Metric-Based Pairwise and Multiple Image Registration

Qian Xie[1], Sebastian Kurtek[2], Eric Klassen[1], Gary E. Christensen[3], and Anuj Srivastava[1]

[1] Florida State University, Tallahassee, Florida, United States
qxie@stat.fsu.edu, klassen@math.fsu.edu, anuj@fsu.edu
[2] Ohio State University, Columbus, Ohio, United States
kurtek.1@stat.osu.edu
[3] University of Iowa, Iowa City, Iowa, United States
gary-christensen@uiowa.edu

Abstract. Registering pairs or groups of images is a widely-studied problem that has seen a variety of solutions in recent years. Most of these solutions are variational, using objective functions that should satisfy several basic and desired properties. In this paper, we pursue two additional properties – (1) invariance of objective function under identical warping of input images and (2) the objective function induces a proper metric on the set of equivalence classes of images – and motivate their importance. Then, a registration framework that satisfies these properties, using the L^2-norm between a novel representation of images, is introduced. Additionally, for multiple images, the induced metric enables us to compute a mean image, or a template, and perform joint registration. We demonstrate this framework using examples from a variety of image types and compare performances with some recent methods.

Keywords: metric-based registration, elastic image deformation, post-registration analysis, mean image, multiple registration.

1 Introduction

The problem of image registration is one of the most widely studied problems in medical image analysis. Given a set of observed images, the goal is to register points across the domains of these images. This problem has many names: registration, matching, correspondence, reparameterization, domain warping, deformation, etc., but the basic problem is essentially the same – which pixel/voxel on an image matches which pixel/voxel on the other image. The registration problem can be subdivided into categories in several ways. One way is to consider how many images are being registered: the *pairwise registration* where two images are matched and the *groupwise* or *multiple registration* where more than two images are being matched. Another possible division is based on modality - *unimodal registration* which is performed within a single modality and *multimodal registration* which is performed for images across multiple imaging modalities. In this

D. Fleet et al. (Eds.): ECCV 2014, Part II, LNCS 8690, pp. 236–250, 2014.

paper we will restrict to unimodal image registration since we are also interested in comparing and statistically analyzing registered images, beyond the problem of registration.

Although the registration problem has been studied for almost two decades, there continue to be some fundamental limitations in the popular solutions that make them suboptimal, difficult to evaluate and limited in scope. To explain these limitations let $\mathcal{F}$ be a certain set of $\mathbb{R}^n$-valued functions on a domain D, made precise later. A pairwise registration between any two images $f_1,\ f_2 \in \mathcal{F}$ is defined as finding a mapping γ, typically a diffeomorphism from D to itself, such that pixels $f_1(s)$ and $f_2(\gamma(s))$ are optimally matched to each other for all $s \in D$. To develop an algorithm for registration one needs: (1) an objective function for formalizing the notion of optimality, and (2) a numerical procedure for finding the optimal γ. Although the numerical techniques for optimization, i.e. item (2), have become quite mature over the last ten years, the commonly-used objective functions themselves have several fundamental shortcomings. It is the choice of objective function that is under the focus in this paper. The registration problems are commonly posed as variational problems, with the most common form of an objective function being

$$\mathcal{L}\left(f_1, f_2 \circ \gamma\right) \equiv \int_D \left\|f_1(s) - f_2(\gamma(s))\right\|^2 ds + \lambda R(\gamma),\ \gamma \in \mathit{\Gamma}, \tag{1}$$

where $\|\cdot\|$ is the Euclidean norm, R is a regularization penalty on γ, typically involving its first and/or second derivatives, λ is a positive constant and $\mathit{\Gamma}$ denotes the space of relevant deformations. Several variations of this functional are also used. We highlight shortcomings of these methods using a broader discussion about desired properties of an objective function in the next section.

1.1 Desired Properties in an Objective Function

We start with the question: What should be the properties of an objective function for registering images? The answer to this question is difficult since we may desire different results in different contexts. In fact, one can argue that we may never have a "perfect" objective function that matches an expert's intuition and solution. Still, there is a fundamental set of properties that is desirable, even essential, in for registration; some of these have been discussed previously in [3,17]. Some of them have been achieved in the previous papers while others have not. We list these properties next, starting with some notation. Note that some of them are overlapping, in the sense that they, individually or jointly with others, imply some others. Let $\mathcal{L}(f_1, (f_2, \gamma))$ denote the objective function for matching f_1 and f_2 by optimizing over γ (here γ is assumed to be applied to f_2 resulting in $(f_2, \gamma) \in \mathcal{F}$). The bracket (f, γ) denotes the group action where $\gamma \in \mathit{\Gamma}$ acts on $f \in \mathcal{F}$ defined by $(f, \gamma)(s) \equiv (f \circ \gamma)(s)$, $\forall s \in D$. Then, the *desired properties* of $\mathcal{L}$ are:

1. **Symmetry.** For any $f_1, f_2 \in \mathcal{F}$, we want

$$\mathcal{L}(f_1, f_2) = \mathcal{L}(f_2, f_1)\ .$$

2. **Positive Definiteness.** For any $f_1, f_2 \in \mathcal{F}$ we want $\mathcal{L}(f_1, f_2) \geq 0$ and

$$\mathcal{L}(f_1, f_2) = 0 \Leftrightarrow \ f_1 = f_2 \ \ a.\ e.\ .$$

3. **Lack of Bias.** If f_1, f_2 are constant functions then for any $\gamma \in \Gamma$,

$$\mathcal{L}(f_1, f_2) = \mathcal{L}(f_1, (f_2, \gamma))\ .$$

4. **Invariance to Identical Warping.** For any $f_1, f_2 \in \mathcal{F}$ and $\gamma \in \Gamma$, we have

$$\mathcal{L}(f_1, f_2) = \mathcal{L}\left((f_1, \gamma), (f_2, \gamma)\right)\ .$$

5. **Triangle Inequality.** For any $f_1,\ f_2,\ f_3 \in \mathcal{F}$,

$$\mathcal{L}(f_1, f_3) \leq \mathcal{L}(f_1, f_2) + \mathcal{L}(f_2, f_3)\ .$$

6. (An additional property of Γ.) Γ **is a group with composition.** For any γ,γ' and γ" $\in \Gamma$,
   i) $\gamma \circ \gamma' \in \Gamma$
   ii) $(\gamma \circ \gamma') \circ \gamma" = \gamma \circ (\gamma' \circ \gamma")$
   iii) there exists an $\gamma_{\text{id}} \in \Gamma$ such that $\gamma_{\text{id}} \circ \gamma = \gamma \circ \gamma_{\text{id}} = \gamma$
   iv) there exits a $h \in \Gamma$ such that $\gamma \circ h = h \circ \gamma = \gamma_{\text{id}}$.

Despite seemingly different appearances, the properties 1 to 4 are the same or closely related to those introduced previously in [17]. Specifically, properties 1, 4 and 6 together imply what was termed "*Symmetry*", and property 4 and 6 imply "*Invariance under* SDiff^+ but are actually stronger. Property 5 is introduced to the list so that properties 1, 2 and 5 altogether imply that $\mathcal{L}$ is a proper *metric* on $\mathcal{F}$.

Property 4: **Invariance to Identical Warping** is listed as a standalone property not only because it is fundamental but also one of the most important. Why? Consider the two images f_1 and f_2 shown in the left panel of Fig. 1. Even though the two images are different, their corresponding pixels are nicely aligned. The middle panel shows an example of a warping function γ to be applied to both images and the right panel shows the warped images $f_1 \circ \gamma$ and $f_2 \circ \gamma$. It is easy to see that the correspondence between pixels across two images remains unchanged. Thus, since $\mathcal{L}$ is a measure of registration or correspondence between images, we need $\mathcal{L}(f_1, f_2) = \mathcal{L}\left((f_1, \gamma), (f_2, \gamma)\right)$. However, if we take the commonly used L^2-norm as an objective function (the first term in Eqn. 1), then function values are not the same, as shown below the images. In summary, an identical warping of any two images keeps their registration unchanged and, hence, in order to properly measure the level of registration, an objective function must have this property of invariance to identical warping. We seek a framework that achieves all of the properties listed above.

To specify the proposed framework, define an equivalence relation between images as follows: let $f \sim g$ iff there is a $\gamma \in \Gamma$ such that $g = (f, \gamma)$, and let

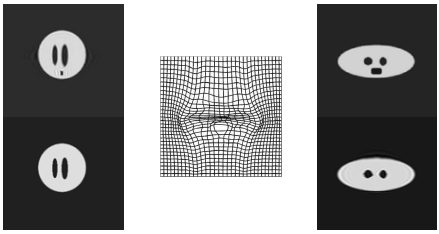

Fig. 1. Illustration of invariance to identical warping. The values of objective functions are not the same: $\|f_1 - f_2\|_{L^2} = 0.1085$ and $\|f_1 \circ \gamma - f_2 \circ \gamma\|_{L^2} = 0.1305$.

$[f] = \{g \mid g \sim f\}$ denote the orbit of an image f. Our goal is to establish a metric on the set of equivalence classes or orbits, the so-called quotient space of $\mathcal{F}$ under the action of Γ. However, since Γ is an open set, the resulting orbits are also open sets and it is difficult to establish a proper metric between these orbits. Thus, instead, we define the quotient space $\mathcal{F}/\Gamma$ as the set of *closures* of all equivalence classes, i.e. $\mathcal{F}/\Gamma = \{\text{closure}([f]) \mid f \in \mathcal{F}\}$. If $\mathcal{L}$ is a proper metric on $\mathcal{F}$ and, additionally, if properties 4 and 6 also hold, then it can be shown that the quantity $\inf_{\gamma\in\Gamma} \mathcal{L}(f_1, (f_2, \gamma))$ defines a proper metric on $\mathcal{F}/\Gamma$. The distance between any two equivalence classes $[f_1]$ and $[f_2]$ is well defined:

$$d([f_1], [f_2]) \equiv \inf_{\gamma\in\Gamma} \mathcal{L}(f_1, (f_2, \gamma)) \; = \inf_{\gamma'\in\Gamma} \mathcal{L}(f_2, (f_1, \gamma')) \; . \tag{2}$$

This metric on the quotient space therefore provides us with a tool to measure the difference between registered images. Note that in case $\mathcal{L}$ is a proper metric, then property 4 implies the action of Γ on $\mathcal{F}$ is by isometries. It can be shown that if $d([f_1], [f_2]) = 0$, then $f_1 \in [f_2]$ and vice-versa.

This setup allows us to study another important property – "**inverse consistency**" – introduced in [3]. It states that for all $f_1, f_2 \in \mathcal{F}$, if

$$\tilde{\gamma} \in \operatorname*{argmin}_{\gamma\in\Gamma} \mathcal{L}(f_1, (f_2, \gamma)) \text{ then } \tilde{\gamma}^{-1} \in \operatorname*{argmin}_{\gamma\in\Gamma} \mathcal{L}(f_2, (f_1, \gamma)).$$

It is natural to have this property since it implies that the optimal registration between two images remains the same even if they are treated in the reverse order. Note that a combination of 1, 4 and 6, along with the definition in Eqn. 2, implies inverse consistency.

The requirement for Γ to be a group is important to derive other properties. While most papers use the full diffeomorphism group for Γ, some papers work with a subset of deformations, e.g. spline-based deformations, and that can be problematic as discussed later.

1.2 Past and Current Literature

While the objective function given in Eqn. 1 is one used most often, several variations have also been applied. For instance, sometimes the first term is replaced by mutual information [23,4], minimum description length [5,21], etc., or

the second term is replaced by the length of a geodesic in the warping space (as in the LDDMM approach [6,20,1,13]). Some methods conceptualize the average image under the large deformation diffeomorphisms setting as an unbiased atlas ([8,12]). However, these methods do not use a formal metric for registration. Another idea is to impose regularization differently, e.g. using Gaussian smoothing of images (diffeomorphic demons [19,22]). Some methods optimize the objective function over a proper subgroup $\Gamma_0 \subset \Gamma$ (e.g. the set of volume-preserving diffeos [17]), some on Γ, some on larger a group Γ_b that contains Γ (e.g. the one including non-diffeomorphic mappings also) and some on even larger deformation spaces that are not groups (e.g. thin plate splines [2,16,7]). Rather than going through individual methods and their properties in details, we summarize their satisfaction of desired properties in Table 1. It is interesting to note that not a single past method satisfies property 4 without drastically restricting the deformation group Γ. The inverse consistency is, similarly, seldom satisfied.

Table 1. Properties of Objective Functions for Registration

Properties	1. Sym.	2. P.D.	3. Lack Bias	4. Inv.	5. Tri. Ineq.	6. Group
L^2	✓	✓	✓	✗	✓	✗
$L^2 + R(\gamma)$	✓	✓	✗	✗	✓	✗
CC [18]	✓	-	✗	✓	-	✓[1]
MI [23]	✓	-	✓	✗	-	✗
Demons [22]	✓	✓	✗	✗	✓	✗
LDDMM [1]	✓	✓	✗	✗	✓	✓
GL_2 [17]	✓	✓	✓	✓	✓	✓[2]
$GL_2 + R(\gamma)$	✓	✓	✗	✗	✓	✓

[1] is invariant to the general linear group.
[2] is invariant to the special group of volume and orientation preserving diffeomorphisms.
- indicates where the property is not proper to be evaluated under the context.

If we have a proper metric on the quotient space $\mathcal{F}/\Gamma$, it leads to additional tools for post-registration analysis. Here one analyzes registered images and applies statistical techniques such as PCA for dimension reduction and modeling. The question is: What should be the metric for these modeling procedures? Currently one performs registration using a certain objective function and then chooses a separate metric to perform post-registration analysis. Ideally, one would like an approach that can *align, compare, average, and model* multiple images in a **unified** framework that leads to efficient algorithms and consistent estimators. The objective function presented in this paper not only satisfies the invariance and the inverse consistency properties but also provides a metric on the quotient space for unified image comparison and analysis. Therefore, we have called our framework a metric-based method for registration, comparison and analysis of images. This idea was prompted by recent work on shape analysis of surfaces [10]. Although we utilize the same idea, the details are different because the representation and the Riemannian metric are not the same.

The rest of this paper is organized as follows. In Section 2, we introduce a new mathematical representation of images and a metric for image registration that satisfies all the desired properties. In Section 2.3, we develop the idea of mean images or templates under the chosen metric, and use these means to perform multiple image registration. Section 3 presents results on synthetic and real data. Finally, in Section 4, we present some concluding remarks.

2 Metric-Based Image Registration

In this section we lay out the framework for joint image registration and comparison under a new objective function which is a proper metric. This method applies to mathematical objects whose range space has dimension at least as high as that of their domain, as $f : D \to \mathbb{R}^n$, where $n \geq m = \dim(D)$. In case of 2D (3D) images, this means that pixels have at least two (three) coordinates which is the case for colored images, or multimodal images (with different imaging modalities as the pixel coordinates). The scalar (gray scaled) images can be transformed as described later to satisfy this condition.

2.1 Image Representation and Pairwise Registration

Let the image space be $\mathcal{F} = \{f : D \to \mathbb{R}^n \mid f \in C^\infty(D)\}$ and $\Gamma = \mathrm{Diff}^+(D)$ is a subgroup of Diff^+ (the orientation-preserving diffeomorphism group) that preserves the boundary of D. Hereafter we will use $\|f\|$ for the L^2-norm of any f and $|A|$ to denote the determinant of a square matrix A unless stated otherwise.

Definition 1. *The right action of Γ on $\mathcal{F}$ is defined by the mapping $\mathcal{F} \times \Gamma \to \mathcal{F}$ given by $(f, \gamma) = f \circ \gamma$.*

It is easy to see that this action is not by isometries under the L^2-metric. That is, for any two $f_1, f_2 \in \mathcal{F}$, and a general $\gamma \in \Gamma$, we have $\|f_1 - f_2\| \neq \|f_1 \circ \gamma - f_2 \circ \gamma\|$. Thus, the important property 4, invariance to simultaneous warping, is not satisfied and, consequently, one cannot work with the L^2-norm in the image space directly. Instead, we will use a mathematical representation of images defined by a mapping called the q-map, that has been prompted by recent work in shape analysis of surfaces [10]. Here, we adopt it for analyzing images as follows.

Let $(x^1, \ldots, x^m) : D \to \mathbb{R}^m$ be coordinates on (a chart of) D and $\mathbf{J}f(s)$ be the Jacobian matrix of f at s with the (j, i)-th element as $\partial f^j / \partial x^i(s)$. Define the "generalized area multiplication factor" of f at s for arbitrary $n \geq m$ as $a(s) = |\mathbf{J}f(s)|_V$ where $|\mathbf{J}f(s)|_V = \|\frac{\partial f}{\partial x^1} \wedge \frac{\partial f}{\partial x^2} \wedge \cdots \wedge \frac{\partial f}{\partial x^m}\|$. Here $\wedge$ denotes the wedge product. The two special cases are: if $m = n = 2$, then $a(s) = |\mathbf{J}f(s)|$; if $m = 2$ and $n = 3$, then $a(s) = \|\frac{\partial f}{\partial x^1}(s) \times \frac{\partial f}{\partial x^2}(s)\|$.

Definition 2. *For an $f \in \mathcal{F}$, define a mapping $Q : \mathcal{F} \to L^2$ such that for any $s \in D$, $Q(f)(s) = \sqrt{a(s)} f(s)$.*

For any $f \in \mathcal{F}$, we will refer to $q = Q(f)$ as its q**-map**; note that $q : D \to \mathbb{R}^m$. Also, we remark that this is a general version of q-map for arbitrary $\mathbb{R}^n$ that extends the work of [10]. Assuming the original set of images to be smooth, the set of all q-maps is a subset of the L^2-space. Intuitively, the q-map leaves uniform regions as zeros while preserving edge information in such a way that it is compatible with change of variables, i.e., stronger edges get higher values. The corresponding action of Γ on L^2 is given as follows.

Lemma 1. *The right action of Γ on the L^2-space, corresponding to the one given in Definition 1, is given by the mapping $L^2 \times \Gamma \to L^2$ as $(q, \gamma) = \sqrt{|\mathbf{J}\gamma|}(q \circ \gamma)$, where $\mathbf{J}\gamma(s)$ denotes the Jacobian matrix of γ at s.*

Note that the mapping Q is equivariant, i.e. it can be shown that for an image f and γ, $Q((f, \gamma)) = (Q(f), \gamma)$, where the action on the left side is given by Definition 1 and the action on the right is given by Definition 1.

This leads to the most important property of this mathematical representation as the following.

Proposition 1. *The reparametrization group Γ acts on the L^2-space by isometries under the L^2-norm, i.e. $\forall q_1, q_2 \in L^2, \forall \gamma \in \Gamma, \|(q_1, \gamma) - (q_2, \gamma)\| = \|q_1 - q_2\|$.*

Proof.

$$\|(q_1, \gamma) - (q_2, \gamma)\|^2 = \int_D |q_1(\gamma(s)) - q_2(\gamma(s))|^2 \, |\mathbf{J}\gamma(s)| \, ds = \|q_1 - q_2\|^2 \ . \quad (3)$$

Setting $q_2 \equiv 0$, $\forall q_1 \in L^2$ and $\gamma \in \Gamma$, we have $\|q_1\| = \|(q_1, \gamma)\|$. Thus, warping of images is actually a unitary operator under this representation.

Definition 3. *Define an objective function between any two images f_1 and f_2, represented by their q-maps q_1 and q_2, as $\mathcal{L}(f_1, (f_2, \gamma)) \equiv \|Q(f_1) - Q((f_2, \gamma))\| = \|q_1 - (q_2, \gamma)\|$.*

The registration is then achieved by minimizing the **objective function**:

$$\gamma^* = \underset{\gamma \in \Gamma}{\operatorname{arginf}} \, \mathcal{L}(f_1, (f_2, \gamma)) = \underset{\gamma \in \Gamma}{\operatorname{arginf}} \, \|q_1 - (q_2, \gamma)\| \ , \quad (4)$$

Upon closer inspection, Proposition 1 is exactly the same as property 4 – invariance to identical warping – in Section 1. In view of this, the L^2-norm between q-maps of images becomes a proper measure of registration between images since it remains the same if the registration is unchanged. This leads to a quantity that will serve as both the registration objective function and an extrinsic distance between registered images as defined in Eqn. 2. We will refer to the method as MBIR, i.e. , the metric-based image registration method.

The objective function $\mathcal{L}$ given in Definition 3 satisfies **all** the properties listed earlier as those desired for registration. Specifically, it satisfies **invariance to simultaneous warping** (property 4). In case our transformation model Γ is a group then this framework satisfies the **inverse consistency** as stated in

Section 1. Additionally, the optimal registration is not affected by scaling and translations of image pixels: let $g_1 = c_1 f_1 + d_1$ and $g_2 = c_2 f_2 + d_2$ with $c_1, c_2 \geq 0$ and $d_1, d_2 \in \mathbb{R}^n$, if $\gamma^* = \arg\inf_\gamma \mathcal{L}(f_1, (f_2, \gamma))$ then $\gamma^* = \arg\inf_\gamma \mathcal{L}(g_1, (g_2, \gamma))$ as well. We point out that there are some unresolved mathematical issues concerning the existence of a unique global solution for γ^*, especially its existence inside Γ rather than being on its boundary. We leave this for a future discussion and focus on a numerical approach that estimates γ^*.

The proposed objective function in Eqn. 4 has only one term (similarity term) and the regularity term appears to be missing. However, the similarity term also has a built-in regularity, since it includes the determinant of the Jacobian, $|J_\gamma|$, in (q, γ). Additional regularity can also be introduced to the framework as is done in the LDDMM framework.

2.2 Gradient Method for Optimization Over Γ

The optimization problem over Γ stated in Eqn. 4 forms the crux of our registration framework and we will use a **gradient descent** method to solve it. Since Γ is a group, we use the gradient to solve for the incremental warping γ, on top of the previous cumulative warping γ_o, as follows. (In this way the required gradient is an element of $T_{\gamma_{\mathrm{id}}}(\Gamma)$, the tangent space of Γ at identity γ_{id}, and one needs to understand only that space.) We define a cost function with respect to γ as the functional

$$E(\gamma) = \|q_1 - \phi_{\tilde{q}_2}(\gamma)\|^2 \,, \tag{5}$$

where $\phi_q : \Gamma \mapsto [q]$ is defined to be $\phi_q(\gamma) = (q, \gamma)$ and $\tilde{q}_2 = (q_2, \gamma_o)$ with γ_o being the current deformation. Given a set of orthonormal basis elements, say $\mathbf{B}$, of $T_{\gamma_{\mathrm{id}}}(\Gamma)$, the gradient at γ_{id} takes the form $\nabla E(\gamma_{\mathrm{id}}) = \sum_{b \in \mathbf{B}} (\nabla_b E(\gamma_{\mathrm{id}}))\, b$, where $\nabla_b E(\gamma_{\mathrm{id}})$ is the directional derivative of E at γ_{id}. Let $\phi_{q,*}$ be the differential of ϕ_q at γ_{id} in the direction of b. Then $\nabla_b E(\gamma_{\mathrm{id}}) = \langle q_1 - \phi_{\tilde{q}_2}(\gamma_{id}), \phi_{q,*}(b) \rangle$. Brackets denote the L^2 inner products. There exists an explicit form of $\phi_{q,*}$ such that for $b \in T_{\gamma_{\mathrm{id}}}(\Gamma)$ and $j = 1, 2, \ldots, n$, the coordinate functions of $\phi_{q,*}(b)$ are given by:

$$\phi^j_{q,*}(b) = \frac{1}{2}(\nabla \cdot b) q^j + (\nabla q^j) b \,, \tag{6}$$

where $\nabla\cdot$ denotes the divergence operator. The form is the same as that is derived for parameterized surfaces in [10]. The basis elements are constructed using ideas of Fourier basis functions (See [24] for details).

2.3 Distance in the Quotient Space

Recall that in the LDDMM framework the regularity part of the objective function comes from a proper distance on Γ, computed as geodesic length under a Riemannian metric. If we study two images within an equivalence class using LDDMM (as A shown in Fig. 2), then the first term will go to zero and only the regularity term will remain. Thus, within equivalence classes, LDDMM provides a proper distance for comparing images. However, when two images are not

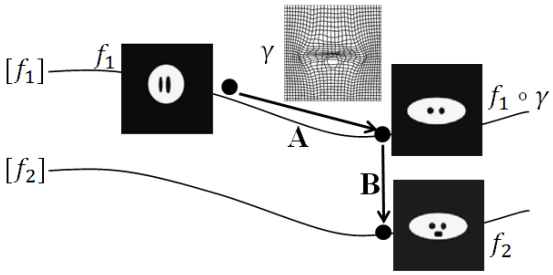

Fig. 2. Illustration of quotient space and orbits

equivalent, the variation left after registration is nonzero and needs a metric for analysis. In other words, one needs a metric between equivalence classes (shown in Fig. 2 as B), or a metric on the orbit space. Our framework naturally induces a metric of that type. In case the variation in Γ is also of importance, we can combine our metric (B) with a metric on A for analysis.

As stated in Eqn. 2, the minimal value of the objective function $\mathcal{L}$ introduced in Definition 3 induces a distance in the quotient space $\mathcal{F}/\Gamma$. To explain further, we define $\overline{[q]}$ to be the set of all warpings of a q-map. Since all elements of $\overline{[q]}$ can be obtained using warpings (including the boundary of the orbit) of the same image, we deem them equivalent from the perspective of registration. Let L^2/Γ be the (quotient) set of all such equivalence classes of q-maps. Define

$$d([f_1],[f_2]) = d([q_1],[q_2]) \equiv \inf_{\gamma\in\Gamma} \mathcal{L}(f_1,(f_2,\gamma)) = \inf_{\gamma\in\Gamma} \|q_1 - (q_2,\gamma)\| \ , \qquad (7)$$

It can be shown that the quantity $d([q_1],[q_2])$ (or $d([f_1],[f_2])$) forms a proper distance on the quotient space L^2/Γ. At the same time since $g \in [f]$ indicates $Q(g) \in [Q(f)]$, this quantity is a proper measure to quantify the level of registration.

Mean Image and Groupwise Registration

An important problem in image analysis, especially medical image analysis, is to compute a "typical" or an "average" of several images from the same class and use it as a template. Then, the individual images can be registered to the sample mean in a pairwise manner, resulting in a group registration. By registering member images to the group mean, one can analyze their variations from the typical template image. Suppose there is a set of N images, $\{f_i\}_{i=1}^N$. Their Karcher mean is defined as the image that minimizes the sum of squares of the distances to the given images, i.e. $[\mu] = \text{argmin}_{f\in\mathcal{F}} \sum_{i=1}^n d^2([f_i],[f])$. The algorithm to find the Karcher mean is a standard one, and helps us find a mean image f_μ deformable to the underlying image, such that $f_\mu \in [\mu]$. With this mean, we can register groups of images to the mean image rather than to an arbitrarily chosen template, as is often done in current methods.

3 Experiments

In this section we present various image registration results in order to validate our method. We provide examples of pairwise registration on both synthetic images and brain MRIs. In order to improve the registration of images with larger deformations, we also show results in landmark-aided registration for a better solution. We demonstrate the utility of our method to compute mean images as templates for registering multiple image. The problem of image classification is also considered using the proposed metric and the results are compared to those from other methods.

Recall that in case of grayscale images, with $n = 1$, our method does not apply directly since $n < dim(D)$. Instead, we make use of image gradients $\nabla f = (f_u, f_v)$ for $(u, v) \in [0, 1]^2$ and register objects in the form of $g = (f, f_u, f_v) \in \mathbb{R}^3$. In other words, the vector-valued image $g : D \to \mathbb{R}^3$ forms the input data for registration. Such image gradients are a type of edge measure and are often used in their own right as robust spatial features for image registration.

3.1 Pairwise Image Registration

We first present some results on synthetic images to demonstrate the use of the registration framework suggested in Eqn. 4. Fig. 3 shows images f_1 and f_2 that are registered twice, first by taking f_1 as the template image and estimating γ_{21} that optimally deforms f_2 using Eqn. 4. Then, the roles are reversed and f_2 is used as the template to obtain γ_{12}. We show the two converged objective functions, $\|(q_1, \gamma_{12}) - q_2\|$ and $\|q_1 - (q_2, \gamma_{21})\|$, associated with the optimal γ_{12} and γ_{21} to verify symmetry. The cumulative diffeomorphisms $\gamma_{21} \circ \gamma_{12}$ and $\gamma_{12} \circ \gamma_{21}$ are also used to demonstrate the inverse consistency of the proposed metric. As mentioned above, the theory shows that γ_{12} and γ_{21} are expected to be inverses of each other. We show the original images f_1 and f_2 with the matching warped images $f_2 \circ \gamma_{21}$ and $f_1 \circ \gamma_{12}$ respectively. The diffeomorphisms γ_{12} and γ_{21} used to register the images are also presented. By composing them in different orders we expect the resulting diffeomorphisms to be the identity map. In order to better visualize that the composed diffeomorphisms are close to identity, we apply them to checkerboard images and observe that the composed diffeomorphisms $\gamma_{21} \circ \gamma_{12}$ and $\gamma_{12} \circ \gamma_{21}$ are close to the identity map.

In Fig. 4, we present registration results using 2D brain MR images. In order to illustrate our method, in each of the two experiments, we show (1) the original images overlapped f_1/f_2 and (2) overlapped images after registration ($f_1/f_2 \circ \gamma_{21}$ and $f_2/f_1 \circ \gamma_{12}$). The overlapped images show image pairs in a common canvas such that red and green denote positive and negative image differences respectively.

Landmark-Aided Registration. Our framework can be extended to incorporate landmark information during registration and all of the nice mathematical properties of the objective function are preserved. Assume that there are a fixed

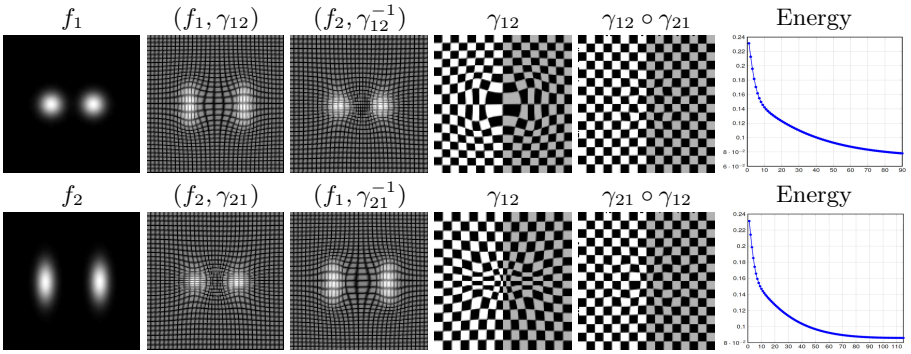

Fig. 3. Registering Synthetic Smooth Grayscale Images. $\gamma_{12} = \text{argmin}_{\gamma \in \Gamma} \mathcal{L}(f_2, (f_1, \gamma))$ and $\gamma_{21} = \text{argmin}_{\gamma \in \Gamma} \mathcal{L}(f_1, (f_2, \gamma))$. $\|q_1 - q_2\| = 0.2312$, $\|q_1 - (q_2, \gamma_{21})\| = 0.0728$ and $\|(q_1, \gamma_{12}) - q_2\| = 0.0859$.

number, say K, of distinct landmark points, $\mathcal{P} = \{p_1, p_2, \ldots, p_K\}$, in the image domain D. They are typically chosen according to the application but fixed within the analysis. The landmark-guided registration is achieved by defining a subgroup of Γ, denoted by $\Gamma_{\mathcal{P}}$, as:

$$\Gamma_{\mathcal{P}} = \{\gamma \in \Gamma \quad | \gamma(p_i) = p_i, \quad i = 1, 2, \ldots, K\} \; . \tag{8}$$

Given two images f_1, f_2 with landmark information $\mathcal{P}$, the images can be registered in two steps in an iterative way:

1) Register the landmark points $\mathcal{P}$ and apply an initial deformation to f_1 to form $f_1^{\mathcal{P}}$ such that the landmarks are at the same locations in $f_1^{\mathcal{P}}$ and f_2.
2) Register $f_1^{\mathcal{P}}$ to f_2 using Eqn. 4 restricting to the subgroup $\Gamma_{\mathcal{P}}$.

Similar technique of forming landmark-constrained basis on $\mathbb{S}^2$ has been used on closed surfaces as described in [11]. In the second step, searching over $\Gamma_{\mathcal{P}}$ ensures correspondences of landmarks established in step 1 are preserved. The registration is refined without moving the landmarks. This search is based on a basis $\mathbf{B}^{\mathcal{P}}$ in $T_{\gamma_{id}}(\Gamma_{\mathcal{P}})$ constructed such that its elements, the vector fields on D, vanish at the landmark locations $\{p_i\}$.

We remark that $\Gamma_{\mathcal{P}}$ forms a subgroup of Γ and, as a result, the desired properties discussed earlier are still satisfied. This approach may be termed *landmark-guided* registration where the landmarks are treated as hard constraints.

We show results on MRIs with landmarks in Fig. 5. In the first example, presented in the first row of Fig. 5, the optimally deformed f_1 is displayed at the end as (f_1, γ). The deformation in the skull is so large that our original method fails to reach a global minimizer of the objective function. By adopting the landmark-aided registration, we at first get a deformed image $f_1^{\mathcal{P}}$, with nicely matched landmarks and the skull deformed correspondingly. Then, $f_1^{\mathcal{P}}$ is further deformed to register the intensity details without moving the landmarks. The final result $(f_1^{\mathcal{P}}, \gamma)$ matches f_2 with no artifacts around the skull. Another

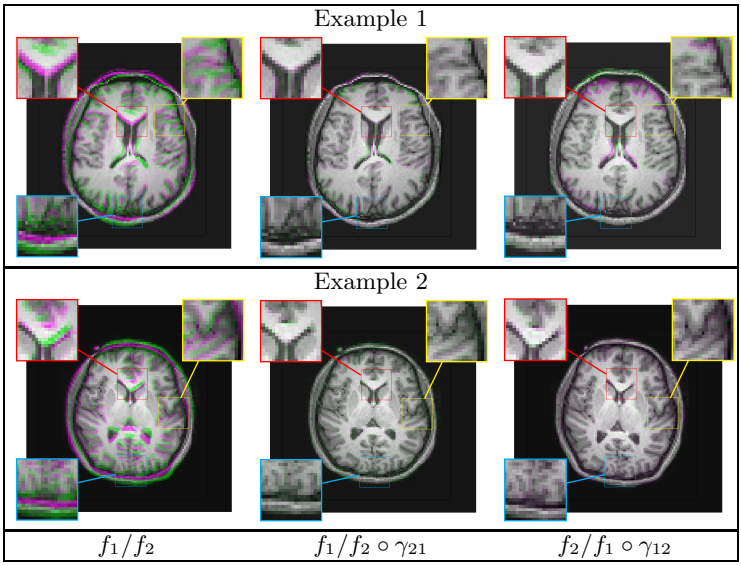

Fig. 4. Two examples of brain MR image registration. First columns show overlapped original images f_1 and f_2; second columns show overlapped images f_1 and deformed f_2; third columns show f_2 and deformed f_1.

example is shown in the bottom two rows of Fig. 5. Generally, the registration with landmarks outperforms the identity map as the initial condition of our procedure, especially when the deformations are large.

3.2 Registering Multiple Images

We use part of the MNIST database of handwritten digits to illustrate our method to compute the mean image and multiple image registration. In Fig. 6 we present the mean image of each digit computed without and with registration, respectively. We also show an example using brain MRIs in Fig. 7. Four brain images without alignment are shown on the top row with the corresponding mean image. This mean without registration appears blurred due to misalignment. On the bottom row, the images are aligned to the Karcher mean as described in Section 2.3. We can see that with multiple registration the mean image imporves the bluriness issue.

3.3 Image Classification

The framework introduced in this paper establishes a proper distance on the quotient space of q-maps of images. These distances can be used for pattern analysis of images such as clustering or classification. To illustrate this idea, we use a subset of the MNIST database of images of handwritten digits from 0 to 9. It contains ten images of handwritings for each digit. In addition to the

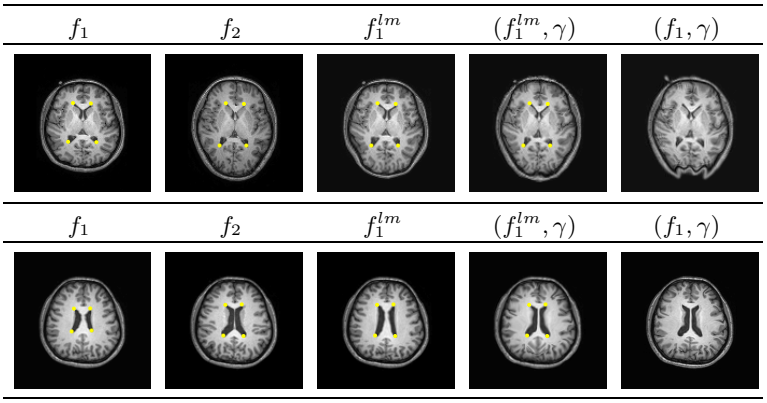

Fig. 5. Two examples of brain image registration with landmarks. First two columns show original images f_1 and f_2. The third column shows the deformed images f_1^{lm} using only landmarks; forth column shows final deformed images $(f_1^{\text{lm}}, \gamma)$ with $f_1^{\mathcal{P}}$ as the initial condition; the last column shows registered images (f_1, γ) without involving landmarks.

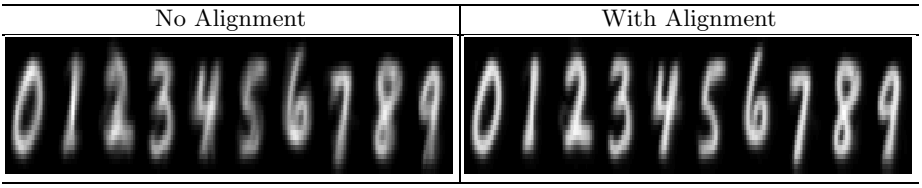

Fig. 6. First row contains the mean image without registration for digit groups (0-9); second row contains corresponding mean images with registration

baseline L^2-distance, or sum of squared distances (SSD) (without any warping), we compare our method to three other methods - diffeomorphic demons [22], FAIR [15] and NiftyReg [14].

For computing distance matrices between all pairs of images, the digits are registered in a pairwise manner using each of the three methods and then the SSD is computed for the registered images as a measure of distance. In the case of our method the distance defined in Eqn. 7 is used. Using the leave-one-out nearest-neighbor (LOO-NN) classifier, the rates of correct classification are listed in Table 2. Provided as a baseline, the L^2-distance without any registration provides a rate of classification of 76% with our method performing the best with 94%.

Table 2. Classification of MNIST Digits

Method	MBIR	Demons	FAIR	NiftyReg
%	94	86	85	83

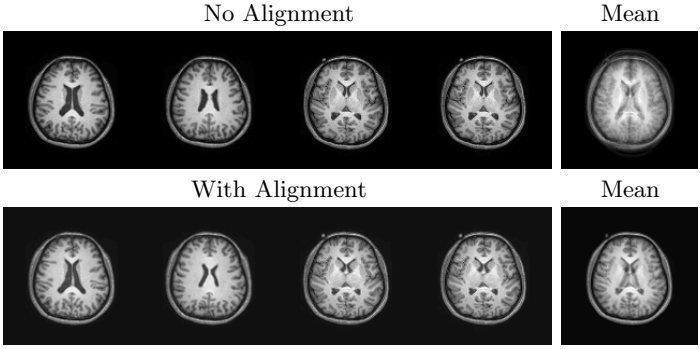

Fig. 7. Mean images of brain MRIs. Upper row: unregistered images and the cross-sectional mean; bottom row: mean with registration and images registered to it.

4 Conclusion

We have proposed a novel framework to register, compare and analyze images in a unified manner. This framework results in an objective function for registration that is both inverse consistent and invariant to random warpings of images. Furthermore, this function forms a proper metric on the quotient space of images, modulo the deformation group, and can be used to define and compute sample means of given images. This last item is based on computing an extrinsic distance between images in the representation space $\mathbb{L}^2$. With this framework, our method gives better results for pairwise registration and comparison, and multiple image registration and analysis. Furthermore, it allows the use of pre-determined (registered) landmark on images to help improve registration performance.

Acknowledgement. This research was supported in part by the NSF grants DMS 1208959, IIS 1217515, and CCF 1319658. We also thank the producers of datasets used here for making them available to public.

References

1. Beg, M., Miller, M., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International Journal of Computer Vision 61, 139–157 (2005)
2. Bookstein, F.L.: Principal warps: Thin-plate splines and the decomposition of deformations. IEEE Transactions on Pattern Analysis and Machine Intelligence 11(6), 567–585 (1989)
3. Christensen, G., Johnson, H.: Consistent image registration. IEEE Transactions on Medical Imaging 20(7), 568–582 (2001)
4. Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P.: Multimodality image registration by maximization of mutual information. IEEE Transactions on Medical Imaging 16(2), 187–198 (1997)
5. Davies, R., Twining, C., Cootes, T., Waterton, J., Taylor, C.: A minimum description length approach to statistical shape modeling. IEEE Transactions on Medical Imaging 21(5), 525–537 (2002)

6. Dupuis, P., Grenander, U.: Variational problems on flows of diffeomorphisms for image matching. Journal Quarterly of Applied Mathematics LVI (3), 587–600 (1998)
7. Eriksson, A., Astrom, K.: Bijective image registration using thin-plate splines. In: International Conference on Pattern Recognition, vol. 3, pp. 798–801 (2006)
8. Joshi, S., Davis, B., Jomier, B.M., Gerig, G.: Unbiased diffeomorphic atlas construction for computational anatomy. Neuroimage 23, 151–160 (2004)
9. Kurtek, S., Klassen, E., Ding, Z., Jacobson, S., Jacobson, J., Avison, M., Srivastava, A.: Parameterization-invariant shape comparisons of anatomical surfaces. IEEE Transactions on Medical Imaging 30, 849–858 (2011)
10. Kurtek, S., Klassen, E., Ding, Z., Srivastava, A.: A novel Riemannian framework for shape analysis of 3D objects. In: 2010 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1625–1632 (2010)
11. Kurtek, S., Srivastava, A., Klassen, E., Laga, H.: Landmark-guided elastic shape analysis of spherically-parameterized surfaces. In: Computer Graphics Forum (Proceedings of Eurographics 2013, vol. 32(2), pp. 429–438 (2013)
12. Lorenzen, P., Davis, B., Joshi, S.: Unbiased atlas formation via large deformations metric mapping. In: Duncan, J.S., Gerig, G. (eds.) MICCAI 2005. LNCS, vol. 3750, pp. 411–418. Springer, Heidelberg (2005)
13. Miller, M., Trouve, A., Younes, L.: On the metrics and Euler-Lagrange equations of computational anatomy. Annual Review of Biomedical Engineering 4, 375–405 (2002)
14. Modat, M., Cardoso, M., Daga, P., Cash, D., Fox, N., Ourselin, S.: Inverse-consistent symmetric free form deformation. In: Dawant, B.M., Christensen, G.E., Fitzpatrick, J.M., Rueckert, D. (eds.) WBIR 2012. LNCS, vol. 7359, pp. 79–88. Springer, Heidelberg (2012)
15. Modersitzki, J.: FAIR: Flexible Algorithms for Image Registration. Society for Industrial and Applied Mathematics (2009)
16. Szeliski, R., Coughlan, J.: Spline-based image registration. International Journal of Computer Vision 22(3), 199–218 (1997)
17. Tagare, H., Groisser, D., Skrinjar, O.: Symmetric non-rigid registration: A geometric theory and some numerical techniques. Journal of Mathematical Imaging and Vision 34(1), 61–88 (2009)
18. Taquet, M., Macq, B., Warfield, S.: A generalized correlation coefficient: application to DTI and multi-fiber DTI. In: Mathematical Methods in Biomedical Image Analysis (2012)
19. Thirion, J.: Image matching as a diffusion process: an analogy with Maxwell's demons. Medical Image Analysis 2(3), 243–260 (1998)
20. Trouve, A.: Diffeomorphisms groups and pattern matching in image analysis. International Journal of Computer Vision 28(3), 213–221 (1998)
21. Twining, C., Marsland, S., Taylor, C.: Groupwise non-rigid registration: The minimum description length approach. In: Proceedings of the British Machine Vision Converence (BMVC), vol. 1, pp. 417–426 (2004)
22. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. NeuroImage 45(suppl. 1), S61–S72 (2009)
23. Viola, P., Wells III, W.: Alignment by maximization of mutual information. In: Fifth International Conference on Computer Vision, pp. 16–23 (June 1995)
24. Xie, Q., Kurtek, S., Christensen, G., Ding, Z., Klassen, E., Srivastava, A.: A novel framework for metric-based image registration. In: Dawant, B.M., Christensen, G.E., Fitzpatrick, J.M., Rueckert, D. (eds.) WBIR 2012. LNCS, vol. 7359, pp. 276–285. Springer, Heidelberg (2012)