The CF, being a non-uniform integral in the native space of the interpolated image, is here minimized by decreasing the weighting, even though resulting in an increase in the mismatch measure. (This toy example is provided to illustrate the disagreement on the decrease/increase in the SSD when the transformation is observed in the two native spaces, without commenting on whether the SSD-minimizing transformation here is the “expected” result. The two red curves in Figure 4 are another such example, yet using a symmetrization approach. For simplicity, we assume the regularization to be small relative to the data-attachment term.)

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 in Section 3.1, 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 (Section 4) 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

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

(as in asymmetric registration, e.g. Figure 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 (Section 2), more detailed description of the method (Section 3), more comprehensive experimental validation (Section 4), and further implementation details (appendices).

2. Spatial non-uniformity problem description

2.1. Definitions

The data attachment term of the standard asymmetric cost function (CF) that is minimized in deformable registration is defined as:

$$I_f(f, g, T) := \int_{\Omega} D(f(x), g \circ T(x)) dx, \quad (1)$$

where $\Omega \subseteq \mathbb{R}^d$ is the d -dimensional image space, $f, g: \Omega \mapsto \mathbb{R}$ are two input images, $T: \Omega \mapsto \Omega$ is the invertible transformation with respect to which the CF is minimized, and $D: \mathbb{R}^2 \mapsto \mathbb{R}_{\geq 0}$ is the (often symmetric) nonnegative intensity distance (mismatch measure); e.g., defining $D(u, v) := (u - v)^2$ results in the SSD CF.⁵ Some global regularization is also typically either added explicitly to this data-attachment CF, or performed at each iteration (see Appendix B). By searching for the transformation T that minimizes the CF, the optimization algorithm is expected to align corresponding regions of the two images.

Since image f is not deformed by T in the integral of Eq. (1), we say that the integral is represented in the *native space* of f . Image g , on the other hand, appears as $g \circ T$ in Eq. (1) and is generally deformed; however, by a simple change of variable $y := T(x)$, the same integral can be represented in the native space of g (see Eq. (3) below). We call I_f the *f -native CF*, since the integral of D is *unweighted* (taken *uniformly*) when represented in the native space of f .

2.2. Pitfalls of non-uniformity

We explain here the pitfalls of having a non-uniform (weighted) CF, which will help later to describe the key problem addressed in this paper. Let us assume that instead of Eq. (1), the following weighted version of it were minimized,

$$\min_T \int_{\Omega} D(f(x), g \circ T(x)) W_T(x) dx, \quad (2)$$

⁵ This is the particular form of the CFs aggregating a local mismatch measure, the most common of which being the SSD. Several other CFs, such as normalized mutual information, correlation coefficient, local correlation, etc., cannot be expressed with Eq. (1).

where $W_T: \Omega \mapsto \mathbb{R}_{\geq 0}$ is an arbitrary nonnegative weight function that generally depends on T . Allowing the image regions to contribute differently to the CF would introduce a multiplicative regional bias that would change throughout the registration process, depending on the input images and the registration parameters. More troublesome, however, would be the dependence of the bias on the transformation, thereby giving the optimization algorithm the ability to control the weighting in a counterproductive manner. More precisely, the algorithm would minimize such a weighted CF by:

- 1) searching for T that makes the transformed image better match the reference, $f(x) \approx g \circ T(x)$, in a larger subset of Ω , and/or
- 2) manipulating T so as to lower the weighting $W_T(x)$, particularly in regions where the image intensities do not match well (i.e., with larger D).

The first point in this list is the *raison d'être* of CF optimization, and what the registration algorithm is supposed to do. The second, however, creates a mathematical loophole in the optimization, allowing the algorithm to not only evade its main task (i.e., decreasing D), but possibly lead the deformation towards even worse alignment (by changing T in a local fashion such that D increases, but both W_T and the product $D \cdot W_T$ decrease). It is important to avoid regionally biasing the CF integral (unless intentionally done for a particular benefit) to prevent such negative consequences and make the algorithm focus solely on improving the alignment. Next, we show that standard deformable registration algorithms have CFs that are non-uniform (weighted) in at least one of the two native spaces, thus suffering the disadvantages described here.

2.3. Intrinsic non-uniformity in standard deformable image registration

As discussed in Section 2.2, the mismatch measure D should be integrated uniformly in order to avoid registration inaccuracy. However, there are *two* native spaces – that of f and that of g – which are equally physically meaningful. The f -native CF (Eq. (1)) satisfies uniformity in the native space of f ; yet, rewriting it in the native space of g via the change of variable $y := T(x)$ leads to the following non-uniform integral:

$$I_f(f, g, T) = \int_{\Omega} D(f \circ T^{-1}(y), g(y)) / J \circ T^{-1}(y) dy, \quad (3)$$

where the function $J: \Omega \mapsto \mathbb{R}_{> 0}$ is the Jacobian determinant of the transformation, $J(x) := \det \partial T(x)$, which we assume to be positive (i.e., T preserves the topology). Given that J represents the relative volume change caused by the transformation and is in general inhomogeneous, we can see that the f -native CF is weighted by $1/J$ and does not satisfy uniformity in our second native space, i.e. that of g . The uniform CF in the native space of g , or the *g-native CF*, has the following form:

$$I_g(f, g, T) := \int_{\Omega} D(f \circ T^{-1}(y), g(y)) dy. \quad (4)$$

The change of variable $y = T(x)$ again reveals that the g -native CF does not have a uniform representation in the native space of f :

$$I_g(f, g, T) = \int_{\Omega} D(f(x), g \circ T(x)) J(x) dx. \quad (5)$$

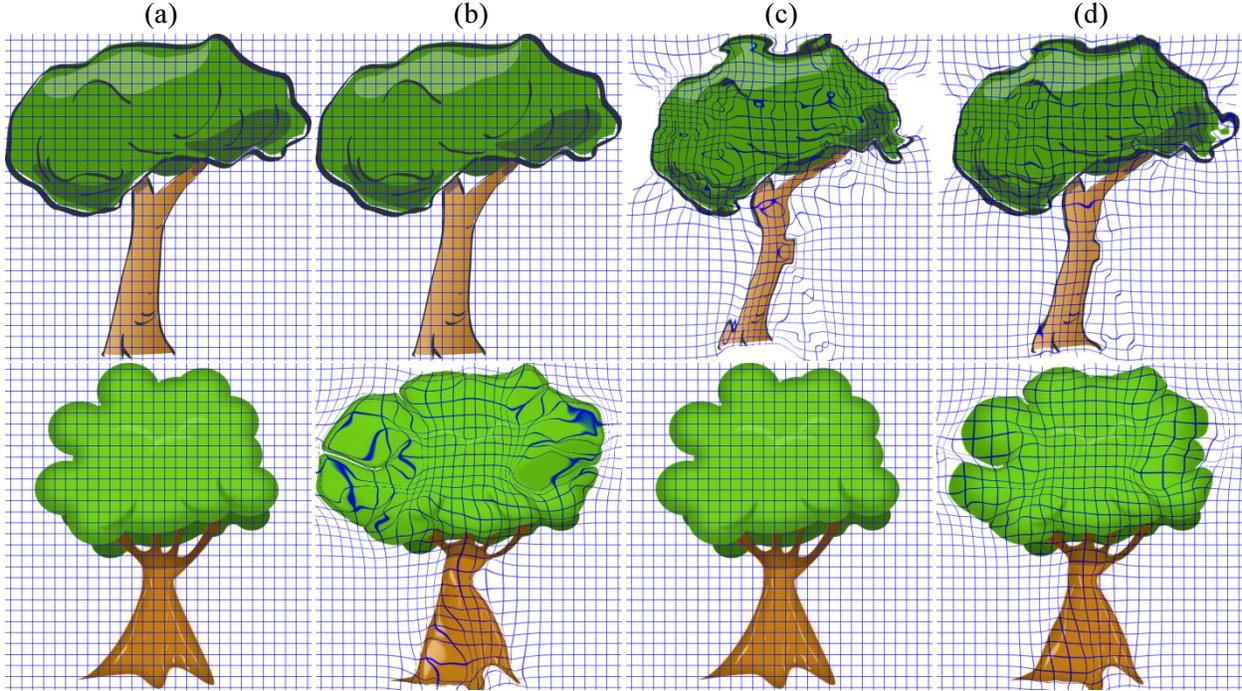


Figure 2. a) Two images f (top) and g (bottom) are given as input to the registration algorithm. The deformed images in an intermediate registration iteration are illustrated in: b) native space of f , c) native space of g , and d) a mid-space. The grid distortion in (b,c,d) demonstrates that no matter what space is chosen to compute the CF integral, at least one of the images is integrated non-uniformly.

It can be seen that, regardless of which native space is picked as the reference space, the undesirable weighting of the integral shows up in at least one of the native spaces (Figure 2a-c), except for the especially restrictive cases where $J(x) = 1$ is imposed, such as in rigid registration, or more generally volume-preserving registration (Elen et al., 2008; Haber and Modersitzki, 2007; Hirota et al., 2000; Li et al., 2009; Loeckx et al., 2005; Mansi et al., 2011; Park et al., 2011; Rohlfing et al., 2003; Saddi et al., 2007; Tanner et al., 2002; Yin et al., 2009). Volume-preserving approaches, however, make the assumption that the deformation causes no local shrinkage or expansion, which may generally not be the case.

The lack of simultaneous uniformity of the CFs in both native spaces, as mathematically revealed above, is an inherent imperfection of standard deformable registration. Specifically, since the weighting $J(x)$ is the local volume change, the algorithm may attempt to enforce artificial expansion or shrinkage – depending on which image is chosen as the reference – to decrease the weighted CF without accurate alignment (see Section 2.2). It is important to stress that this is not an esoteric argument that rarely occurs in practice; instead, it is inevitable in any nonlinear registration in any region in which there is an initial mismatch (see Figure 1 for an example). In these regions a biased CF integral will lead the algorithm to attempt to reduce $J(x)$ (along with D) as much as possible within the limits of the imposed regularization. In fact, standard registration has been shown to privilege either expanding or shrinking regions depending on the choice of the reference (Cachier and Rey, 2000; Rey et al., 2002). The added global regularization restricts this artificial volume change, but only to some extent; therefore we need to further prevent this weight factor from interfering in the data-attachment term’s driving of the registration towards the desired optimal alignment.

An implication of the non-uniform weighting in exactly one of the native spaces is that in general $I_f \neq I_g$. Also, since:

$$I_g(f, g, T) = I_f(g, f, T^{-1}), \quad (6)$$

one can see that the native CFs are inverse-inconsistent, i.e. $I_f(f, g, T) \neq I_f(g, f, T^{-1})$, and $I_g(f, g, T) \neq I_g(g, f, T^{-1})$. The primary approach in the literature to address this inverse-inconsistency is the minimization of the following *average CF* (Cachier and Rey, 2000; Christensen and Johnson, 2001; Tagare et al., 2009; Trouvé and Younes, 2000)⁶:

$$I_{ave} := \frac{I_f(f, g, T) + I_g(f, g, T)}{2}. \quad (7)$$

Substituting for I_g from Eq. (6), one can see that I_{ave} is a symmetrized version of either of the native CFs and is therefore inverse-consistent, i.e. $I_{ave}(f, g, T) = I_{ave}(g, f, T^{-1})$. However, by rewriting it in the native spaces of the two images, once using Eqs. (1,5) and again using Eqs. (3,4),

$$\begin{aligned} I_{ave}(f, g, T) &= \int_{\Omega} D(f(x), g \circ T(x)) \frac{1 + J(x)}{2} dx \\ &= \int_{\Omega} D(f \circ T^{-1}(y), g(y)) \frac{1 + 1/J \circ T^{-1}(y)}{2} dy, \end{aligned} \quad (8)$$

it turns out that I_{ave} is a weighted integral in both native spaces (Figure 2d), and so the registration still incurs the regional bias errors (Section 2.2), in spite of its inverse-consistency.

In an alternative approach to restore symmetry, the mismatch error is integrated uniformly in an abstract *mid-space*, chosen to be equally close to both native spaces (Avants et al., 2008; Beg and Khan, 2007; Lorenzen et al., 2004; Noblet et al., 2008). Provided that the mid-space is chosen “in between” the two native spaces, changing the ordering of the input images will not affect the CF, hence the inverse-consistency. Nevertheless, it can be easily shown similarly as above that rewriting such a mid-space CF in the native spaces of the images, once again, results in the Jacobian determinant of the transformation (that takes the mid-space to the native space) to appear as a weighting in the native-space representations of the CF. Therefore, the integral of such a mid-space CF is non-uniform in both of the physically meaningful native spaces, and only uniform in the virtual mid-space, which is physically nonexistent and possibly changing during the registration process. As a consequence, the problems associated with the regional bias (Section 2.2) remain unresolved in these approaches. For instance, since the mid-space and the weighting often depend on T , the optimization algorithm has the ability to alter the mid-space so as to decrease the CF without necessarily making the two images more similar. An example of this behavior is when the algorithm changes T in such a way that regions with mismatching image intensities are shrunk in the mid-space. This results in a decrease in the (mid-space) CF but not necessarily in the native CFs, since the images will look more similar in the mid-space but not necessarily similar in the two native spaces. Constraints that are used to keep the mid-space “in between” the native spaces of the two images usually

⁶ More references are provided in Section 1.

prevent such regions from totally shrinking, but depending on the strength of these constraints, partial shrinkage may still happen.

It should be noted that minimizing I_{ave} or a mid-space CF results in transformations that are not necessarily minimizers of the native CFs, and for which, I_f and I_g are not necessarily equal. Here, we define the *native symmetry* property of an image registration algorithm, as the quality of producing transformations minimizing the native CFs, while keeping them equal ($I_f = I_g$). Rigid, and more generally any volume-preserving registration is natively symmetric (disregarding the asymmetry due to resampling artifacts (Aganj et al., 2013b; Reuter et al., 2010; Sabuncu et al., 2009)), since with $J(x) = 1$ everywhere, the two native CFs (Eqs. (1,5)) are always equal. Native symmetry is a stronger form of symmetry, indicating that both native CFs agree with the progress of registration. Neither asymmetric (minimizing either I_f or I_g) nor symmetrized (minimizing I_{ave} or a mid-space CF) deformable registration is guaranteed to be natively symmetric.

3. Proposed methods

3.1. Quasi-volume-preserving constraint

In this work we introduce a constraint on the deformation field that keeps the non-uniformity error in the CF arbitrarily small. Examining Eqs. (3,5,8), we can see that the condition $J(x) = 1, \forall x \in \Omega$ would be sufficient for I_f, I_g , and I_{ave} to be free of non-uniformity error. Nevertheless, such a restrictive volume-preserving constraint is not a necessary condition, as J can indeed be different from 1 where the mismatch measure D is zero, and no error will yet arise from non-uniformity. Accordingly, by equating the integrand to its unweighted version, i.e. $D(f(x), g \circ T(x))J(x) = D(f(x), g \circ T(x))$, the following less restrictive constraint (here written in the native space of f) can also be seen to guarantee the elimination of the non-uniformity error:

$$D(f(x), g \circ T(x))(J(x) - 1) = 0 \quad , \quad \forall x \in \Omega. \quad (9)$$

Nonetheless, in practice, exact intensity matching, $D(f(x), g \circ T(x)) = 0$, and absolute volume preservation, $J(x) = 1$, are seldom achieved. Thus, to allow for large deformations, particularly in the presence of noise and poor intensity match, here we propose the following relaxation of the above constraint, which we call the *quasi-volume-preserving* (QVP) constraint:

$$D(f(x), g \circ T(x))|J(x) - 1| < \epsilon \quad , \quad \forall x \in \Omega, \quad (10)$$

where the constant ϵ is a positive threshold on the local non-uniformity error. In fact, the prefix *quasi-* indicates the fact that volume preservation is imposed only in regions with disagreeing image intensities, as well as the relaxation of Eq. (9) to Ineq. (10). Roughly speaking, for a small ϵ , the QVP constraint ensures that there is almost no volume change ($J(x) \cong 1$), except for regions where the image intensities almost match ($D(f(x), g \circ T(x)) \cong 0$).⁷ The performance of enforcing the QVP constraint with different values

⁷ Although the strict Eq. (9) is symmetric, its relaxation leading to Ineq. (10) introduces a level of asymmetry in the optimization space, which grows with ϵ . Symmetrization of (10) would require the additional constraint $D(f(x), g \circ T(x))|J^{-1}(x) - 1| < \epsilon, \forall x \in \Omega$, which, however, is overly restrictive, especially that even in diffeomorphic registration diffeomorphism is sometimes violated leading to $J(x) \leq 0$ for some voxels. Therefore, we only enforce Ineq. (10) in our experiments, as it fulfills the task of decreasing the non-uniformity error.

of ϵ has been evaluated in 3D brain image registration in Section 4.2. As we shall see, smaller values of ϵ , which further restrict the deformation, actually lead to improvements in the overall registration quality, indicating that the QVP constraint guides the registration away from problematic regions in the energy landscape, yielding better CF optima. A higher value for ϵ , however, might be needed to accommodate larger deformations and regions with poorer intensity match and higher noise (e.g., in longitudinal brain imaging when a tumor both grows *and* changes intensity).

By constraining the deviation of the volume change $J(x)$ from 1 as in the QVP inequality (10), the total non-uniformity error in the native CFs will be bounded by $\epsilon|\Omega|$ (with $|\Omega|$ the size of the space), which can be made arbitrarily small by choosing a small enough ϵ .⁸ Remaining in this (almost) error-free QVP zone ensures that the optimization algorithm is not misled by a biased CF gradient that can potentially steer it away from the global minimum, for example, by encouraging complex warps in regions where no matching of image intensities is possible. In addition, by combining Eqs. (1,5,10), the error in native symmetry can similarly be seen to be bounded:

$$\begin{aligned} |I_f - I_g| &= \left| \int_{\Omega} D(f(x), g \circ T(x))(J(x) - 1) dx \right| \\ &\leq \int_{\Omega} D(f(x), g \circ T(x)) |J(x) - 1| dx \\ &< \epsilon |\Omega|, \end{aligned} \tag{11}$$

with I_{ave} being close to both native CFs (using Eqs. (7,11)):

$$|I_{ave} - I_{f/g}| < \frac{\epsilon |\Omega|}{2}. \tag{12}$$

The local nature of the QVP constraint also ensures regional native symmetry; i.e., satisfying Ineq. (10) keeps native CFs defined in any local $\Omega' \subset \Omega$ within $\epsilon|\Omega'|$ of each other.⁹

By keeping the integrals of the CFs (almost) uniform in the native spaces of *both* images except for regions where non-uniformity barely perturbs the CF, the QVP constraint ensures that CF minimization is uniquely aimed at better alignment, and that no regional bias irrelevant to registration can be exploited to minimize the CF (Section 2.2).

3.2. Adaptive regularization interpretation

Enforcing the QVP constraint restricts $J(x)$ to remain close to 1 in regions where the difference between the two images is large, but relaxes this restriction where the image intensities are similar. An additional practical advantage of the proposed method is its role as adaptive spatially-variant regularization (Cahill et al., 2009; Freiman et al., 2012; Glocker et al., 2009; Loeckx et al., 2004; Ou et al., 2011; Pace et al., 2013;

⁸ Using the change of variable $y = T(x)$, this bound can be seen to exist on the non-uniformity error of both f - and g -native CFs.

⁹ Note that the closeness of the native CFs (native symmetry) is a nice byproduct, but not the goal of QVP constraining. Otherwise, simpler approaches could be used to keep I_f and I_g close to each other without necessarily achieving the goal of QVP, which is limiting the local non-uniformity error. In fact, since the QVP constraint is enforced everywhere in the image, it leads to one constraint per voxel after discretization, which may be several orders of magnitude more constraints than merely one bound on the native symmetry error.

Risholm et al., 2009; Staring et al., 2007; Stefanescu et al., 2004; Suárez et al., 2002; Tang et al., 2010). By restricting the deformation in regions where there is not a good match and letting the surrounding area to mostly drive the displacement, the QVP constraint allows the deformed image to “pass through” regions of poor match with restricted volume change, to potentially find regions where a good registration is possible, after which the constraint is naturally relaxed. Intuitively, the algorithm avoids fitting noise in regions where a good overall match cannot be found. Examples of such behavior are provided in sections 4.1.1 and 4.1.2.

The factor D in Ineq. (10) is the major contrast between our proposed approach and the volume-preserving methods in the literature (Section 2.3); our algorithm preserves volume *only* in regions where it finds no good match between the images. This means that the transformation may still contain compression and expansion in areas where the two images are locally similar, allowing good matching in regions where it is feasible. Although the QVP constraint is automatically relaxed in such similar-intensity regions, the inclusion of the global regularization still keeps the deformation field regular.¹⁰

3.3. Incorporation of the QVP constraint

The QVP constraint, which can be seen as a restriction on the feasible region of the optimization search space, can be applied to any deformable registration tool that optimizes an aggregate objective function, i.e. sum of a local similarity measure (such as SSD). To enforce the QVP inequality (10), at each iteration of the main registration algorithm we reduce the non-uniformity error:

$$e(x) := D(f(x), g \circ T(x)) |J(x) - 1|, \quad (13)$$

until $e(x) < \epsilon$ for all $x \in \Omega$. To that end, $e(x)$ may be directly reduced by iteratively minimizing the integral $\int_{\Omega} R(e(x) - \epsilon) dx$ to zero, with $R(t) := (t + |t|)/2$ the ramp function, thereby satisfying $e(x) < \epsilon$ for all x . The dependence of $e(x)$ on both $T(x)$ and its derivative $J(x)$, however, makes this direct minimization complicated.¹¹ Thus, for simplicity and numerical stability, we choose a different approach to take $T(x)$ to the QVP space.

To estimate the derivative of $e(x)$ with respect to $T(x)$ and moving $T(x)$ in its descent direction, we first note that $e(x)$ is the product of two factors: the mismatch measure $D(f(x), g \circ T(x))$ and the Jacobian deviation $|J(x) - 1|$. According to the product rule, $d(uv) = vdu + udv$, the derivative of $e(x)$ has two terms, in each of which only one of the two factors is derived. The first term has the derivative of the

¹⁰ Note that the primary motivation for the proposed work, however, is limiting the non-uniformity error and achieving more accurate registration. Therefore the QVP constraint is not a deformation model and the adaptive regularization is only a secondary motivation.

¹¹ For instance for the SSD CF, where $D(u, v) := (u - v)^2$, one can verify the following complex gradient direction for $T(x)$, which is the variation of $\int_{\Omega} R(e(x) - \epsilon) dx$ with respect to $T(x)$:

$$2(f(x) - g \circ T(x)) [(\epsilon \delta(e(x) - \epsilon) + u(e(x) - \epsilon)) (J(x) (\nabla g) \circ T(x) - C(x) \nabla f(x)) \text{sign}(J(x) - 1) - u(e(x) - \epsilon) |J(x) - 1| (\nabla g) \circ T(x))] - \delta(e(x) - \epsilon) (f(x) - g \circ T(x))^4 C(x) \sum_k H^k(x) \left(C_{k,:}(x) \right)^T,$$

where $u(\cdot)$ and $\delta(\cdot)$ are the Heaviside step and the Dirac delta functions, respectively, and the cofactor matrix $C(x)$ and the Hessian tensor $H(x)$ are defined in Appendix A and Appendix B, respectively.

mismatch measure $D(f(x), g \circ T(x))$, and is a gradient term similar to that of the main optimization algorithm, simply leading to additional CF minimization. Since the mismatch measure is already minimized by the main optimization algorithm, we do not include this image-aligning term in the QVP constraining step,¹² and rather focus on the second term of the derivative.

To reduce the Jacobian deviation factor $|J(x) - 1|$ of $e(x)$, whose derivative appears in the second term of the derivative of $e(x)$, we take advantage of the fact that the descent direction for the following measure of regularity (where $K(x)$ is a weighting function),

$$\int_{\Omega} K(x) \|\partial T(x) - \mathbb{I}\|_F^2 dx, \quad (14)$$

is the well-known inhomogeneous diffusion, applied to the displacement field,

$$T_{t+1}(x) = T_t(x) + \gamma \nabla \cdot [K(x)(\nabla T_t(x) - \vec{1})], \quad (15)$$

where t , γ , $K(x)$, and $\vec{1}$ are the diffusion iteration number, the step size, the space-variant diffusion coefficient, and the vector of all 1s, respectively. Therefore, the above diffusion process reduces the deviation of the Jacobian ∂T from the identity matrix (the higher $K(x)$, the more strongly). We have shown in Appendix C that such small updates, while bringing the Jacobian ∂T closer to identity, also make the Jacobian determinant J asymptotically closer to 1. Since we want the diffusion to be much stronger in the areas with higher error, we make the diffusion coefficient grow quadratically with $e(x)$ as follows:

$$K(x) := \alpha e^2(x). \quad (16)$$

It is important for $K(x)$ to be a smooth function of $e(x)$, as its derivative is computed in the diffusion process. We perform the inhomogeneous diffusion of Eq. (15) after both the main gradient-descent/Newton iteration and the standard demons regularization have been performed. To speed up the diffusion process, we apply a number of (10 in our experiments) iterations of diffusion, then update $e(x)$ and $K(x)$ with the new $T(x)$, and if the QVP constraint $e(x) < \epsilon, \forall x \in \Omega$ is not yet satisfied we repeat with further diffusion.¹³ The stopping criterion (satisfaction of QVP) makes this operation idempotent.

As we adopt the demons deformable registration (Thirion, 1998) in this work, the inhomogeneous blurring of the displacement field (primarily in areas with higher $e(x)$) at each registration iteration can be seen to be related to the Gaussian blurring (equivalent to homogeneous diffusion, see Appendix B) in the demons method, which blurs the displacement field at each iteration. However, a major contrast here is that the (space-variant) blurring depends on the mismatch measure $D(f(x), g \circ T(x))$ that includes image data,

¹² As we can see in the results (Section 4), although we have no such additional CF minimization, we obtain a lower CF value after the same number of gradient descent or Newton iterations (as well as better label matching and lower distortion). Devising a more direct strategy to enforce the QVP constraint is a subject of the ongoing research.

¹³ A drawback of this implementation is that we do not know beforehand how many blurring iterations are required, thus we cannot precisely predict the runtime of the algorithm. The runtime increases as ϵ decreases; however, it also depends on the input images and other parameters, such as the optimization and diffusion step sizes. A specific parameterization of the deformation field intrinsically satisfying the QVP constraint is a subject for future work. The operator introduced here is also not symmetric, as we intend to minimize the non-uniformity error, rather than to explicitly enforce symmetry. Note that the QVP constraint can be incorporated in non-rigid registration algorithms in various ways, the implementation of which is not the primary focus of this paper.

since both (i) the spatial weighting of the diffusion $K(x)$, and (ii) the important stopping criterion for the diffusion, i.e. the QVP inequality (10), depend on $D(f(x), g \circ T(x))$.

We chose $\alpha = 70$ in our experiments. Furthermore, to keep the diffusion stable, we choose the diffusion step size γ to be at most $2^{-(d+1)}$ (Johnson et al., 2014), and threshold $K(x)$ so it remains below a certain value (2 in our experiments). Also to make sure that local noise does not result in an overly restrictive constraint, we denoise $K(x)$ by a Gaussian low-pass filter with the standard deviation of 1 voxel, before using it in the diffusion step.

4. Results and discussion

To show that the proposed QVP constraint can be incorporated into and improve different implementations of deformable registration techniques, we validate it on both demons (in Matlab) and diffeomorphic demons (in Insight Toolkit (ITK)) registration. In all of the experiments, the image intensity values were initially scaled to be between 0 and 1.

4.1. Demons registration of 2D images

We first validate our method using 2D (non-diffeomorphic) demons registration by comparing the performance of the following three approaches: minimization of I_f , minimization of I_{ave} , and minimization of I_{ave} with the proposed QVP constraint (with the heuristically-determined optimum ϵ). Since with no explicit diffeomorphism constraints the topology is not necessarily preserved, for each method and experiment, we choose the optimum general demons regularization parameter λ (see Appendix B for definition) heuristically to achieve the best convergence with no observable topology break. For each experiment, we plot I_f and I_g with respect to the iteration number to assess the registration native symmetry by the extent to which these two native CFs agree with each other. We evaluate this throughout all the iterations, as opposed to only at the convergence, since this error is an aggregate measure of non-uniformity, which in turn can cause inaccurate gradient descent directions at every iteration and lead the optimization algorithm to undesirable solutions.¹⁴ In these non-diffeomorphic experiments, we also evaluate the overall results by the number of iterations that each algorithm takes to converge, and the best achievable ultimate CF value that preserves the topology, which is an indicator of the *CF-distortion tradeoff*. The experiments in this subsection are performed using our *unidirectional* Matlab implementation of CF minimization (Appendix A) and global regularization (Appendix B), meaning that only the forward CF needs to be explicitly evaluated.

¹⁴ Since the QVP is incorporated in an already inverse-consistent method (minimization of I_{ave}), we do not expect to see any further improvement in the inverse-consistency (contrary to native symmetry) here. Note that the QVP constraint is not employed for the purpose of inverse-consistency, but to improve the registration through minimization of the non-uniformity errors.

4.1.1. Synthetic data

We first compare the three methods on synthetic data, by registering two letters ‘B’ with noticeably different shapes. As the plot in Figure 3 demonstrates, minimizing I_f (dark and light green) results in a relatively large difference $I_g - I_f > 0$, which is expected, since I_g is not in any form accounted for in the optimization. Minimizing I_{ave} (dark and light red) reduces the gap between I_f and I_g , as equal weight is given to both of them in the optimization. Nevertheless, this approach – not being natively symmetric – does not guarantee the equality of the two native CFs. The proposed QVP approach with $\epsilon = .2$ (dark and light blue), however, enforces a bound on $|I_f - I_g|$ and therefore, as illustrated in the plot, produces transformations on the cost of which both native CFs agree (within the error margin ϵ). In addition, this method converges at an earlier gradient descent iteration (\sim Itr. 750, see the orange rectangles in Figure 3 (top)) compared to the two others (\sim Itr. 1800), while resulting in lower final native CF values than the standard asymmetric and the symmetrized approaches do.¹⁵

The better convergence of the QVP method may be related to the fact that its optimum general regularization parameter $\lambda = .2$ is naturally lower compared to the other two techniques ($\lambda = .3$ and $\lambda = .7$ for the standard and the symmetrization methods), given that it also performs a separate inhomogeneous regularization (Eq. (15)).

¹⁵ It should be noted that the number of iterations is not indicative of the computational cost in our implementation, since each iteration is more expensive in the QVP method than in the other two methods, given the extra step of diffusion. As mentioned before, the final CF values are the lowest achieved that preserve the topology, thereby indicating the CF-distortion tradeoff. The fact that the optimum λ was different for each method makes direct comparison of these convergence rates difficult. However, one can see that the best achievable convergence over a range of λ values was improved by enforcing the QVP constraint.

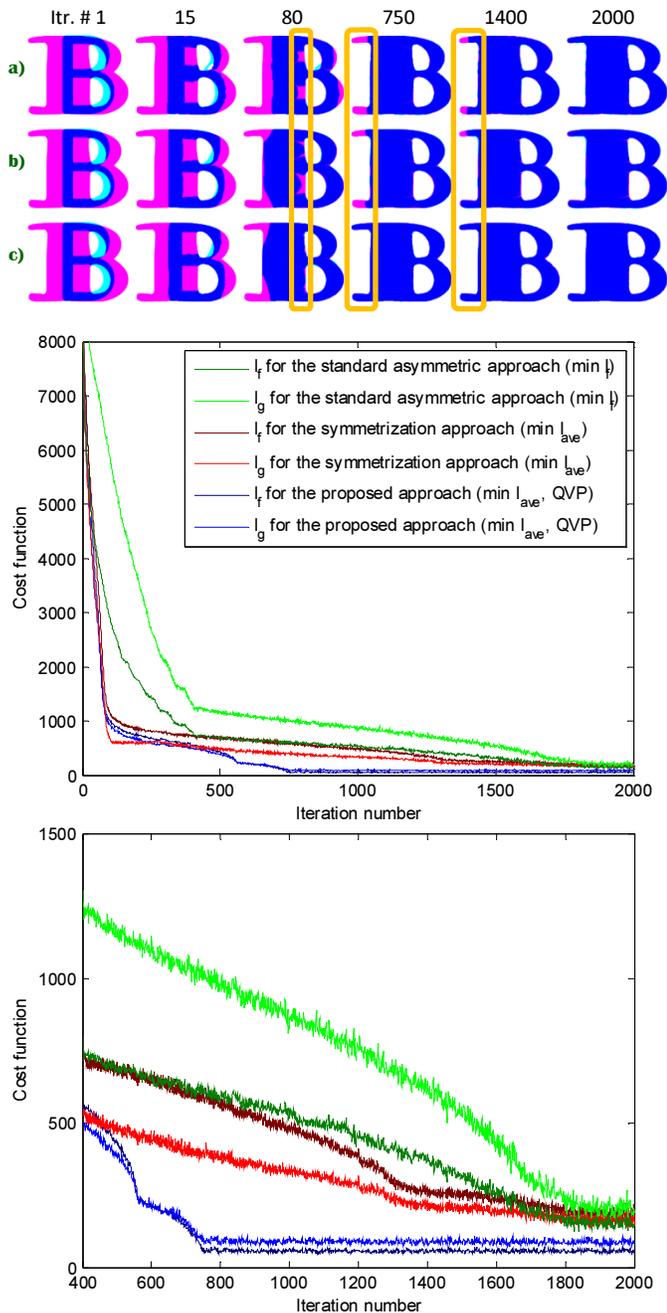


Figure 3. Registration of synthetic data by a) minimizing I_f , b) minimizing I_{ave} , and c) minimizing I_{ave} with the proposed QVP constraint. Top: Cyan and magenta indicate respectively the reference and the interpolated images, with blue being their intersection. Middle/Bottom: The native CFs for each method, and a zoomed version.

This method allows the transformation to freely absorb and compress one ‘B’ into the other, since the corresponding volume change happens in regions where the intensities match, for which the inhomogeneous regularization, and consequently the total regularization is weak. The inhomogeneous regularization, in contrast, prevents a topology break by regularizing more aggressively in areas where the difference between the two images (and therefore the SSD force) is large. Conversely, the other two methods need a higher λ to avoid topology break, and yet apply this strong general regularization to the entire image (including regions with matching intensities and no gradient force), hence a slower convergence and higher final CF (worse CF-distortion tradeoff).

4.1.2. X-ray images

Next, we compared the three methods on a pair of public jaw X-ray images acquired before (Figure 4a) and after (Figure 4b) an orthognathic surgery. Prior to deformable registration, high-resolution versions of the two images were superimposed, the upper jaws were rigidly aligned manually, and the images were low-pass-filtered and resampled to the lower resolution that is shown in Figure 4. As in the previous case, minimization of the standard asymmetric forward CF, I_f (Figure 4, the two green curves), soon creates a considerable gap

between the native CFs. Minimizing I_{ave} (Figure 4, the two red curves) results in a steady optimization of the *average* of the native CFs. However, since the native CFs are not individually considered, not only does the gap between them still mostly exist, but a jump that abruptly changes them – yet not their average – occurs around iteration 2000, which is possibly a sign of the potential susceptibility of the algorithm to local optima. In fact, an optimization algorithm might decide that it has converged and stop the procedure prematurely before iteration 2000, thus leaving the teeth region misaligned (as was actually the case here; the misalignment was observed in the deformed image for iterations before 2000). This is another example (similar to Figure 1) showing a transformation that decreases the SSD in one native space but increases it in the other native space, despite the fact that a symmetrized CF has been used. Conversely, the proposed

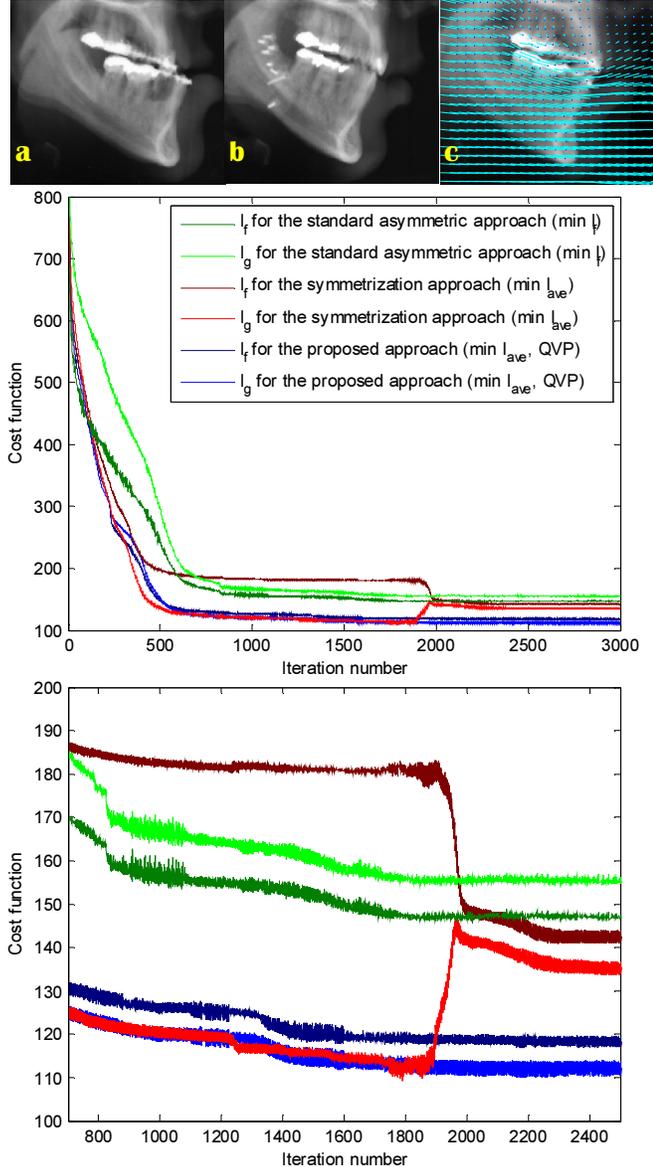


Figure 4. Top: X-ray images taken a) before and b) after surgery, with c) registration results using the proposed QVP approach. Middle/Bottom: Native CFs for the three methods, and a zoomed version.

QVP method with $\epsilon = .2$ (Figure 4, the two blue curves) produces the smallest native CF gap for iterations before 2000, while obtaining the lowest final CF values (best CF-distortion tradeoff) compared to the standard asymmetric and the symmetrized approaches. The warped image obtained by the QVP approach is depicted in Figure 4c, along with the computed deformation field. Results from the two other methods were visually similar, except for the fact that the displacement in the upper jaw, which is supposed to be small because of the initial rigid alignment of the upper jaws, was lowest in the deformation obtained by the QVP approach (2.5 pixels, averaged in an upper jaw mask), and therefore least affected by the general regularization, compared to the standard (3.9 pixels) and the symmetrization (3.2 pixels) approaches. This is likely again due to the lower value of λ required by QVP, as explained in Section 4.1.1.

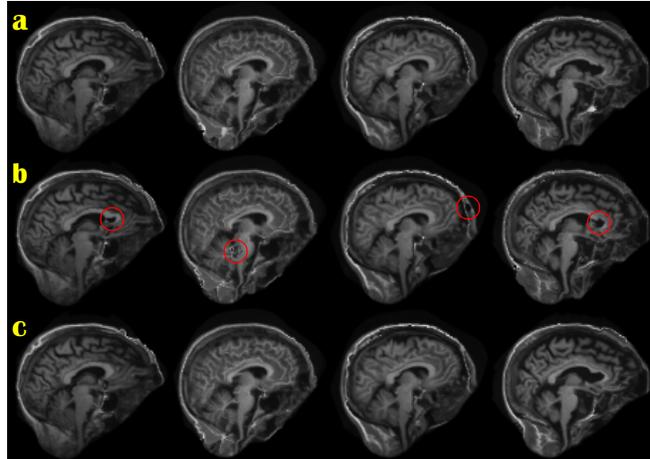


Figure 5. a) QVP registration results on four subjects with optimal values of λ . b) Results of I_{ave} minimization using the same λ as in (a). Note the distortion due to under-regularization in the encircled areas. c) Results of I_{ave} minimization using optimal (higher) values of λ , thus with higher CF values.

4.1.3. Brain images

We also tested the three algorithms on the mid-sagittal planes of 20 brain images taken from the publicly available OASIS database (Marcus et al., 2007), which we pre-processed in FreeSurfer (Dale et al., 1999; Fischl, 2012; Fischl et al., 1999). The intensity-normalized and resampled volumes (1-mm³ isotropic voxel size) were made upright by robust rigid registration (Reuter et al., 2010) of each volume to its left-right mirrored version. The sagittal slice located four voxels to the right of the mid-sagittal plane was extracted from each volume, and to adjust for any nodding rotation, was rigidly registered to that slice of the first volume, and resampled to the size 128×128. Out of the 20 sagittal slices corresponding to the 20 subjects, the one closest (in L_2 norm) to the rest was chosen as the reference, and non-rigidly registered to the other 19 subjects individually. For each subject and method, we ran the registration with 21 different values for general regularization parameter λ . We observed throughout our experiments that a lower value of λ that still preserved the topology consistently resulted in a lower final CF. Therefore, in order for each algorithm to present its best CF-distortion tradeoff, we inspected the results (19×3×21=1197 images), and for each subject and method we chose the one obtained with the lowest λ containing no visible topology break. The optimally chosen λ was the same for the symmetrization and the QVP ($\epsilon = .17$) methods in 13 subjects; however, for the other 6 subjects, QVP passed the visual inspection test at a lower λ (see Figure 5 for examples). We computed the native asymmetry error as the mean absolute value of the difference of the two native CFs through all iterations. When averaged across subjects, this error was 161% and 16% higher for the asymmetric and the symmetrization techniques, respectively, compared to the proposed QVP approach. We also compared the final values of the native CFs, (I_f, I_g) , among the methods, which were (30%, 56%) and (5%, 4%) higher for the asymmetric and the symmetrization approaches, respectively, compared to the QVP method. We hypothesized that the QVP approach results in a better CF-distortion

tradeoff by producing ultimate native CF values (for the lowest λ preserving the topology) that are lower than those of the two other techniques. A left-tailed Student's t -test with a .05 significance level rejected the null hypothesis when comparing the QVP with both the asymmetric ($p = 10^{-6}$) and the symmetrization ($p = .02$) approaches.

4.2. Diffeomorphic demons registration of 3D brain images

Next, we examined the effects of QVP constraining on the accuracy of registration using 3D labeled data, by incorporating the QVP constraint into the diffeomorphic demons registration module (Vercauteren et al., 2008a) of ITK. We randomly selected 50 pairs of images, with each image chosen from a set of 40 3D brain volumes that had 37 neuroanatomical structures manually labeled (Fischl et al., 2002). These pairs were registered with and without the QVP constraint for various values of λ (here the global regularization parameter of diffeomorphic demons) for 50 Newton iterations (enough for convergence). When QVP was used, four values of $\epsilon = .05, .1, .2, .3$ were tried (with $\epsilon = .05$ tried only for three values of λ because of its high computational cost).¹⁶

As illustrated in Figure 6, the inclusion of the QVP constraint consistently improved the results in every way that we quantified them. 1) QVP resulted in better alignment of manual labelings, i.e. lower number of differently-classified corresponding voxels (Figure 6, top). We then considered the CF-distortion tradeoff in the results. 2) QVP achieved a lower final CF value, despite the fact that it further regularizes the deformation field (Figure 6, 2nd row). This provides evidence that the QVP constraint leads the algorithm into a better optimum. 3) QVP resulted in transformations with lower distortion in terms of the elastic energy, computed following (Hagemann et al., 1999; Yanovsky et al., 2008a), and the number of voxels with Jacobian determinant values smaller than 10^{-3} or larger than 10^3 (Figure 6, two bottom rows). This improvement is expected, as QVP increases the regularization.

¹⁶ The standard symmetric-gradient implementation of ITK is not equivalent to minimizing I_{ave} . Given that this subsection focuses on registration accuracy rather than symmetry, to keep the results fair and reproducible, we only included the QVP constraint and did not modify the ITK implementation any further. Contrary to our experiments in Section 4.1, where we showed improvement by QVP compared to symmetrized registration, in Section 4.2 we assess the inclusion of QVP in the ITK implementation that is not symmetric. Therefore, the improvements observed in this subsection may be as a result of the symmetry as well.

We also attempted to incorporate the QVP constraint into the more recent symmetric log-domain diffeomorphic demons registration (Vercauteren et al, 2008b). However, since this method parameterizes the transformation implicitly as a velocity field, and its public implementation does not allow for explicit modification of the deformation field, we could not apply the inhomogeneous diffusion of Section 3.3 directly to the deformation field. We tried applying the diffusion to the velocity field instead, but it performed poorly, since the QVP constraint is naturally very local, and any local modification of the velocity field propagates into a larger area in the deformation field. A specifically log-domain enforcement of the QVP constraint is a subject of the future research.

Please note that a fair and accurate assessment of QVP constraining using any other registration methods would require explicit implementation of QVP in their source codes with their specific deformation parameterization; otherwise, directly comparing our diffeomorphic demons implementation of QVP with other (non-demons, non-QVP) methods would cause the difference between the registration methods to dominate the QVP effect.

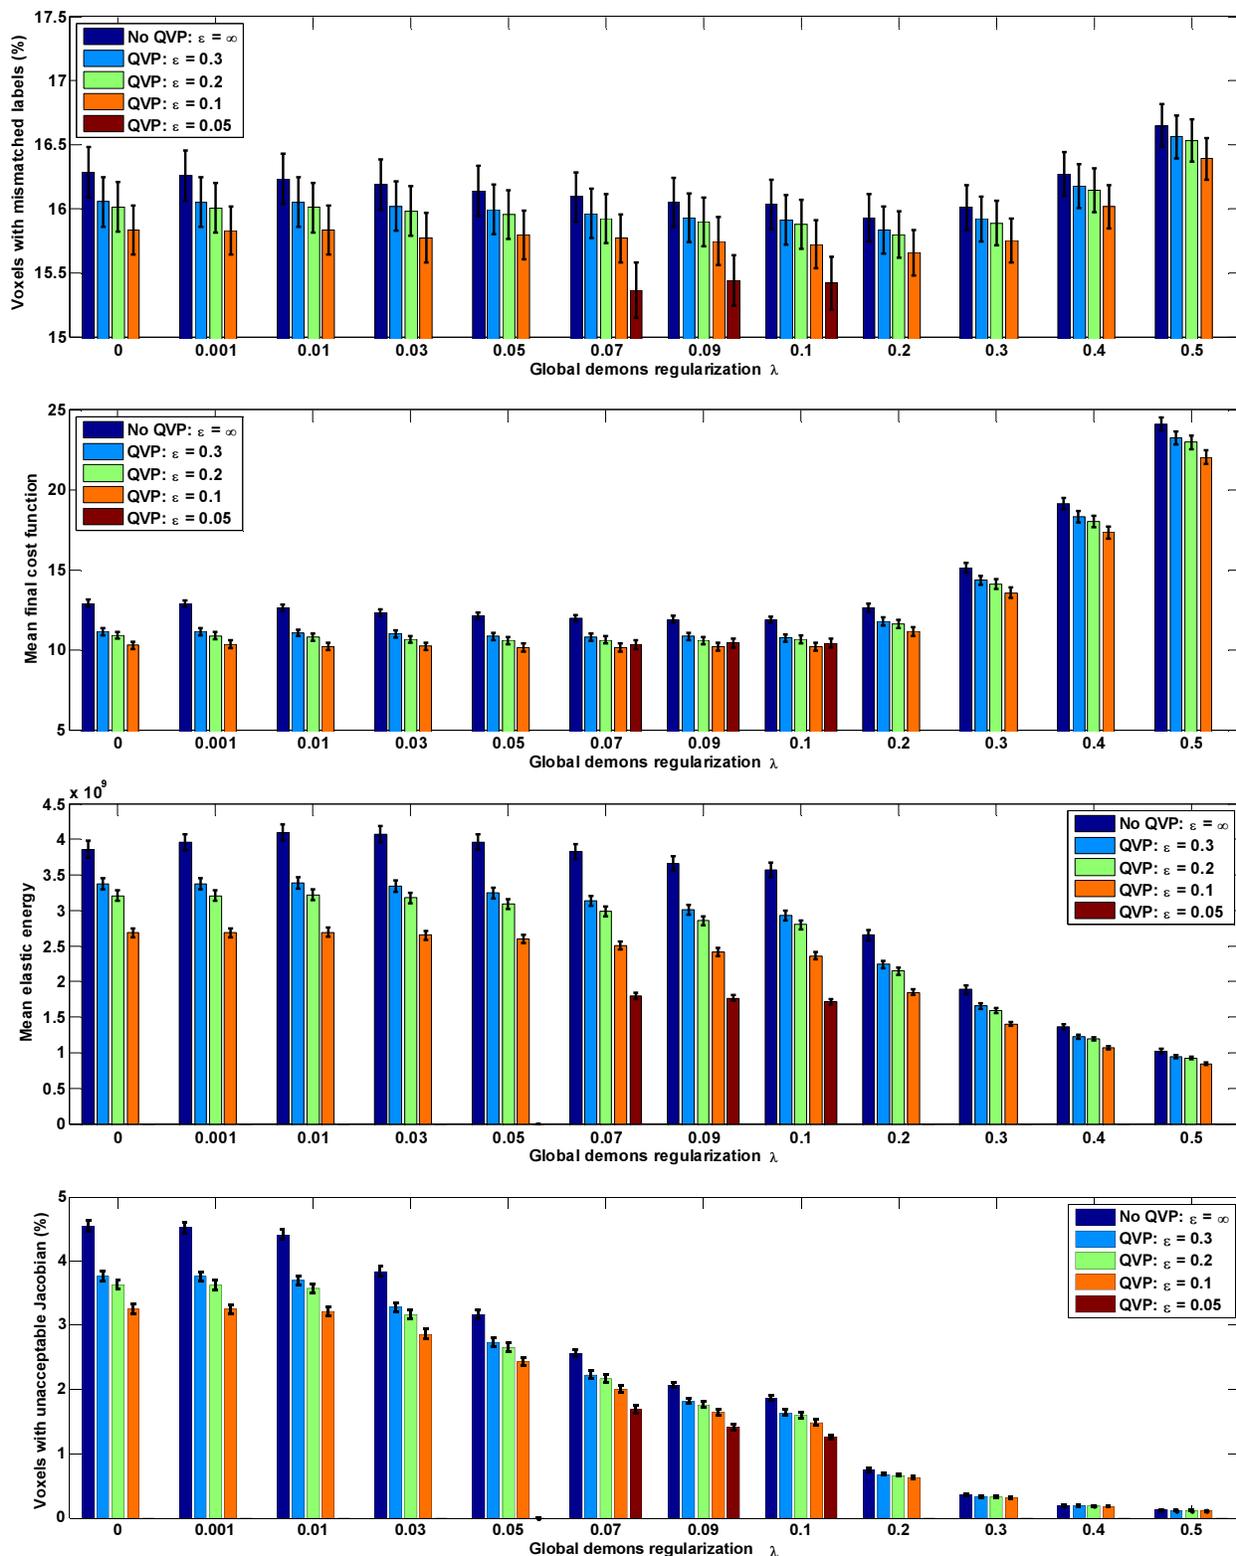


Figure 6. Adding the QVP constraint to the diffeomorphic demons algorithm results in improved alignment of manual labelings (top), better intensity matching (2nd row), and less distortion (two bottom rows). Quantities are averaged across 50 experiments, and the error bars indicate the standard error of the mean. This is the optimum range for λ ; we tried higher values of it, which degraded the performance of both methods, with QVP still outperforming. 19

Figure 7 depicts the CF-distortion tradeoff for the experiments.¹⁷ Indeed, a tighter QVP constraint (smaller ϵ) always improved both the label matching and the CF-distortion tradeoff. (This is true for $\epsilon = .05$ as well, even though it did not reduce the final CF as much as $\epsilon = .1$.) Disabling the global regularization ($\lambda = 0$) did not produce the lowest CF, indicating that too much flexibility increases the chances of being trapped in a local optimum. We stress that while it is impossible to exactly match the amount of regularization in the QVP and non-QVP case, in all of the aforementioned registration experiments with a wide range of regularization levels the incorporation of the QVP constraint yielded better accuracy and image matching, and lower levels of distortion.

In practice, there may be a limit to the improvement that can be achieved by lowering ϵ . In particular, the number of necessary diffusion iterations (for each Newton iteration) goes to infinity as ϵ approaches zero, and for the extreme choice of $\epsilon = 0$ the QVP constraint is never satisfied. In our experiments, an order of 1, 10, and 100 rounds of diffusion were needed for $\epsilon = .3, .2, \epsilon = .1$, and $\epsilon = .05$, respectively, to satisfy QVP for all the voxels in the image.

Since the non-uniformity error is the culprit in breaking the symmetry of registration, it is expected that limiting this error by QVP constraining would make the registration more inverse-consistent. For the representative regularization value of $\lambda = .1$, we repeated the experiments above by swapping the order of input images and registering them again once without QVP, and a second time with QVP ($\epsilon = .3$). We then computed the inverse-consistency error as $\int_{\Omega} \|T_1 \circ T_2(x) - x\|_2^2 dx / |\Omega|$, where T_1 and T_2 are the two transformations achieved by forward and backward registrations. As demonstrated by the histogram in Figure 8, in all of the experiments, enforcing the QVP constraint reduced the inverse-consistency error (mean \pm SD: $5.2\% \pm 3.3\%$). A paired left-sided Wilcoxon signed rank test revealed the significance value of $p = 3 \times 10^{-8}$.

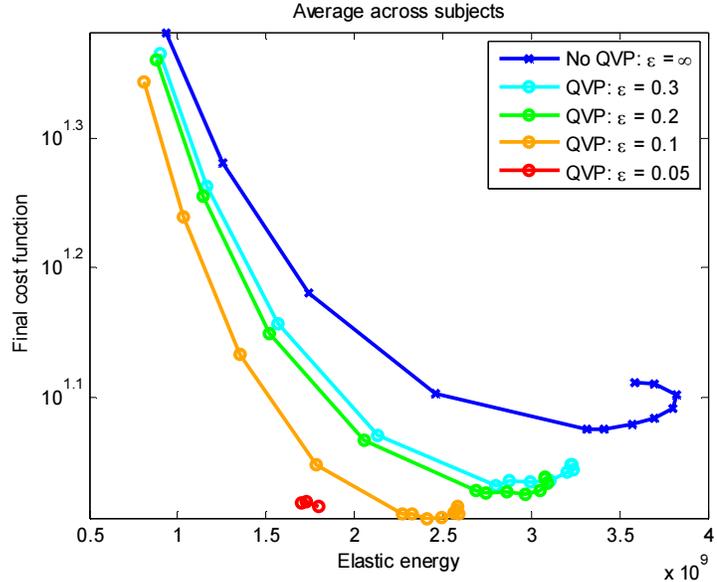


Figure 7. CF-distortion tradeoff in diffeomorphic registration experiments. Each curve corresponds to a unique value of ϵ , and each point on the curve represents a different value of λ . For each point on the curve, the CF and distortion values have been averaged across subjects.

¹⁷ Although individual comparison of the final CF value or the distortion among methods might not be meaningful, here we compare the CF-distortion tradeoff among methods, which can be used as a criterion to assess registration algorithms.

5. Conclusions

We have proposed a quasi-volume-preserving (QVP) constraint on the deformation field in deformable registration that keeps the optimization algorithm away from the regions violating the uniformity of the cost function integral and consequently the symmetry. The main focus of this work is addressing the inaccuracies arising from the non-uniformity of the cost functions, which is the major distinction between our method and the symmetrization and mid-space techniques. In that regard, we have shown theoretically and through experiments that the QVP constraint improves, not only native symmetry, but the overall registration results. Enforcing the QVP constraint led to cross-subject registration with increased accuracy in the alignment of

manually defined neuroanatomical structures, better tradeoff between the intensity error and distortion in the warp field, improved native symmetry, and reduced susceptibility to local optima. Subjects of the future research include: intrinsic enforcement of QVP via specific parameterization of the deformation field (without the need for blurring the displacement field), optimal ϵ determination based on image properties, proving the convergence of and further stabilizing the diffusion step, and assessing the incorporation of QVP in the mid-space approaches.

Acknowledgements

Support for this research was provided in part by the National Center for Research Resources (U24 RR021382), the National Institute for Biomedical Imaging and Bioengineering (NIBIB, P41-EB015896, R01EB006758), the National Institute on Aging (AG022381, 5R01AG008122-22, R01 AG016495-11), the National Center for Alternative Medicine (RC1 AT005728-01), the National Institute for Neurological Disorders and Stroke (R01 NS052585-01, 1R21NS072652-01, 1R01NS070963, R01NS083534), and was made possible by the resources provided by Shared Instrumentation Grants 1S10RR023401, 1S10RR019307, and 1S10RR023043. Additional support was provided by The Autism & Dyslexia Project funded by the Ellison Medical Foundation, and by the NIH Blueprint for Neuroscience Research (5U01-MH093765), part of the multi-institutional Human Connectome Project. IA was additionally supported by a grant from the Massachusetts Alzheimer’s Disease Research Center (5 P50 AG005134), the MGH Neurology Clinical Trials Unit, and the Harvard NeuroDiscovery Center. MS received support from an NIBIB K25 grant (1K25EB013649-01) and a BrightFocus Alzheimer’s disease pilot grant (A2012333). In addition, BF has a financial interest in CorticoMetrics, a company whose medical pursuits focus on brain imaging and measurement technologies. BF’s interests were reviewed and are managed by Massachusetts General Hospital and Partners HealthCare in accordance with their conflict of interest policies.

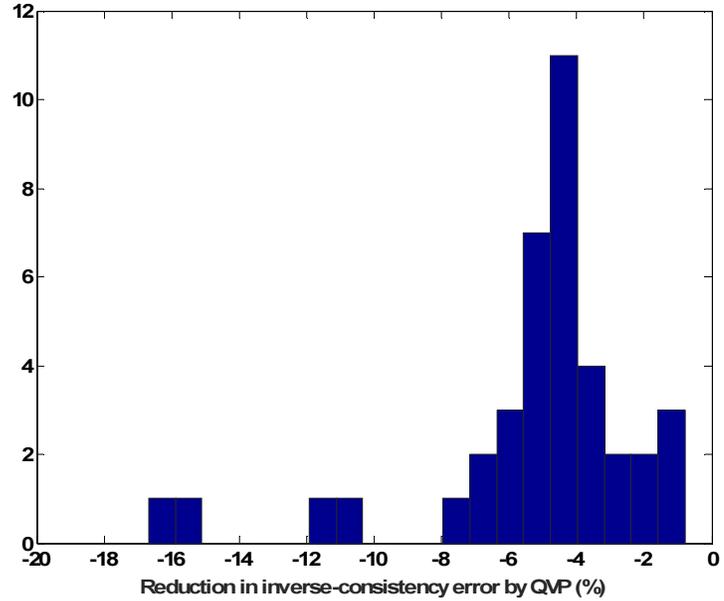


Figure 8. The histogram of the decrease in inverse-consistency error by including the QVP constraint ($\epsilon = .3$) confirms that reducing the non-uniformity error results in improvement in inverse-consistency.

Appendix A. Unidirectional minimization of the average cost function

We minimize the average CF (Eq. (7)) in the experiments of Section 4.1 via gradient descent in the same unidirectional manner as in (Leow et al., 2007), with differences that are outlined here. Basically, we calculate the variation of the average CF with respect to the transformation, $\delta I_{ave}/\delta T(x)$, and update $T(x)$ towards its opposite (descent) direction.

I_{ave} is the sum of two terms, with the first term, $I_f/2$, being written in terms of $T(x)$ as in Eq. (1). The variation of I_f with respect to $T(x)$ is straightforward to compute:

$$\frac{\delta I_f}{\delta T(x)} = D_2(f(x), g \circ T(x))(\nabla g) \circ T(x), \quad (17)$$

where D_k is the partial derivative of D with respect to the k^{th} argument, and ∇ is the image gradient operator. The second term of I_{ave} is $I_g/2$, which is written in terms of $T^{-1}(y)$ as in Eq. (4) and, yet, is not an explicit function of $T(x)$.¹⁸ Here, instead of direct inversion of T , we compute the variation of I_g with respect to $T(x)$ indirectly. First we note that the variation of I_g with respect to $T^{-1}(y)$ can be computed similarly to Eq. (17), as:

$$\frac{\delta I_g}{\delta T^{-1}(y)} = D_1(f \circ T^{-1}(y), g(y))(\nabla f) \circ T^{-1}(y). \quad (18)$$

Let us assume that a small variation $\Delta T(x)$ changes the transformation as $T^*(x) = T(x) + \Delta T(x)$. This change results in a small variation in $T^{-1}(y)$, as $T^{-1*}(y) = T^{-1}(y) + \Delta T^{-1}(y)$. Keeping the first-order terms, the relationship $T^{-1*} \circ T^*(x) = x$ leads to:

$$\Delta T^{-1}(T(x)) = -\partial T(x)^{-1} \Delta T(x), \quad (19)$$

where $\partial T(x)$ is the Jacobian matrix of $T(x)$ (Leow et al., 2007). Furthermore, by definition, I_g will vary as follows:

$$\Delta I_g = \int_{\Omega} \left[\frac{\delta I_g}{\delta T^{-1}(y)} \right]^T \Delta T^{-1}(y) dy. \quad (20)$$

Making the change of variable $y = T(x)$ and substituting for $\Delta T^{-1}(T(x))$ from Eq. (19) gives:

$$\Delta I_g = \int_{\Omega} \left(-J(x) \partial T(x)^{-1T} \left[\frac{\delta I_g}{\delta T^{-1}(y)} \right]_{y=T(x)} \right)^T \Delta T(x) dx, \quad (21)$$

with $J(x) := \det \partial T(x)$. The variation of I_g with respect to $T(x)$ is therefore derived as:

$$\frac{\delta I_g}{\delta T(x)} = -C(x) \left[\frac{\delta I_g}{\delta T^{-1}(y)} \right]_{y=T(x)}, \quad (22)$$

¹⁸ Although I_g can be written with respect to $T(x)$ as in Eq. (5), the inclusion of $J(x)$ in the integral makes the computation of the variation complicated.

where $C(x) = J(x)\partial T(x)^{-1T}$ is the cofactor matrix of $\partial T(x)$, defined as $C_{ij} = (-1)^{i+j}M_{ij}$ with M the minor of ∂T (i.e., M_{ij} is the determinant of the matrix that results from deleting row i and column j of ∂T). Substituting from Eq. (18),

$$\frac{\delta I_g}{\delta T(x)} = -D_1(f(x), g \circ T(x))C(x)\nabla f(x), \quad (23)$$

and combining Eqs. (7,17,23), the variation of I_{ave} with respect to $T(x)$ is derived as follows:

$$\frac{\delta I_{ave}}{\delta T(x)} = \frac{D_2(f(x), g \circ T(x))(\nabla g) \circ T(x) - D_1(f(x), g \circ T(x))C(x)\nabla f(x)}{2}. \quad (24)$$

For instance, in case of the SSD CF, i.e. $D(u, v) := (u - v)^2$, Eq. (24) reduces to:

$$\frac{\delta I_{ave}}{\delta T(x)} = -(f(x) - g \circ T(x))(C(x)\nabla f(x) + (\nabla g) \circ T(x)). \quad (25)$$

Note that the descent direction for I_{ave} was here computed by deriving both its terms $-I_f$ and I_g – consistently with respect to $T(x)$. On the contrary, in (Leow et al., 2007), the descent directions are computed separately for I_f and I_g with respect to $T(x)$ and $T^{-1}(y)$, respectively, with the latter direction being subsequently translated in terms of change in $T(x)$.¹⁹

Next, we explain how to incorporate the demons regularization symmetrically.

Appendix B. Unidirectional symmetric demons regularization

In addition to QVP, we used a global regularization in the experiments, which is an essential part of the demons registration (Thirion, 1998). In our 2D implementation in Section 4.1, we included a general Tikhonov term $\frac{1}{2}\|\partial T(x) - \mathbb{I}\|_F^2$ (weighted by λ , with \mathbb{I} the identity matrix) inside the integral of the f -native CF (Eq. (1)), whereas in Section 4.2 we used ITK’s built-in regularization. The variation of $R_f := \int_{\Omega} \frac{1}{2}\|\partial T(x) - \mathbb{I}\|_F^2 dx$ with respect to $T(x)$ can be derived as:

$$\frac{\delta R_f}{\delta T(x)} = -\nabla_x^2 T(x), \quad (26)$$

where ∇_x^2 is the Laplacian with respect to x performed on each element of T . Updating the deformation field in the opposite direction of this variation is equivalent to solving the diffusion equation, which in turn results in a Gaussian blurring in the displacement field, i.e. what is done directly in the original demons algorithm.²⁰ The average CF in Eq. (7), however, includes the g -native CF (Eq. (4)), which also needs such a term as $R_g := \int_{\Omega} \frac{1}{2}\|\partial T^{-1}(y) - \mathbb{I}\|_F^2 dy$ for the symmetry to be preserved. To minimize this term

¹⁹ Replacing $C(x)$ with $\partial T(x)$ in Eqs. (24,25) leads to the results in (Leow et al., 2007). In the case of 1D images, for instance, $C_{1D}(x) = 1$ making Eq. (25) similar to the demons registration with symmetric gradients, whereas $\partial T_{1D}(x) = T_{1D}'(x)$.

²⁰ Using the displacement $T(x) - x$ instead of the deformation $T(x)$ field will not alter the regularization; the term x , being linear, is eliminated by the Laplacian, and being anti-symmetric, is cancelled out by the symmetric Gaussian blurring.

unidirectionally, we need to compute its variation with respect to $T(x)$. As in Appendix A, we first compute its variation with respect to $T^{-1}(y)$,

$$\frac{\delta R_g}{\delta T^{-1}(y)} = -\nabla_y^2 T^{-1}(y), \quad (27)$$

and then derive its variation with respect to $T(x)$ using the change of variable $y = T(x)$ and Eq. (22), as:

$$\begin{aligned} \frac{\delta R_g}{\delta T(x)} &= -C(x) \left[\frac{\delta R_g}{\delta T^{-1}(y)} \right]_{y=T(x)} \\ &= C(x) [\nabla_y^2 x]_{y=T(x)}. \end{aligned} \quad (28)$$

So the problem boils down to the computation of $\nabla_y^2 x = \nabla_y \cdot \nabla_y x$, represented in x . The gradient of x with respect to y is simply the inverse Jacobian matrix:

$$\nabla_y x = \partial T(x)^{-1}, \quad (29)$$

which we shall denote by $Q := \partial T^{-1}$.²¹ We now compute the divergence of the above:

$$\nabla_y \cdot \nabla_y x = \sum_k \frac{\partial Q_{:,k}}{\partial y_k}, \quad (30)$$

where the subscript $(:, k)$ means the k^{th} column of the matrix. The chain rule implies:

$$\begin{aligned} \frac{\partial}{\partial y_k} &= \sum_l \frac{\partial x_l}{\partial y_k} \frac{\partial}{\partial x_l} \\ &= \sum_l Q_{l,k} \frac{\partial}{\partial x_l}, \end{aligned} \quad (31)$$

and combining Eqs. (30,31) leads to:

$$\begin{aligned} \nabla_y^2 x &= \sum_{k,l} \frac{\partial Q_{:,k}}{\partial x_l} Q_{l,k} \\ &= \sum_l \left[\frac{\partial Q}{\partial x_l} Q^T \right]_{:,l}. \end{aligned} \quad (32)$$

We now compute $\frac{\partial Q}{\partial x_l}$ by taking the derivative of the relationship $Q\partial T = \mathbb{I}$, as follows:

$$\frac{\partial}{\partial x_l} (Q\partial T) = 0, \quad (33)$$

²¹ For brevity, we omit (x) in most of the rest of this appendix.

$$\frac{\partial Q}{\partial x_l} \partial T + Q \frac{\partial}{\partial x_l} \partial T = 0, \quad (34)$$

$$\begin{aligned} \frac{\partial Q}{\partial x_l} &= -Q \frac{\partial}{\partial x_l} \partial T \partial T^{-1} \\ &= -Q \frac{\partial}{\partial x_l} \partial T Q. \end{aligned} \quad (35)$$

Equation (32) thus becomes:

$$\begin{aligned} \nabla_y^2 x &= - \sum_l \left(Q \frac{\partial}{\partial x_l} \partial T Q Q^T \right)_{:,l} \\ &= -Qh, \end{aligned} \quad (36)$$

with the vector h defined as:

$$h := \sum_l \left(\frac{\partial}{\partial x_l} \partial T Q Q^T \right)_{:,l}. \quad (37)$$

The n^{th} element of h can be calculated as follows:

$$\begin{aligned} h_n &= \sum_l \left(\frac{\partial}{\partial x_l} \partial T Q Q^T \right)_{n,l} \\ &= \sum_l \sum_k \left(\frac{\partial}{\partial x_l} \partial T \right)_{n,k} (Q Q^T)_{k,l} \\ &= \sum_{k,l} \frac{\partial^2 T_n}{\partial x_k \partial x_l} (Q Q^T)_{k,l} \\ &= \sum_{k,l} H_{k,l}^n (Q Q^T)_{k,l} \\ &= \text{tr}(H^n Q Q^T) \\ &= \text{tr}(Q^T H^n Q), \end{aligned} \quad (38)$$

where T_n is the n^{th} element of T , and H^n is the Hessian matrix of T_n . By combining Eqs. (28,36), the variation of R_g with respect to $T(x)$ is obtained as:

$$\frac{\delta R_g}{\delta T(x)} = -C(x)Q(x)h(x). \quad (39)$$

Therefore, we first compute the elements of the vector $h(x)$, followed by the variation of R_g with respect to $T(x)$, which allows for $T^{-1}(y)$ to be regularized without being explicitly calculated.

An issue that we faced in the computation of Eq. (39) is the calculation of the inverse Jacobian matrix Q where diffeomorphism is not preserved. As a remedy, we only performed the regularization of T^{-1} in regions where the Jacobian matrix had a positive determinant and singular values greater than a threshold (.25 in our experiments).

Appendix C. Deviation of the Jacobian determinant

Using the following Theorem, we can see that for a diffeomorphic d -dimensional transformation, $J(x) - 1$ is bounded from above and below by two functions of $\|\partial T(x) - \mathbb{I}\|_F$, as follows:

$$\max(1 - \|\partial T(x) - \mathbb{I}\|_F, 0)^d - 1 \leq J(x) - 1 \leq \left(1 + \frac{1}{\sqrt{d}} \|\partial T(x) - \mathbb{I}\|_F\right)^d - 1. \quad (40)$$

Both of these bounds approach 0 monotonically as $\|\partial T(x) - \mathbb{I}\|_F \rightarrow 0$, thereby taking $J(x) - 1$ to 0 following the squeeze theorem.²²

Theorem. For a $d \times d$ matrix M with $\det M \geq 0$, we have the following bounds for $\det M$:

$$\max(1 - \|M - \mathbb{I}\|_F, 0)^d \leq \det M \leq \left(1 + \frac{1}{\sqrt{d}} \|M - \mathbb{I}\|_F\right)^d, \quad (41)$$

where \mathbb{I} is the $d \times d$ identity matrix.

Proof. We have $\det M \leq |\det M| = \prod_{k=1}^d \sigma_k$, where σ_k is the k^{th} singular value of M . According to the triangular inequality for the Frobenius norm,

$$\|M - \mathbb{I}\|_F \geq \|M\|_F - \|\mathbb{I}\|_F = \sqrt{\sum_{k=1}^d \sigma_k^2} - \sqrt{d}. \quad (42)$$

Combining this with the *root mean square-geometric mean* inequality, $\sqrt{\frac{1}{d} \sum_{k=1}^d \sigma_k^2} \geq (\prod_{k=1}^d \sigma_k)^{\frac{1}{d}}$, yields:

$$\det M \leq \left(1 + \frac{1}{\sqrt{d}} \|M - \mathbb{I}\|_F\right)^d, \quad (43)$$

which proves the right-hand side of Ineq. (41). Note that this upper bound is a supremum, since for a given $\|M - \mathbb{I}\|_F$, equality can be achieved, e.g., for $M = \left(1 + \frac{1}{\sqrt{d}} \|M - \mathbb{I}\|_F\right) \mathbb{I}$. We did not need $\det M \geq 0$ to prove this upper bound.

Regarding the lower bound, one can see that for $\|M - \mathbb{I}\|_F \geq 1$ the lower limit $\det M = 0$ can be reached, for instance for a diagonal M whose first diagonal entry is 0, and all the rest of its $d - 1$ diagonal entries are equal to $1 + \sqrt{\frac{1}{d-1} (\|M - \mathbb{I}\|_F^2 - 1)}$. Concerning the case with $\|M - \mathbb{I}\|_F < 1$, we use the inequality $\sigma_{\min}(A - B) \leq \sigma_{\min}(A) + \sigma_{\max}(B)$ (Horn and Johnson, 1994), with $\sigma_{\min}(A)$ and $\sigma_{\max}(A)$ the minimum and maximum singular value of A , respectively. Letting $A = M$ and $B = M - \mathbb{I}$ gives:

$$1 = \sigma_{\min}(\mathbb{I}) \leq \sigma_{\min}(M) + \sigma_{\max}(M - \mathbb{I}). \quad (44)$$

Since $\sigma_{\max}(M - \mathbb{I}) \leq \|M - \mathbb{I}\|_F$, we have:

²² This approach may reduce also the divergence of the displacement field.

$$\sigma_{\min}(M) \geq 1 - \|M - \mathbb{I}\|_F. \quad (45)$$

Considering that $\det M = |\det M| = \prod_{k=1}^d \sigma_k(M) \geq \sigma_{\min}(M)^d$ and the assumption $1 - \|M - \mathbb{I}\|_F > 0$, we raise the equation above to the power of d , resulting in:

$$\det M \geq (1 - \|M - \mathbb{I}\|_F)^d. \quad (46)$$

Finally, we rewrite the two cases of $\|M - \mathbb{I}\|_F \geq 1$ and $\|M - \mathbb{I}\|_F < 1$ in the following compact form, which proves the left-hand side of Ineq. (41):

$$\max(1 - \|M - \mathbb{I}\|_F, 0)^d \leq \det M. \quad (47)$$

□

References

- Aganj, I., Reuter, M., Sabuncu, M.R., Fischl, B., 2013a. Symmetric non-rigid image registration via an adaptive quasi-volume-preserving constraint. Proceedings of the IEEE International Symposium on Biomedical Imaging, San Francisco, CA, USA, pp. 234-237.
- Aganj, I., Yeo, B.T.T., Sabuncu, M.R., Fischl, B., 2013b. On Removing Interpolation and Resampling Artifacts in Rigid Image Registration. Image Processing, IEEE Transactions on 22, 816-827.
- Alvarez, L., Deriche, R., Papadopoulos, T., Sánchez, J., 2007. Symmetrical Dense Optical Flow Estimation with Occlusions Detection. International Journal of Computer Vision 75, 371-385.
- Ashburner, J., Andersson, J.L.R., Friston, K.J., 1999. High-Dimensional Image Registration Using Symmetric Priors. NeuroImage 9, 619-628.
- Ashburner, J., Andersson, J.L.R., Friston, K.J., 2000. Image registration using a symmetric prior—in three dimensions. Human brain mapping 9, 212-225.
- Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C., 2008. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. Medical image analysis 12, 26-41.
- Basri, R., Costa, L., Geiger, D., Jacobs, D., 1998. Determining the similarity of deformable shapes. Vision Research 38, 2365-2385.
- Beg, M.F., Khan, A., 2007. Symmetric data attachment terms for large deformation image registration. Medical Imaging, IEEE Transactions on 26, 1179-1189.
- Bondar, L., Hoogeman, M.S., Osorio, E.M.V., Heijmen, B.J.M., 2010. A symmetric nonrigid registration method to handle large organ deformations in cervical cancer patients. Medical Physics 37, 3760-3772.
- Cachier, P., Rey, D., 2000. Symmetrization of the non-rigid registration problem using inversion-invariant energies: Application to multiple sclerosis. In: Delp, S., DiGoia, A., Jaramaz, B. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2000. Springer Berlin / Heidelberg, pp. 697-708.
- Cahill, N., Noble, J.A., Hawkes, D., 2009. A Demons Algorithm for Image Registration with Locally Adaptive Regularization. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2009. Springer Berlin Heidelberg, pp. 574-581.
- Chen, Y., Ye, X., 2010. Inverse Consistent Deformable Image Registration. In: Alladi, K., Klaunder, J.R., Rao, C.R. (Eds.), The Legacy of Alladi Ramakrishnan in the Mathematical Sciences. Springer New York, pp. 419-440.
- Christensen, G., Johnson, H., 2003. Invertibility and transitivity analysis for nonrigid image registration. Journal of Electronic Imaging 12, 106-117.
- Christensen, G.E., Johnson, H.J., 2001. Consistent image registration. Medical Imaging, IEEE Transactions on 20, 568-582.
- Chui, H., 2001. Non-rigid point matching: Algorithms, extensions and applications. Yale University.
- Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical Surface-Based Analysis: I. Segmentation and Surface Reconstruction. NeuroImage 9, 179-194.
- Dedeoglu, G., Kanade, T., 2005. On the Source of Asymmetry in Image Registration Problems. Carnegie Mellon University.

Elen, A., Hon Fai, C., Loeckx, D., Hang, G., Claus, P., Suetens, P., Maes, F., D'Hooge, J., 2008. Three-Dimensional Cardiac Strain Estimation Using Spatio-Temporal Elastic Registration of Ultrasound Images: A Feasibility Study. *Medical Imaging, IEEE Transactions on* 27, 1580-1591.

Feng, W., Reeves, S.J., Denney, T.S., Lloyd, S., Dell'Italia, L., Gupta, H., 2009. A new consistent image registration formulation with a B-spline deformation model. *Biomedical Imaging: From Nano to Macro, 2009. ISBI '09. IEEE International Symposium on*, pp. 979-982.

Fischl, B., 2012. FreeSurfer. *NeuroImage* 62, 774-781.

Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der, K.A., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., Dale, A.M., 2002. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* 33, 341-355.

Fischl, B., Sereno, M.I., Dale, A.M., 1999. Cortical surface-based analysis II: Inflation, flattening, and a surface-based coordinate system. *NeuroImage* 9, 195-207.

Fox, N.C., Ridgway, G.R., Schott, J.M., 2011. Algorithms, atrophy and Alzheimer's disease: Cautionary tales for clinical trials. *NeuroImage* 57, 15-18.

Freiman, M., Voss, S., Warfield, S., 2012. Abdominal Images Non-rigid Registration Using Local-Affine Diffeomorphic Demons. In: Yoshida, H., Sakas, G., Linguraru, M. (Eds.), *Abdominal Imaging. Computational and Clinical Applications*. Springer Berlin Heidelberg, pp. 116-124.

Geng, X., 2007. Transitive inverse-consistent image registration and evaluation. University of Iowa.

Gholipour, A., Kehtarnavaz, N., Yousefi, S., Gopinath, K., Briggs, R., 2010. Symmetric deformable image registration via optimization of information theoretic measures. *Image and Vision Computing* 28, 965-975.

Glocker, B., Komodakis, N., Navab, N., Tziritas, G., Paragios, N., 2009. Dense Registration with Deformation Priors. In: Prince, J., Pham, D., Myers, K. (Eds.), *Information Processing in Medical Imaging*. Springer Berlin Heidelberg, pp. 540-551.

Haber, E., Modersitzki, J., 2007. Image registration with guaranteed displacement regularity. *International Journal of Computer Vision* 71, 361-372.

Hagemann, A., Rohr, K., Stiehl, H.S., Spetzger, U., Gilsbach, J.M., 1999. Biomechanical modeling of the human head for physically based, nonrigid image registration. *Medical Imaging, IEEE Transactions on* 18, 875-884.

He, J., Christensen, G., 2003. Large Deformation Inverse Consistent Elastic Image Registration. In: Taylor, C., Noble, J. (Eds.), *Information Processing in Medical Imaging*. Springer Berlin / Heidelberg, pp. 438-449.

Hirota, G., Maheshwari, R., Lin, M.C., 2000. Fast volume-preserving free-form deformation using multi-level optimization. *Computer-Aided Design* 32, 499-512.

Horn, R.A., Johnson, C.R., 1994. *Topics in Matrix Analysis*. Cambridge University Press.

Hua, X., Gutman, B., Boyle, C.P., Rajagopalan, P., Leow, A.D., Yanovsky, I., Kumar, A.R., Toga, A.W., Jack Jr, C.R., Schuff, N., Alexander, G.E., Chen, K., Reiman, E.M., Weiner, M.W., Thompson, P.M., 2011. Accurate measurement of brain changes in longitudinal MRI scans using tensor-based morphometry. *NeuroImage* 57, 5-14.

Johnson, H.J., McCormick, M.M., Ibanez, L., Consortium, I.S., 2014. *The ITK Software Guide, Book 1: Introduction and Development Guidelines*, 4 ed.

Joshi, S., Davis, B., Jomier, M., Gerig, G., 2004. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage* 23, Supplement 1, S151-S160.

Leow, A.D., Yanovsky, I., Ming-Chang, C., Lee, A.D., Klunder, A.D., Lu, A., Becker, J.T., Davis, S.W., Toga, A.W., Thompson, P.M., 2007. Statistical Properties of Jacobian Maps and the Realization of Unbiased Large-Deformation Nonlinear Image Registration. *Medical Imaging, IEEE Transactions on* 26, 822-832.

Li, X., Dawant, B.M., Welch, E.B., Chakravarthy, A.B., Freehardt, D., Mayer, I., Kelley, M., Meszoely, I., Gore, J.C., Yankeelov, T.E., 2009. A nonrigid registration algorithm for longitudinal breast MR images and the analysis of breast tumor response. *Magnetic Resonance Imaging* 27, 1258-1270.

Loeckx, D., Drisis, S., Maes, F., Vandermeulen, D., Marchal, G., Suetens, P., 2005. Plaque and Stent Artifact Reduction in Subtraction CT Angiography Using Nonrigid Registration and a Volume Penalty. In: Duncan, J., Gerig, G. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2005*. Springer Berlin Heidelberg, pp. 361-368.

Loeckx, D., Maes, F., Vandermeulen, D., Suetens, P., 2004. Nonrigid Image Registration Using Free-Form Deformations with a Local Rigidity Constraint. In: Barillot, C., Haynor, D., Hellier, P. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2004*. Springer Berlin Heidelberg, pp. 639-646.

Lorenzen, P., Davis, B., Joshi, S., 2004. Model based symmetric information theoretic large deformation multi-modal image registration. *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*, pp. 720-723 Vol. 721.

Lorenzi, M., Ayache, N., Frisoni, G.B., Pennec, X., 2013. LCC-Demons: A robust and accurate symmetric diffeomorphic registration algorithm. *NeuroImage* 81, 470-483.

Mansi, T., Pennec, X., Sermesant, M., Delingette, H., Ayache, N., 2011. iLogDemons: A Demons-Based Registration Algorithm for Tracking Incompressible Elastic Biological Tissues. *International Journal of Computer Vision* 92, 92-111.

Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L., 2007. Open Access Series of Imaging Studies (OASIS): Cross-sectional MRI data in young, middle aged, nondemented, and demented older adults. *Journal of Cognitive Neuroscience* 19, 1498-1507.

Modat, M., Cardoso, M., Daga, P., Cash, D., Fox, N., Ourselin, S., 2012. Inverse-Consistent Symmetric Free Form Deformation. In: Dawant, B., Christensen, G., Fitzpatrick, J., Rueckert, D. (Eds.), *Biomedical Image Registration*. Springer Berlin / Heidelberg, pp. 79-88.

Mohagheghian, F., Ahmadian, A., Saberi, H., Alirezaie, J., 2010. Symmetric multi-scale image registration. *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*, pp. 5931-5934.

Noblet, V., Heinrich, C., Heitz, F., Armspach, J.-P., 2008. Symmetric nonrigid image registration: Application to average brain templates construction. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2008*. Springer Berlin / Heidelberg, pp. 897-904.

Ou, Y., Sotiras, A., Paragios, N., Davatzikos, C., 2011. DRAMMS: Deformable registration via attribute matching and mutual-saliency weighting. *Medical image analysis* 15, 622-639.

Pace, D.F., Aylward, S.R., Niethammer, M., 2013. A Locally Adaptive Regularization Based on Anisotropic Diffusion for Deformable Image Registration of Sliding Organs. *Medical Imaging, IEEE Transactions on* 32, 2114-2126.

- Park, S., Kim, B., Lee, J., Goo, J.M., Shin, Y.-G., 2011. GGO Nodule Volume-Preserving Nonrigid Lung Registration Using GLCM Texture Analysis. *Biomedical Engineering, IEEE Transactions on* 58, 2885-2894.
- Reuter, M., Rosas, H.D., Fischl, B., 2010. Highly accurate inverse consistent registration: A robust approach. *NeuroImage* 53, 1181-1196.
- Reuter, M., Schmansky, N.J., Rosas, H.D., Fischl, B., 2012. Within-subject template estimation for unbiased longitudinal image analysis. *NeuroImage* 61, 1402-1418.
- Rey, D., Subsol, G., Delingette, H., Ayache, N., 2002. Automatic detection and segmentation of evolving processes in 3D medical images: Application to multiple sclerosis. *Medical image analysis* 6, 163-179.
- Risholm, P., Samset, E., Talos, I.-F., Wells, W., 2009. A Non-rigid Registration Framework That Accommodates Resection and Retraction. In: Prince, J., Pham, D., Myers, K. (Eds.), *Information Processing in Medical Imaging*. Springer Berlin Heidelberg, pp. 447-458.
- Rogelj, P., Kovačič, S., 2006. Symmetric image registration. *Medical image analysis* 10, 484-493.
- Rohlfing, T., Maurer, C.R., Jr., Bluemke, D.A., Jacobs, M.A., 2003. Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint. *Medical Imaging, IEEE Transactions on* 22, 730-741.
- Sabuncu, M., Yeo, B., Van Leemput, K., Vercauteren, T., Golland, P., 2009. Asymmetric image-template registration. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2009*. Springer Berlin / Heidelberg, pp. 565-573.
- Saddi, K.A., Chafd'hotel, C., Cheriet, F., 2007. Large deformation registration of contrast-enhanced images with volume-preserving constraint. 651203-651203.
- Škrinjar, O., Bistoquet, A., Tagare, H., 2008. Symmetric and Transitive Registration of Image Sequences. *International Journal of Biomedical Imaging*.
- Sotiras, A., Davatzikos, C., Paragios, N., 2013. Deformable Medical Image Registration: A Survey. *Medical Imaging, IEEE Transactions on* 32, 1153-1190.
- Staring, M., Klein, S., Pluim, J.P.W., 2007. Nonrigid registration with tissue-dependent filtering of the deformation field. *Physics in medicine and biology* 52, 6879.
- Stefanescu, R., Pennec, X., Ayache, N., 2004. Grid powered nonlinear image registration with locally adaptive regularization. *Medical image analysis* 8, 325-342.
- Suárez, E., Westin, C.-F., Rovaris, E., Ruiz-Alzola, J., 2002. Nonrigid Registration Using Regularized Matching Weighted by Local Structure. In: Dohi, T., Kikinis, R. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2002*. Springer Berlin Heidelberg, pp. 581-589.
- Tagare, H., Groisser, D., Skrinjar, O., 2009. Symmetric non-rigid registration: A geometric theory and some numerical techniques. *Journal of Mathematical Imaging and Vision* 34, 61-88.
- Tang, L., Hamarneh, G., Abugharbieh, R., 2010. Reliability-Driven, Spatially-Adaptive Regularization for Deformable Registration. In: Fischer, B., Dawant, B., Lorenz, C. (Eds.), *Biomedical Image Registration*. Springer Berlin Heidelberg, pp. 173-185.
- Tanner, C., Schnabel, J.A., Degenhard, A., Castellano-Smith, A.D., Hayes, C., Leach, M.O., Hose, D.R., Hill, D.L.G., Hawkes, D.J., 2002. Validation of Volume-Preserving Non-rigid Registration: Application to Contrast-Enhanced MR-Mammography. In: Dohi, T., Kikinis, R. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2002*. Springer Berlin Heidelberg, pp. 307-314.

Tao, G., He, R., Datta, S., Narayana, P.A., 2009. Symmetric inverse consistent nonlinear registration driven by mutual information. *Computer Methods and Programs in Biomedicine* 95, 105-115.

Thirion, J.P., 1998. Image matching as a diffusion process: An analogy with Maxwell's demons. *Medical image analysis* 2, 243-260.

Thompson, W.K., Holland, D., 2011. Bias in tensor based morphometry Stat-ROI measures may result in unrealistic power estimates. *NeuroImage* 57, 1-4.

Trouvé, A., Younes, L., 2000. Diffeomorphic matching problems in one dimension: Designing and minimizing matching functionals. *Computer Vision - ECCV 2000*. Springer Berlin / Heidelberg, pp. 573-587.

Vercauteren, T., Pennec, X., Perchant, A., 2008a. Ayache N. Diffeomorphic Demons Using ITK's Finite Difference Solver Hierarchy. *The Insight Journal*.

Vercauteren, T., Pennec, X., Perchant, A., Ayache, N., 2008b. Symmetric Log-Domain Diffeomorphic Registration: A Demons-Based Approach. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2008*. Springer Berlin Heidelberg, pp. 754-761.

Yang, D., Li, H., Low, D.A., Deasy, J.O., Naqa, I.E., 2008. A fast inverse consistent deformable image registration method based on symmetric optical flow computation. *Physics in medicine and biology* 53, 6143.

Yanovsky, I., Le Guyader, C., Leow, A.D., Thompson, P.M., Vese, L., 2008a. Nonlinear Elastic Registration with Unbiased Regularization in Three Dimensions. *MICCAI Workshop on Computational Biomechanics for Medicine*.

Yanovsky, I., Thompson, P.M., Osher, S., Leow, A.D., 2008b. Asymmetric and symmetric unbiased image registration: Statistical assessment of performance. *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW '08*. IEEE Computer Society Conference on, pp. 1-8.

Ye, X., Chen, Y., 2009. A New Algorithm for Inverse Consistent Image Registration. In: Bebis, G., Boyle, R., Parvin, B., Koracin, D., Kuno, Y., Wang, J., Wang, J.-X., Pajarola, R., Lindstrom, P., Hinkenjann, A., Encarnação, M., Silva, C., Coming, D. (Eds.), *Advances in Visual Computing*. Springer Berlin / Heidelberg, pp. 855-864.

Yeung, S.-K., Tang, C.-K., Shi, P., Pluim, J.P.W., Viergever, M.A., Chung, A.C.S., Shen, H.C., 2008. Enforcing stochastic inverse consistency in non-rigid image registration and matching. *Computer Vision and Pattern Recognition, 2008. CVPR 2008*. IEEE Conference on, pp. 1-8.

Yin, Y., Hoffman, E.A., Lin, C.-L., 2009. Mass preserving nonrigid registration of CT lung images using cubic B-spline. *Medical Physics* 36, 4213-4222.

Yushkevich, P.A., Avants, B.B., Das, S.R., Pluta, J., Altinay, M., Craige, C., 2010. Bias in estimation of hippocampal atrophy using deformation-based morphometry arises from asymmetric global normalization: An illustration in ADNI 3 T MRI data. *NeuroImage* 50, 434-445.

Zeng, Q., Chen, Y., 2008. Accurate Inverse Consistent Non-rigid Image Registration and Its Application on Automatic Re-contouring. In: Mandoiu, I., Sunderraman, R., Zelikovsky, A. (Eds.), *Bioinformatics Research and Applications*. Springer Berlin / Heidelberg, pp. 293-304.

Zhang, Z., Jiang, Y., Tsui, H., 2006. Consistent multi-modal non-rigid registration based on a variational approach. *Pattern Recognition Letters* 27, 715-725.