Slice-to-Volume Image Registration Models for MRI-Guided Cardiac Procedures

L.W. Lorraine Ma and Mehran Ebrahimi(✉)

Faculty of Science, University of Ontario Institute of Technology,
2000 Simcoe Street North, Oshawa, ON L1H 7K4, Canada
{lok.ma,mehran.ebrahimi}@uoit.ca

Abstract. A mathematical formulation for intensity-based slice-to-volume registration is proposed. The approach is flexible and accommodates various regularization schemes, similarity measures, and optimizers. The framework is evaluated by registering 2D and 3D cardiac magnetic resonance (MR) images obtained *in vivo*, aimed at real-time MR-guided applications. Rigid-body and affine transformations are used to validate the parametric model. Target registration error (TRE), Jaccard, and Dice indices are used to evaluate the algorithm and demonstrate the accuracy of the registration scheme on both simulated and clinical data. Registration with the affine model appeared to be more robust than with the rigid model in controlled cases. By simply extending the rigid model to an affine model, alignment of the cardiac region generally improved, without the need for complex dissimilarity measures or regularizers.

1 Introduction

Recently, there has been increased interest in using magnetic resonance imaging (MRI) for image-guided procedures that have traditionally been guided by X-ray imaging. In some patients with a history of myocardial infarction (MI), electrical activity in the heart may be disrupted by substrate formed from a previous MI, triggering arrhythmia. Treatment options include catheter ablation, where the offending substrate is surgically ablated to correct the arrhythmia. Catheter ablation is traditionally guided by X-ray fluoroscopy, however, possible concerns over radiation exposure have led to MRI being proposed as an alternative. Besides being a non-invasive imaging modality, other advantages of MRI include superior soft tissue contrast to better image anatomical features in and around the heart, the ability to capture depth information without multiple projections and to easily adjust the positions of imaging planes to access areas of interest. MRI-guided procedures necessitate fast imaging techniques to capture images in real-time. Fortunately, MR sequences for real-time 2D visualization exist and have been used to guide cardiac procedures [1].

The short acquisition time required by real-time MRI means that the 2D images lack in quality compared to slices of cine MR volumes that are acquired

M. Pop and G.A. Wright (Eds.): FIMH 2017, LNCS 10263, pp. 139–151, 2017.
DOI: 10.1007/978-3-319-59448-4_14

prior to surgical intervention without the constraints of producing images at real-time frame rates and are therefore less noisy. 3D pre-operative images provide detailed anatomical information while 2D intra-operative images provide live positional updates. Ideally, one can register the pre- and intra-operative images together to combine the advantages of both.

Various methods to perform slice-to-volume registration have been proposed. Of interest to us is the application of slice-to-volume registration to MRI-guided procedures. Registration between pre- and intra-operative images has been studied for non-cardiac applications in humans [2], but on a more relevant note, work has been done *in vivo* on swine to register 2D intra-operative cardiac MR image slices to pre-operative MR image volumes [3]. Xu et al. also present registration of high-quality pre-operative MR image volumes to live cardiac MR images on human volunteers *in vivo* with applications to MRI-guided radiofrequency ablation of substrate in the heart, but also focuses on registration incorporating rigid-body transformations [4].

In most of the studies mentioned above, rigid-body registration was employed. While rigid registration is generally employed to reduce computational cost and to speed up the registration process, it risks oversimplifying the displacement of body tissues, which are generally not rigid. The highly deformable nature of the heart and displacement at various stages of the breathing cycle make registration of the cardiac region more challenging. Deformable registration may be more accurate, but is computationally much more expensive.

Work involving slice-to-volume registration has been recent and not nearly as numerous as projective 2D-3D registration, especially with respect to applications in MRI-guided procedures. In addition, there seems to be a lack of a precise model in the literature, in contrast to 2D-2D or 3D-3D registration [5].

We propose a general mathematical framework for slice-to-volume registration which can accommodate parametric and non-parametric transformation models. A rigid transformation model can be used in this framework, but the user can easily adapt a different parametric transformation model.

We will demonstrate this framework on parametric models, specifically, using this framework to extend existing 2D-3D rigid registration to affine registration. Although the number of parameters in an affine parametric model (12 parameters) is twice the number of parameters in a rigid model (6 parameters), the figure dwarfs in comparison to the number of parameters dealt with in deformable registration, and thus is still a computationally inexpensive method that accounts for some non-rigid deformations. The intensity-based registration framework is flexible and can accommodate various models and parameters. We demonstrate by registering high-resolution 3D MR images to noisier 2D real-time MR images, using rigid and affine parametric models, and investigate the ill-posedness of 2D-3D registration as an inverse problem.

2 Model

Consider the registration problem of a 3D 'template' image $\mathcal{T}$ to a 2D 'reference' image $\mathcal{R}$, where $\mathcal{R}$ is a realization of $\mathcal{T}$ deformed via a transformations y and

sliced at a certain location z. The reference and template images are represented by mappings $\mathcal{R} : \Omega \subset \mathbb{R}^2 \to \mathbb{R}$ and $\mathcal{T} : \Omega \times \mathcal{Z} \subset \mathbb{R}^3 \to \mathbb{R}$ of compact support. Considering a slice location z, the goal is to find the transformation $y : \mathbb{R}^3 \to \mathbb{R}^3$ such that $\mathcal{L}_z(\mathcal{T}[y])$ is similar to $\mathcal{R}$, in which $\mathcal{T}[y]$ is the transformed template image and $\mathcal{L}_z : \mathbb{L}^2(\Omega \times \mathcal{Z}) \to \mathbb{L}^2(\Omega)$ is the slicing operator at level $z \in \mathcal{Z} \subset \mathbb{R}$, where $\mathcal{L}_z(\mathcal{T}(x^1, x^2, x^3)) := \mathcal{T}(x^1, x^2, z)$ for $(x^1, x^2, x^3) \in \mathbb{R}^3$. A formulation of the 2D-3D image registration of a template image $\mathcal{T}$ to a reference image $\mathcal{R}$ can be written as the following problem.

2D-3D Image Registration Problem: Given two images $\mathcal{R} : \Omega \subset \mathbb{R}^2 \to \mathbb{R}$ and $\mathcal{T} : \Omega \times \mathcal{Z} \subset \mathbb{R}^3 \to \mathbb{R}$ and an arbitrary given slice location $z \in \mathbb{R}$, find a transformation $y : \mathbb{R}^3 \to \mathbb{R}^3$ that minimizes the objective functional

$$\mathcal{J}[y] := \mathcal{D}[\mathcal{L}_z(\mathcal{T}[y]), \mathcal{R}] + \mathcal{S}[y - y^{\text{ref}}]. \tag{1}$$

Here, $\mathcal{D}$ is a distance that measures the dissimilarity of $\mathcal{L}_z(\mathcal{T}[y])$ and $\mathcal{R}$, and $\mathcal{S}$ is a regularization expression on the transformation y that penalizes transformations "away" from y^{ref}.

2.1 Parametric 2D-3D Registration

It is possible that y can be parametrized via parameters w. For example if y is an affine transformation, the transformation on a point $x = (x^1, x^2, x^3)$ can be expressed as

$$y(w; x) = \begin{pmatrix} w_1 & w_2 & w_3 \\ w_5 & w_6 & w_7 \\ w_9 & w_{10} & w_{11} \end{pmatrix} \begin{pmatrix} x^1 \\ x^2 \\ x^3 \end{pmatrix} + \begin{pmatrix} w_4 \\ w_8 \\ w_{12} \end{pmatrix}.$$

In general, for the parametric registration problem we equivalently aim to minimize

$$\mathcal{J}[w] := \mathcal{D}[\mathcal{L}_z(\mathcal{T}[y(w)]), \mathcal{R}] + \mathcal{S}[w - w^{\text{ref}}]. \tag{2}$$

Here we assume sum of squared distances (SSD) is the dissimilarity measure $\mathcal{D}$

$$\mathcal{D}[\mathcal{L}_z(\mathcal{T}), \mathcal{R}] = \mathcal{D}^{\text{SSD}}[\mathcal{L}_z(\mathcal{T}), \mathcal{R}] := \frac{1}{2} \int_\Omega (\mathcal{L}_z(\mathcal{T}(x)) - \mathcal{R}(x))^2 \, dx.$$

Furthermore, the regularization functional $\mathcal{S}$ can be defined as

$$\mathcal{S}[w - w^{\text{ref}}] := \frac{1}{2} \times (w - w^{\text{ref}})^T \, \mathbf{M} \, (w - w^{\text{ref}}) \tag{3}$$

for a symmetric positive definite weight matrix $\mathbf{M}$ that acts as a regularizer (see [5]). If no regularization is imposed on w, for any pair of given images $\mathcal{R}$ and $\mathcal{T}$ the above model is ill-posed. Therefore, to yield a unique w, we require a regularizer $\mathcal{S}$. The following theorem proves this claim.

Theorem 1. *Consider a given z. Any two affine transformations w^A and w^B that satisfy the following conditions yield $\mathcal{L}_z(\mathcal{T}[y(w^A;x)])=\mathcal{L}_z(\mathcal{T}[y(w^B;x)])$:*

$$\begin{pmatrix} w_1^A \\ w_5^A \\ w_9^A \end{pmatrix} = \begin{pmatrix} w_1^B \\ w_5^B \\ w_9^B \end{pmatrix}, \begin{pmatrix} w_2^A \\ w_6^A \\ w_{10}^A \end{pmatrix} = \begin{pmatrix} w_2^B \\ w_6^B \\ w_{10}^B \end{pmatrix}, \begin{pmatrix} w_3^A - w_3^B \\ w_7^A - w_7^B \\ w_{11}^A - w_{11}^B \end{pmatrix} z + \begin{pmatrix} w_4^A - w_4^B \\ w_8^A - w_8^B \\ w_{12}^A - w_{12}^B \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

This suggests that if no regularization is imposed, the first two columns of w^A and w^B have to match. In addition, for any given third columns of w^A and w^B, a given z, and a given fourth column of w^A, we can always compute the fourth column of w^B that yields the same sliced result. This suggests that if we impose no regularization, the parameters of w have to be reduced to 9 instead of 12. In practice, since we typically have information about the reference w^{ref}, we impose regularization and keep the number of parameters as 12 in the parametric affine case. Furthermore, regardless of how many parameters we choose for w, the registration problem may be ill-posed in theory due to the intensities of images $\mathcal{R}$ and $\mathcal{T}$. For example, if $\mathcal{R}$ is image of a disk in 2D and $\mathcal{T}$ is image of a sphere in 3D, the problem yields infinitely many solutions since infinitely many cross-sections of a sphere can yield a disk. Due to the structure of the employed input images, this does not happen in practice. That being said, we regularize the affine transformation w in all cases.

2.2 Discretization

Here we employ a discretize-then-optimize paradigm (see the FAIR software [5] for details) to minimize the functional in Eq. (2).

Discretizing Ω into n pixels and $\mathcal{Z}$ into l pixels, we can define grids $\mathtt{x}_R = [x_k^1, x_k^2]_{k=1,\dots,n}$ and $\mathtt{x}_T = [x_j^1, x_j^2, x_j^3]_{j=1,\dots,n\times l}$ relating to $\mathcal{R}$ and $\mathcal{T}$, respectively, to be the discretizations of Ω and $\Omega \times \mathcal{Z}$. Furthermore, $\mathtt{y} \approx y(\mathtt{w}, \mathtt{x}_T)$, $\mathtt{w} = w$, the cell-centered-discretized images are $T \approx \mathcal{T}(\mathtt{x}_T)$ and $R \approx \mathcal{R}(\mathtt{x}_R)$ (containing nl and n pixels, respectively), and discretization of the operators $\mathcal{D}$ and $\mathcal{S}$ are represented by D and S (see [5]). For a given z, the discretization of the operator $\mathcal{L}_z$, denoted by L_z can be computed as

$$L_z = I_{n\times nl} := I_{n\times n} \otimes \overbrace{[0,\dots,0,\underbrace{1}_{\lceil l(z+\omega)/2\omega \rceil\text{-th component}},0,\dots,0]}^{1\times l \text{ size}} \tag{4}$$

in which we have assumed $\mathcal{Z}$ is the interval $(-\omega, \omega)$. The discretized problem is now to minimize the functional

$$J[\mathtt{w}] := D[L_z(T(\mathtt{y}(\mathtt{w}))), R] + S(\mathtt{w} - \mathtt{w}^{\text{ref}}). \tag{5}$$

2.3 Optimization

We compute the derivative and Hessian of J denoted by dJ and H_J respectively in a Gauss-Newton approach described in [6]. For simplicity, we allow ourselves

to interchangeably refer to derivatives of real-valued functions as Jacobians as well. The Hessian and Jacobian of the regularization S are denoted respectively as dS and H_S. To proceed, we represent the Jacobian of the objective function J as $dJ := \frac{\partial J}{\partial \mathtt{w}}$. Now define $L := L_z(T(\mathtt{y}(\mathtt{w})))$ and $r := L - R$. Choosing the SSD distance measure and defining $\Psi(r) := \frac{1}{2} r^T r = D[L_z(T(\mathtt{y}(\mathtt{w}))), R]$ yields $J[\mathtt{w}] = \Psi + S(\mathtt{w} - \mathtt{w}^{\mathrm{ref}})$. Hence using the chain rule

$$\begin{aligned}\frac{\partial J}{\partial \mathtt{w}} &= \left(\frac{\partial \Psi}{\partial r}\right)\left(\frac{\partial r}{\partial L}\right)\left(\frac{\partial L}{\partial T}\right)\left(\frac{\partial T}{\partial \mathtt{y}}\right)\left(\frac{\partial \mathtt{y}}{\partial \mathtt{w}}\right) + \left(\frac{\partial S}{\partial \mathtt{w}}\right) \\ &= r^T \times I_{n\times n} \times I_{n\times nl} \times dT \times d\mathtt{y} \quad + \quad dS \\ &= r^T \times I_{n\times n} \times I_{n\times nl} \times dT \times d\mathtt{y} \quad + (\mathtt{w} - \mathtt{w}^{\mathrm{ref}})^T \mathbf{M}\end{aligned}$$

in which $dT := \frac{\partial T}{\partial \mathtt{y}}$ represents the derivative of the interpolant and $d\mathtt{y} := \frac{\partial \mathtt{y}}{\partial \mathtt{w}}$ is the derivative of the transformation $\mathtt{y}$ with respect to $\mathtt{w}$. Derivatives $d\mathtt{y}$ and dT are both available in FAIR [5]. Finally, the Hessian of J denoted by H_J can be approximated as

$$H_J = d^2\Psi + H_S \approx dr^T dr + H_S = dr^T dr + \mathbf{M}, \tag{6}$$

where

$$dr = \left(\frac{\partial r}{\partial L}\right)\left(\frac{\partial L}{\partial T}\right)\left(\frac{\partial T}{\partial \mathtt{y}}\right)\left(\frac{\partial \mathtt{y}}{\partial \mathtt{w}}\right) = I_{n\times n} \times I_{n\times nl} \times dT \times d\mathtt{y} = I_{n\times nl} \times dT \times d\mathtt{y}. \tag{7}$$

In practice, to speed up the computations, matrix-free implementation of the algorithm can be applied. We also consider different discrete representations of the image registration problem, and address the discrete problems sequentially in the so-called multi-level approach.

3 Experiments and Results

3.1 Data

3D pre-procedural and 2D real-time cardiac MRI were acquired from 6 volunteers using a 1.5T MRI scanner (GE Healthcare, Waukesha, WI).

3.1.1 Prior 3D (Cine) Images

Each pre-procedural 3D volume consists of a stack of 12 to 14 short-axis (SAX) slices of the heart with a resolution of $1.37 \times 1.37 \times 8\,\mathrm{mm}^3$ and a field of view (FOV) of $350 \times 350\,\mathrm{mm}^2$. The images were acquired at end-expiration breath-hold with an electrocardiogram (ECG) triggered GE FIESTA pulse sequence and only end-diastolic images were used.

3.1.2 Real-Time Images

2D real-time images were acquired at the same slice locations as in the pre-procedural scans, but under free-breathing conditions. The images were obtained continuously with a fast spiral balanced steady state free precession (bSSFP) sequence at a frame rate of 8 fps, an in-plane resolution of 2.2×2.2 mm^2, slice thickness 8 mm, and a FOV of 350×350 mm^2. The images were ECG-gated and only images acquired end-diastole were used in the following experiments. It should be noted that stacked images do not produce meaningful volumes as there is no synchronization between different slices.

3.2 Validation of Results

If registration between two (non-identical) images is successful, a slice obtained from transforming the 3D template image with the transformation parameters obtained from registration and then slicing at a predetermined slice location would yield a 2D image similar to the 2D reference image. While a look at the end-result images can give us a subjective impression of whether registration was successful and transformation parameters returned have aligned objects in the image well, no 'ground truth' is available in general. We can, however, evaluate end-result images for their purpose in application.

The images in question are cardiac MR images, where the region of interest is the left ventricle (LV). One way of measuring how well two images have been aligned by registration is to measure how much the LVs in the template and reference images overlap before and after registration. Overlap can be quantified by the Dice coefficient and the Jaccard index, which are, respectively, defined as defined as: $\mathrm{Dice}(A, B) = \frac{2|A \cap B|}{|A|+|B|}$ and $\mathrm{Jaccard}(A, B) = \frac{|A \cap B|}{|A \cup B|}$.

The LV also contains papillary muscles which can be used as landmarks. Alignment of the landmarks can be quantified by computing the distance between corresponding landmarks in the reference and template images, before and after registration. This quantity, called the target registration error (TRE), is the l^2-normed distance between landmarks in the template image and the corresponding landmarks in the reference image.

The LV and landmarks in the cine volume were manually selected. The endocardium of the LV was outlined for each slice, and the in-plane segmentations stacked to form a 3D segmentation mask. To obtain a 2D segmentation mask of L after registration, the 3D segmentation mask is transformed using the parameters obtained from registration, and then sliced. In the real-time images, the LV and landmarks were also segmented, by an expert. The coordinates for the landmarks in the image are 2D, but knowing the location where the slice was taken from allows us to append an approximate third coordinate to the landmarks.

3.3 Cine/Cine Controlled Experiments

Before demonstrating the affine model on registration between a cine (pre-operative) image and a real-time (intra-operative) image, we first perform

controlled experiments. Controlled registration experiments were performed between a 3D cine volume (template) and a 2D image (reference) that is a slice of a transformed version of the 3D volume. Since that initial transformation is known, ground truth is available. For all experiments following, domains $\Omega = (-175, 175) \times (-175, 175)\,\text{mm}^2$ and $\mathcal{Z} = (-48, 48)\,\text{mm}$, and discretizations $n = 128^2 = 16384$ and $l = 12$. For affine registration, we will also assume the regularizer $\mathbf{M}$ is a diagonal 12×12 matrix with unit entries on the main diagonal except for locations 3, 7, and 11 (third column of matrix) where entries are 10^6, i.e., large. If $\mathtt{w}^{\text{ref}}$ is chosen to be the identity transformation, the regularizer ensures the computed parameters $[\mathtt{w}_3, \mathtt{w}_7, \mathtt{w}_{11}]$ to be close to $[0, 0, 1]$; see Theorem 1. Linear interpolation and an Armijo line search [6] were used in the multi-level Gauss-Newton optimization framework.

3.3.1 Affine Initial Transformation

If the initial transformation applied to generate the reference image is affine, successful affine registration should produce a transformed template slice that aligns with the reference image. Recall that the motivation behind using an affine model as opposed to a rigid was to more accurately represent the deformable nature of organs in the body. To demonstrate that the rigid model does indeed fail when the nature of the deformation applied to the reference image R is not rigid, we individually perturbed each entry of the identity transform $\mathtt{w} = [1, 0, 0, 0, 0, 1, 0, 0, 0, 0, 1, 0]$, and applied the perturbed set of parameters in the initial transformation to obtain R. For each R that was obtained, rigid and affine registration was performed. Due to Theorem 1, perturbing entries w_3, w_7, and w_{11} is equivalent (in terms of producing the same template image slice) to perturbing w_4, w_8, and w_{12} but scaled by a factor of z, the location of the slicing operator, so only 9 entries of $\mathtt{w}$ need to be perturbed; w_4, w_8, and w_{12} were not perturbed. For each of the 9 entries, an ϵ between -0.5 to 0.5 was added to the entry to produce a set of initial parameters used to obtain R.

Figures 1, 2, and 3 show the effects of perturbing the entries of w_2 on the TRE before registration, after rigid registration, and after affine registration. As expected, affine registration improves results over rigid registration. Similar results were found for entries w_1, w_5, w_6, w_9, and w_{10} (first two columns of matrix), but for the sake of brevity, no figures will be shown for those entries. In all figures, the box represents the 25th–75th percentile, and the line in the box marks the median. Figures 4, 5, and 6 show the effects of perturbing the entries of w_3 on the TRE before registration, after rigid registration, and after affine registration. Due to Theorem 1, perturbations in w_3, w_7, and w_{11} can be compensated for when the image is reduced from 3D to 2D by changing the values of w_4, w_8, and w_{12}, which are translation parameters and thus forms a rigid transformation, if no other shear terms are present. Rigid registration was therefore comparable to affine transformation for perturbations on w_3, as seen in Figs. 5 and 6. The same was found for entries w_7 and w_{11} (third column of matrix), but for the sake of brevity, no figures will be shown for those entries.

Recall that an indicator of good alignment is a simultaneously large Jaccard index and small TRE. Although not shown, LV overlap was quantified with the Jaccard index as well. For perturbations on w_2 (and $w_1, w_5, w_6, w_9, w_{10}$ as well), affine registration generally performed well, increasing Jaccard indices and reducing TREs. Rigid registration did not improve results; Jaccard indices after rigid registration became more varied and generally appear to worsen. For perturbations on w_3 (and w_7, w_{11}), the results of rigid registration were comparable to affine registration. This can be explained by Theorem 1 – variations in w_3, w_7, and w_{11} can be compensated for by changing the values of w_4, w_8, and w_{12} to obtain the same 2D slice of a 3D volume.

3.4 Real-Time/Cine Experiments

No initial 3D transformation was applied to obtain the reference image as was done on the test cases, since the reference images here are 2D real-time images. It is also not meaningful to perform a 2D transformation on a real-time image to obtain the reference image for registration, since a modified image no longer represents a clinical setting. Because the slice locations in the real-time and cine cardiac MRI are already rather aligned initially in the z-direction, registration between images from same the slice prescription would align things mostly within the xy-plane, and give little indication of how well the algorithm works when the images are taken at different slice locations. Performing registration between different slices would be a better indicator of how well the algorithm improves alignment in the z-direction. For the following example (Fig. 7), the real-time slice was taken at spatial location $z = -4\,\text{mm}$ while the slicing operation was applied on the template image at $z = -36$ mm, so the initial slice of the 3D template is at $z = -36\,\text{mm}$. To register the images successfully, the registration algorithm must return transformation parameters that translate the template image by approximately 32 mm (the physical distance between the spatial locations of the reference image and slicing operator) in the z-direction, along with appropriate alignments in the x- and y-directions. Figure 7 and Table 1 show the results of one experiment. The affine model appears to produce slightly better results for this experiment, due to its ability to deform, apparent in the LV overlap after registration (Fig. 7g). In most clinical applications, initial misalignment will not be as large and the two images registered will be slices in close proximity to one another. Affine and rigid registration was performed on real-time images from 6 data sets, each contributing 1 cine image and between 17 to 29 real-time images, to a total of 143 real-time images across 6 data sets. Each real-time image was registered to the cine image of the same subject at the same slice location and cardiac phase. Although the slice prescriptions are identical, there may be small motion normal to the image plane. The results are listed in Table 2. With the exception of Data Set 3 and Data Set 5, rigid registration improves or leaves results unchanged. Affine registration improves results for all data sets except Data Set 3. For Data Set 3, rigid registration returned values worse than what was initially given and affine registration performed even worse. This was due to local deformation within the cardiac region, consistent throughout the

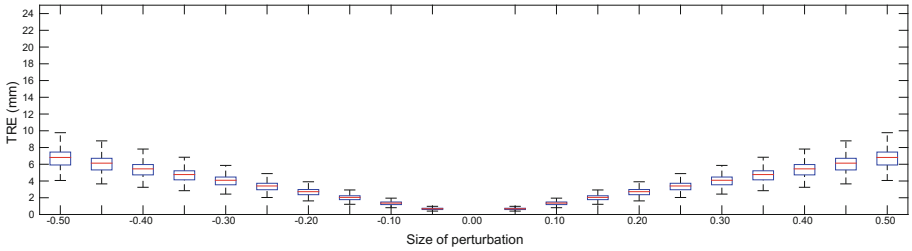

Fig. 1. The TRE as a function of the perturbation on w_2 before registration for all data sets. Reference image obtained by an affine transformation that is the identity transformation except for the addition of the perturbation to w_2.

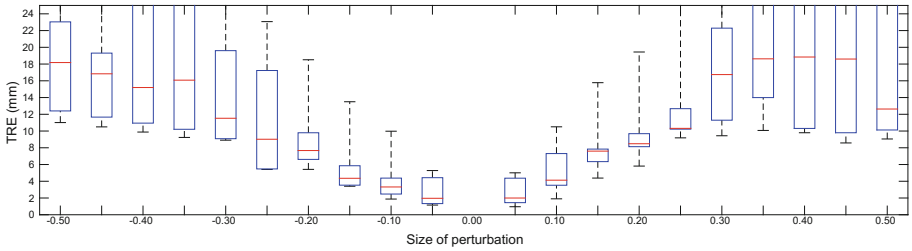

Fig. 2. The TRE as a function of the perturbation on w_2 after rigid registration for all data sets. Reference image obtained by an affine transformation that is the identity transformation except for the addition of the perturbation to w_2.

data available for Data Set 3. Since the body cavity is considerably larger than the cardiac region and comprises most of the content in each image, the algorithm accounted for the body cavity, not the heart, thus the LV becomes more misaligned after registration. For Data Set 5, rigid registration returned slightly worse values than what the algorithm initially started with, but affine registration produced values that were a slight improvement over the initial data. From the values for the rest of the data sets, however, affine registration returns better results in general compared to rigid registration. The TRE was not calculated in this set of experiments because the images are from the same slice locations, image resolution in the z-direction (the direction normal to the short-axis slices) is much coarser than the in-plane resolution, i.e. slice thickness is larger than pixel size; because there is no ground truth available for us to obtain more precise landmark locations, thus z-direction uncertainty would dominate and render the results meaningless.

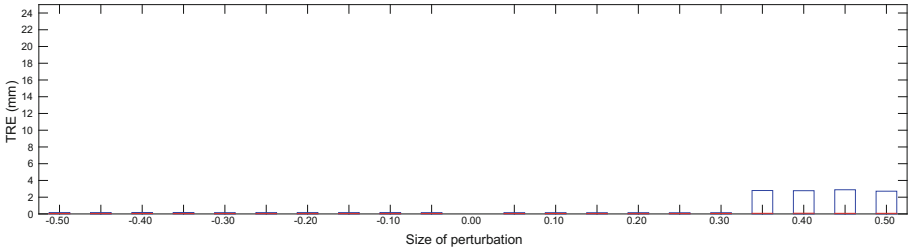

Fig. 3. The TRE as a function of the perturbation on w_2 after affine registration for all data sets. Reference image obtained by an affine transformation that is the identity transformation except for the addition of the perturbation to w_2.

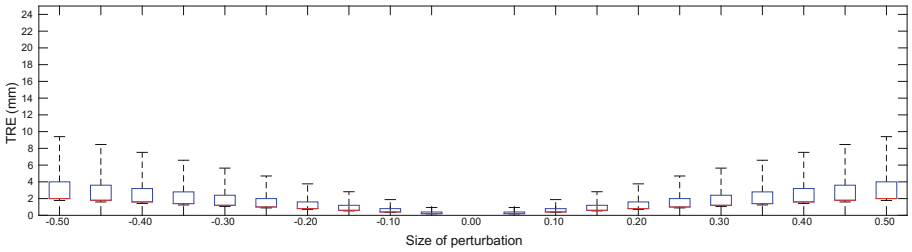

Fig. 4. The TRE as a function of the perturbation on w_3 before registration for all data sets. Reference image obtained by an affine transformation that is the identity transformation except for the addition of the perturbation to w_3.

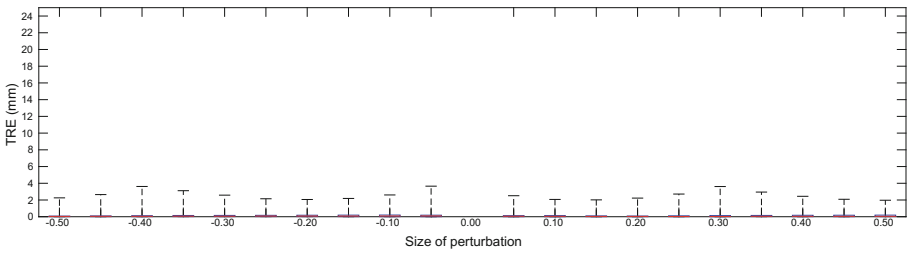

Fig. 5. The TRE as a function of the perturbation on w_3 after rigid registration for all data sets. Reference image obtained by an affine transformation that is the identity transformation except for the addition of the perturbation to w_3.

Table 1. Jaccard indices and Dice coefficients of left ventricle overlap before and after registration between a 3D cine volume and a real-time image taken from $z = -4$ mm. Slicing operations performed at $z = -36$ mm.

	Jaccard	Dice	TRE (mm)
Before registration	0.67	0.80	32.8 ± 0.2
After rigid registration	0.71	0.83	6.6 ± 0.8
After affine registration	0.87	0.93	4.5 ± 0.1

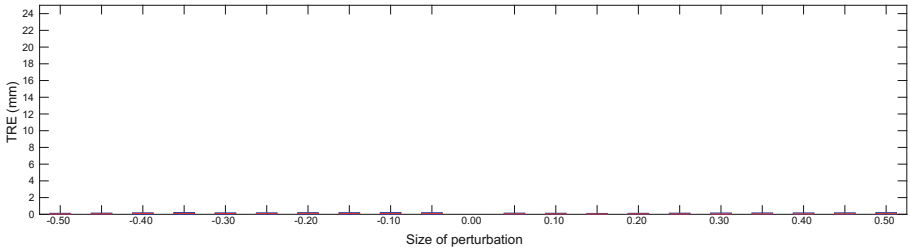

Fig. 6. The TRE as a function of the perturbation on w_3 after affine registration for all data sets. Reference image obtained by an affine transformation that is the identity transformation except for the addition of the perturbation to w_3.

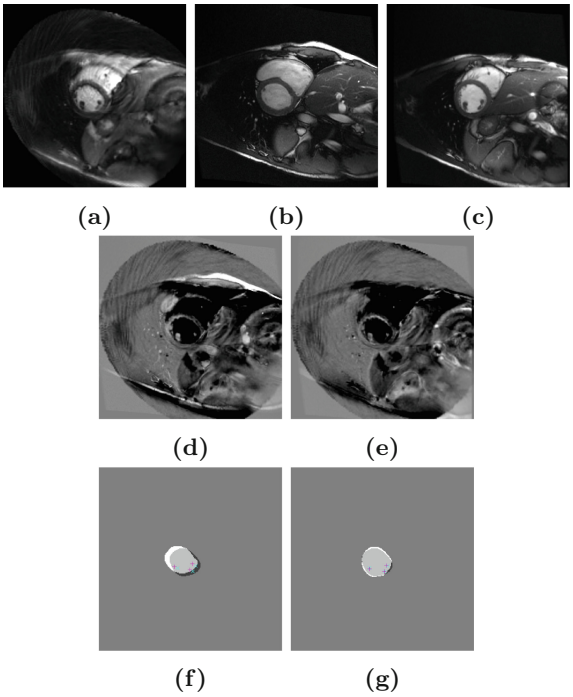

Fig. 7. Results of affine registration between a 3D cine image and a 2D real-time image on the same subject as in the controlled experiment, with an initial misalignment of approximately 32 mm in the z-direction (through the image plane). (a) Reference image R. (b), (c) Template slice L before and after registration. (d), (e) Difference between the reference image and template slice ($L - R$) before and after registration. (f), (g) Segmentation masks showing left ventricle overlap before and after registration, with in-plane reference image landmarks ($\times$) and out-of-plane template image landmarks projected onto image ($+$).

Table 2. LV overlap before registration, after rigid registration, and after affine registration between a pre-operative 3D cine volume and a noisier, lower-resolution intra-operative 2D real-time image, as in a clinical setting.

	Data set	Before registration	After rigid registration	After affine registration
Jaccard	1	0.86 ± 0.06	0.87 ± 0.07	0.92 ± 0.02
Dice		0.92 ± 0.04	0.93 ± 0.04	0.96 ± 0.01
Jaccard	2	0.75 ± 0.02	0.86 ± 0.02	0.87 ± 0.02
Dice		0.86 ± 0.02	0.93 ± 0.01	0.93 ± 0.01
Jaccard	3	0.77 ± 0.06	0.47 ± 0.13	0.21 ± 0.08
Dice		0.87 ± 0.04	0.63 ± 0.13	0.34 ± 0.12
Jaccard	4	0.49 ± 0.08	0.66 ± 0.05	0.73 ± 0.08
Dice		0.65 ± 0.07	0.80 ± 0.04	0.84 ± 0.06
Jaccard	5	0.80 ± 0.04	0.73 ± 0.03	0.83 ± 0.04
Dice		0.89 ± 0.03	0.84 ± 0.02	0.91 ± 0.02
Jaccard	6	0.76 ± 0.09	0.76 ± 0.09	0.80 ± 0.09
Dice		0.86 ± 0.06	0.86 ± 0.06	0.88 ± 0.06

4 Discussions and Conclusions

In controlled experiments where the reference image is a transformed and sliced version of the template image, it was demonstrated that rigid registration did not sufficiently account for deformations that are affine in nature. We can conclude that the affine model performs better than, or is at least comparable to, the rigid model for controlled experiments, but at the expense of extra computational time. In registration between real-time images and cine images of the same slice location and cardiac phase, affine registration generally performed better than rigid registration, presumably due to its greater flexibility over the rigid model, allowing it to deform the cine image to more closely match the real-time image. We can conclude that between images of the same modality, the proposed multi-level parametric 2D-3D registration scheme can align images well for misalignments within reasonable limits encountered in clinical applications, such as motion due to respiration. Despite different acquisition methods in the real-time and prior cine MR images, the registration algorithm improved alignment with the SSD dissimilarity measure.

Affine registration was found to be a generally more robust model than rigid registration in this framework. This suggests that in attempting to improve results for applications employing 2D-3D rigid registration with the SSD, one can first consider simply expanding the transformation model to an affine one before considering more complex dissimilarity measures and regularizers. The advantage of the affine model is its simplicity, allowing more accurate registration at a small cost. For multi-modality registration where intensities of the

template and reference images differ more drastically, one can consider using other dissimilarity measures and/or optimizers [5,7] that can fit well within the context of the general proposed model.

Acknowledgments. This research was supported in part by a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant for M. Ebrahimi. We would like to thank Drs. Graham Wright and Robert Xu of Sunnybrook Research Institute, Toronto, Canada, for valuable discussions and providing the MR data.

References

1. Pushparajah, K., Tzifa, A., Razavi, R.: Cardiac MRI catheterization: a 10-year single institution experience and review. Intervent. Cardiol. **6**, 335–346 (2014)
2. Helen, X., Lasso, A., Fedorov, A., Tuncali, K., Tempany, C., Fichtinger, G.: Multi-slice-to-volume registration for MRI-guided transperineal prostate biopsy. Int. J. CARS **10**(1), 563–572 (2015)
3. Smolíková, R., Wachowiak, M.P., Drangova, M.: Registration of fast cine cardiac MR slices to 3D preprocedural images: toward real-time registration for MRI-guided procedures. In: Proceedings of SPIE, vol. 5370, pp. 1195–1205, May 2004
4. Xu, R., Wright, G.A.: Registration of real-time and prior imaging data with applications to MR guided cardiac interventions. In: Camara, O., Mansi, T., Pop, M., Rhode, K., Sermesant, M., Young, A. (eds.) STACOM 2014. LNCS, vol. 8896, pp. 265–274. Springer, Cham (2015). doi:10.1007/978-3-319-14678-2_28
5. Modersitzki, J.: FAIR: Flexible Algorithms for Image Registration. SIAM, Philadelphia (2009)
6. Nocedal, J., Wright, S.J.: Numerical Optimization, 2nd edn. Springer, New York (2006)
7. Ardeshir Goshtasby, A.: 2-D and 3-D Image Registration. Wiley Press, New York (2005)