# Multimodal Registration via Mutual Information Incorporating Geometric and Spatial Context

Jonghye Woo, Member, IEEE, Maureen Stone, and Jerry L. Prince, Fellow, IEEE

Abstract-Multimodal image registration is a class of algorithms to find correspondence from different modalities. Since different modalities do not exhibit the same characteristics, finding accurate correspondence still remains a challenge. To deal with this, mutual information (MI)-based registration has been a preferred choice as MI is based on the statistical relationship between both volumes to be registered. However, MI has some limitations. First, MI-based registration often fails when there are local intensity variations in the volumes. Second, MI only considers the statistical intensity relationships between both volumes and ignores the spatial and geometric information about the voxel. In this work, we propose to address these limitations by incorporating spatial and geometric information via a 3D Harris operator. In particular, we focus on the registration between a high-resolution image and a low-resolution image. The MI cost function is computed in the regions where there are large spatial variations such as corner or edge. In addition, the MI cost function is augmented with geometric information derived from the 3D Harris operator applied to the high-resolution image. The robustness and accuracy of the proposed method were demonstrated using experiments on synthetic and clinical data including the brain and the tongue. The proposed method provided accurate registration and yielded better performance over standard registration methods.

Index Terms—Multimodal image registration, mutual information, Harris operator.

## I. INTRODUCTION

ULTIMODAL image registration is a basic yet M important operation for many applications. Clinical research studies often involve large numbers of volumes taken from multiple modalities. Different modalities provide diverse features, which can be used in a variety of theoretical and practical applications, such as relating function to anatomy or image guided surgery [1], [2].

Multimodal image registration still remains a challenge partly because different modalities to be registered exhibit

Manuscript received May 16, 2014; revised September 30, 2014; accepted December 10, 2014. Date of publication December 31, 2014; date of current version January 13, 2015. This work was supported by the National Institutes of Health under Grant R01CA133015 and Grant R00DC012575. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Quanzheng Li.

J. Woo is with the Department of Radiology, Massachusetts General Hospital, Boston, MA 02114 USA, and also with the Harvard Medical School, Boston, MA 02115 USA (e-mail: jwoo@mgh.harvard.edu).

M. Stone is with the Department of Neural and Pain Sciences, University of Maryland at Baltimore, Baltimore, MD 21250 USA (e-mail: stone@umaryland.edu).

J. L. Prince is with the Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD 21218 USA (e-mail: prince@jhu.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIP.2014.2387019

Fig. 1. Examples of multimodal images: (a) an MR-T2 image of the brain,

(b) a high-resolution MRI (hMRI) of the tongue that was acquired at rest, (c) a PET image of the brain, and (d) the first time frame of cine-MRI of the tongue. These images differ considerably in spatial resolution.

different characteristics. For instance, the widely used modalities, structural imaging, such as Computed Tomography (CT), and high-resolution magnetic resonance imaging (MRI) (see Fig. 1(a), (b)) depict detailed anatomical information. Functional or dynamic imaging, such as Positron Emission Tomography (PET) (see Fig. 1(c)), Single Photon Emission Computed Tomography (SPECT), functional MRI (fMRI) contain information about changes in blood flow, metabolism, and regional chemical composition. Cine-MRI (see Fig. 1(d)) provides surface motion of anatomical structures. The benefits of using multimodal data have created a need for the development of highly accurate and robust multimodal image registration. In particular, multimodal image registration of high-resolution images (e.g., high-resolution MRI) with low-resolution images (e.g., PET) is often needed. Such registration is the main interest of this paper, as our method specifically exploits geometric information in the higher resolution image in the computation of the similarity measure.

Multimodal image registration is a class of algorithms to find correspondences between multiple datasets from the same subject, acquired using different imaging modalities [3]. The task of aligning two images is cast as an

1057-7149 © 2014 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications\_standards/publications/rights/index.html for more information.





Fig. 2. When the two circles in (a) are to be aligned, the conventional MI objective function in (b) fails to show a peak at the correct point of alignment, whereas the proposed method in (c) has the desired peak. Note that the values in (b) and (c) are plotted with respect to the x and y translation of the dark circle (top (a)) against the gray circle (bottom (a)).

optimization problem: a common approach to registration is to deform one of the images so as to maximize its similarity to the other image while maintaining a "smoothness" in the estimated deformation field. Most of the similarity measures can be classified into two categories [4]: feature-based and intensity-based. The former usually requires four steps: (1) feature detection, (2) feature matching, (3) transform model estimation, and (4) resampling. In general, feature-based methods are computationally efficient compared to intensity-based methods; but manual intervention is often required to improve accuracy and efficiency.

Intensity-based methods have been shown to be more accurate than feature-based methods [5], but special assumptions on the intensities of the images are often required in order to achieve successful registration. For instance, while unimodal image registration is based on the assumption that corresponding pixels have similar intensity values, the same assumption does not hold in multimodal registration problems. This is because the two modalities may assign different intensities to the same structure. Mutual information (MI), therefore, has become the established intensity similarity measure in multimodal registration because it accommodates different intensities between the modalities provided that they are relatively consistent within each modality [6], [7]. MI has been extensively used for many applications including diagnosis [8], surgical planning [9], and radiation therapy [10].

Although MI is considered to be the gold standard similarity measure for multimodal image registration, there are two problems with the traditional method. First, the performance of MI degrades when there are local intensity variations; this happens because the joint histogram computation is adversely affected [11], [12]. Second, the conventional similarity measure only incorporates intensity information, which means that the spatial information that may provide additional cues about the optimal registration is entirely ignored [13]. Fig. 2 illustrates the problem with local intensity variations and demonstrates the improvement provided by our proposed approach.

Several previous methods incorporated spatial information in the computation of MI. Pluim et al. [13] combined spatial information by multiplying the MI with an external local gradient, and both gradient magnitude and orientation were incorporated into the calculation of MI. Rueckert et al. [14] proposed second-order MI to the problem of 2D registration by using a 4D joint histogram that considered the six nearest neighbors

of each voxel to calculate MI. However, Rueckert et al. [14] required a large number of samples to compute the highdimensional histogram, and therefore is not easily extended to 3D for computational reasons. As an extension of the secondorder MI, Russakoff et al. [15] performed regional MI taking into account both corresponding pixels and their neighbors. This method also used a high-dimensional histogram, which may not be reliable when the number of samples is small. Yi et al. [16] proposed the inclusion of spatial variability via a weighted combination of normalized mutual information (NMI) and local matching statistics. Loeckx et al. [17] introduced conditional MI, which incorporates both intensity and spatial dimensions to express the location of the joint intensity pair. Zhuang et al. [18] proposed to unify spatial information into the computation of the joint histogram. This method used a hierarchical weighting scheme to differentiate the contribution of sample points to a set of entropy measures, which are associated to spatial variable values. In related developments, Myronenko et al. [19] presented a novel similarity measure, the residual complexity, that accounts for complex spatially-varying intensity distortions in mono-modal settings. However, this approach may not work well in multi-modal settings.

Our work shares the spirit of these past works in the sense that we include spatial information in the computation of MI, thereby incorporating both spatial and geometric information. However, one of the main contributions of this work is that our method involves a new approach that specifically exploits the higher resolution image. Specifically, we compute a structure matrix at each voxel and use the 3D Harris operator to decompose the image into three disjoint and geometrically distinct regions. These classes were then used to determine the relative contribution of each voxel's intensity in the computation of the joint histogram. In this way, geometric structure from the highresolution image influences the matching computation between the two images. A preliminary version [20] of this method was reported. Here we report the completed algorithm design and provide new validations on both synthetic and in vivo data including the tongue and the brain.

The remainder of this paper is organized as follows. Section II provides a background about maximization of MI and Harris corner detector. The proposed registration method with the 3D Harris operator is described in Sec. III, followed by experimental results presented in Sec. IV. Finally, a discussion and concluding remarks are given in Secs. V and VI, respectively.

# **II. PRELIMINARIES**

## A. Maximization of Mutual Information

In this section, we describe the maximization of MI for multimodal image registration. We first define terms and notation used in this work. The images  $I_1 : \Omega_1 \subset \mathbb{R}^3 \to \mathbb{R}$ and  $I_2 : \Omega_2 \subset \mathbb{R}^3 \to \mathbb{R}$ , defined on the open and bounded domains  $\Omega_1$  and  $\Omega_2$ , are the template and target images, respectively. Given two images, a deformation field is defined by the mapping  $u(\mathbf{x}; \boldsymbol{\mu}) : \Omega_2 \mapsto \Omega_1$ . In our work, u is a B-spline transformation with associated parameters  $\boldsymbol{\mu} = (\mu_1, \mu_2, \mu_3, \cdots)$ , containing the B-spline coefficients. The goal of registration is to find a deformation field at each pixel location **x** such that the deformed template  $I_1(u(\mathbf{x}))$ aligns the underlying anatomy as closely as possible with  $I_2(\mathbf{x})$  satisfying the given criterion. Since  $I_1$  and  $I_2$  are considered to be different modalities, we focus on the MI criterion. The main idea is to find the deformation field  $\hat{u}$ that maximizes the mutual information  $\mathcal{M}$  contained in the distribution of paired image intensities of the aligned images. Accordingly,

$$\hat{u} = \arg\max_{u} (\mathcal{M}(I_1(u(\mathbf{x})), I_2(\mathbf{x}))), \qquad (1)$$

where

$$\mathcal{M}(I_1(u(\mathbf{x})), I_2(\mathbf{x})) = \iint p_u(i_1, i_2) \log \frac{p_u(i_1, i_2)}{p_{I_1}(i_1)p_{I_2}(i_2)} di_1 di_2.$$
(2)

Here,  $i_1$  and  $i_2$  are the image intensity values in  $I_1(u(\mathbf{x}))$ and  $I_2(\mathbf{x})$ , respectively, and  $p_{I_1}(i_1)$  and  $p_{I_2}(i_2)$  are their marginal probability distributions while  $p_u(i_1, i_2)$  is their joint probability distribution. All densities are computed using a Parzen window approximation [21] and the joint density is computed only over the overlap region.

## B. Harris Corner Detector

The Harris corner detector [22] was first introduced to detect corner features that contain high intensity changes in both the horizontal and vertical directions. The Harris corner detector is a well established technique using linear filtering of an image. Given an image I, the autocorrelation matrix of the point l = (x, y) in the neighborhood N of l is given by

$$\mathcal{N}(l) = \begin{pmatrix} \sum_{m \in N} I_x^2(m) & \sum_{m \in N} I_x(m)I_y(m) \\ \sum_{m \in N} I_x(m)I_y(m) & \sum_{m \in N} I_y^2(m), \end{pmatrix}$$
(3)

where  $I_x$  and  $I_y$  denote the partial derivatives of I in the x and y directions, respectively.

The Harris corner indicator  $H_2$  is then given by

$$H_2 = \det(\mathcal{N}) - r(\operatorname{trace}(\mathcal{N}))^2, \qquad (4)$$

where r is an arbitrary constant.

#### **III. PROPOSED APPROACH**

In this section, we describe our proposed method. Our method is based on an iterative framework of computing MI incorporating spatial information and geometric cues. The underlying idea is to split the image into a set of non-overlapping regions using the 3D Harris operator derived from the higher resolution image and to perform registration on spatially meaningful regions. Additionally, we exploit structural information describing the gradient of the local neighborhood of each voxel to define structural similarity for MI computation.

## A. Volume Labeling Using 3D Harris Operator

In this work, we extend the 2D Harris detector to three dimensions so that it can be used to define regions over which MI is more heavily weighted. The Harris operator is derived from the local autocorrelation function of the intensity. The autocorrelation function at a point (x, y, z) is defined as

$$f(x, y, z) = \sum_{p} \sum_{q} \sum_{r} w(p, q, r) \times [I(p+x, q+y, r+z) - I(p, q, r)]^{2}, \quad (5)$$

where  $I(\cdot, \cdot, \cdot)$  is the 3D image, (p, q, r) denote a neighborhood of (x, y, z) in the Gaussian function  $w(\cdot, \cdot, \cdot)$  centered on (x, y, z). Using a first-order Taylor expansion, I(p + x, q + y, r + z) is approximated by

$$I(p + x, q + y, r + z) \approx I(p, q, r) + xI_x(p, q, r) + yI_y(p, q, r) + zI_z(p, q, r).$$
(6)

f(x, y, z) can then be given by

$$f(x, y, z) = \sum_{p} \sum_{q} \sum_{r} w(p, q, r)$$

$$\times [x I_{x}(p, q, r) + y I_{y}(p, q, r) + z I_{z}(p, q, r)]^{2} \quad (7)$$

$$f(x, y, z) \approx \sum \sum \sum w(p, q, r)$$

$$\times \left\{ \begin{bmatrix} I_x(p,q,r)I_y(p,q,r)I_z(p,q,r) \begin{bmatrix} x \\ y \\ z \end{bmatrix} \right\}^2$$
$$= [x \ y \ z] \mathcal{C}(x,y,z) \begin{bmatrix} x \\ y \\ z \end{bmatrix}.$$
(8)

Here the local structure matrix C(x, y, z), which captures the intensity structure of the local neighborhood, is defined as

$$C(x, y, z) = \sum_{p} \sum_{q} \sum_{r} w(p, q, r) \begin{pmatrix} I_{x}^{2} & I_{x}I_{y} & I_{x}I_{z} \\ I_{x}I_{y} & I_{y}^{2} & I_{y}I_{z} \\ I_{x}I_{z} & I_{y}I_{z} & I_{z}^{2} \end{pmatrix}$$
(9)

where  $I_x$ ,  $I_y$ , and  $I_z$  denote the partial derivatives of I in the x, y, and z directions, respectively. In analogy to the 2D Harris operator [22], we define the 3D Harris operator as

$$H_3 = \det(\mathcal{C}) - k(\operatorname{trace}(\mathcal{C}))^3, \tag{10}$$

where k is an arbitrary constant. Each voxel can then be classified as one of three types using a threshold T and the following definitions

- Type 1:  $H_3 \ge T$ , voxel has significant local variation
- Type 2:  $H_3 \leq -T$ , voxel has moderate local variation
- Type 3:  $-T \le H_3 \le T$ , voxel has small local variation

We assume that Type 1 and Type 2 regions have more structural and characteristic information compared to Type 3 (homogeneous) region to calculate local statistics. Thus we use Type 1 and Type 2 regions to calculate MI. One example result of the voxel labeling is shown in Fig. 3. It is worth noting that we perform the 3D Harris operator on higher resolution volumes such as hMRI only because the higher resolution image has more detailed anatomical information.



Fig. 3. Examples of voxel labeling using the 3D Harris operator. (a) An axial image of hMRI. In (b)–(d), white voxels represent (b) Type 1 - high variability, (c) Type 2 - moderate variability, and (d) Type 3 - small variability.

# B. Mutual Information Using the Local Structure Matrix

As shown in Eq. (2), MI provides a measure of image similarity using only the marginal and joint probability densities of the image intensities. Here, we incorporate spatial and geometric information into the calculation of MI by computing a weighted joint distribution similar to the work by Luan et al. [23] as follows

$$p_{u}^{\mathcal{C}}(i_{1}, i_{2}) = \frac{1}{|V|} \int_{V} \gamma(\mathbf{x}) \varphi\left(\frac{i_{1} - I_{1}(u(\mathbf{x}))}{\rho}\right) \varphi\left(\frac{i_{2} - I_{2}(\mathbf{x})}{\rho}\right) d\mathbf{x},$$
(11)

where  $\varphi$  is a Gaussian kernel, the overlap region  $V = \Omega_2 \cap u^{-1}(\Omega_1)$ ,  $\rho$  controls the width of window, and  $\gamma$  (**x**) is a weighting function that is large when local structure matrices are similar and is small otherwise. Specifically, we define this weighting function as

$$\gamma(\mathbf{x}) = \exp(-\frac{\Delta(\mathcal{C}_{i_1}(\mathbf{x}), \mathcal{C}_{i_2}(\mathbf{x}))}{m}), \qquad (12)$$

where  $\Delta(C_{i_1}(\mathbf{x}), C_{i_2}(\mathbf{x}))$  is a distance between two matrices, *m* is a normalization constant, and  $C_{i_1}(\mathbf{x})$  and  $C_{i_2}(\mathbf{x})$  are the local structure matrices of the corresponding pixels in  $I_1(u(\mathbf{x}))$  and  $I_2(\mathbf{x})$ , respectively.

Since local structure matrices do not live in a vector space, the distance function  $\Delta$  cannot be defined by a conventional Euclidean metric. The matrices are symmetric and positive semidefinite (like covariance matrices), however, and therefore belong to a connected Riemannian manifold that is locally Euclidean [24]. Accordingly, we can define the distance between two structure matrices as

$$\Delta(\mathcal{C}_{i_1}(\mathbf{x}), \mathcal{C}_{i_2}(\mathbf{x})) = \sqrt{\sum_{n=1}^{N} (\ln \lambda_n(\mathcal{C}_{i_1}(\mathbf{x}), \mathcal{C}_{i_2}(\mathbf{x})))^2}, \quad (13)$$

where N is the number of rows and columns in each matrix and  $\lambda_n$  are the generalized eigenvalues of  $C_{i_1}(\mathbf{x})$  and  $C_{i_2}(\mathbf{x})$ defined by

$$\lambda_n \mathcal{C}_{i_1} x_n = \mathcal{C}_{i_2} x_n, \quad n = 1, \dots, N,$$
(14)

where  $x_n$ , n = 1, ..., N, are the corresponding generalized eigenvectors. This definition of distance satisfies the metric



Fig. 4. Distance calculated between structure matrices in hMRI and cine-MRI. (a) 2D slice of cine (left) and high-resolution (right) MR images. (b) First order derivatives with respect to x and y axes of 2D cine (top) and high-resolution (bottom) MR slices. (c) Distance between structure matrices defined in Eq. (13).

properties including symmetry, positivity, and the triangle inequality. An example of this distance, computed between aligned 2D cine-MRI and hMRI images, is shown in Fig. 4(c). This shows that the distance metric is small when there are regions of high geometric similarity, such as edges.

We write a modified MI criterion using the above weighting scheme as follows

$$\mathcal{M}^{\mathcal{C}}(I_{1}(u(\mathbf{x})), I_{2}(\mathbf{x})) = \iint p_{u}^{\mathcal{C}}(i_{1}, i_{2}) \log \frac{p_{u}^{\mathcal{C}}(i_{1}, i_{2})}{\hat{p}_{I_{1}}(i_{1})\hat{p}_{I_{2}}(i_{2})} di_{1} di_{2}, \quad (15)$$

where  $\hat{p}_{I_1}(i_1)$  and  $\hat{p}_{I_2}(i_2)$  are recomputed based on  $p_u^C(i_1, i_2)$  as in [25]. Using this modified MI, the local structure matrices provide a geometric similarity measure while the image intensities continue to provide an appearance measure, thereby allowing us to find correspondence more reliably and address the limitation of the traditional MI-based registration. In summary, our registration approach seeks to maximize the image similarity given by

$$\mathcal{D}(I_1(u(\mathbf{x})), I_2(\mathbf{x})) = \mathcal{M}^{\mathcal{C}}(I_1(u(\mathbf{x})), I_2(\mathbf{x})), \quad \mathbf{x} \in M(\mathbf{x}),$$
(16)

where  $M(\mathbf{x})$  is a mask comprising only the Type 1 and Type 2 voxels in the higher resolution image.

## C. Transformation Model

Having defined an image similarity, we must now define a model for the spatial transformation that we will use. Here we follow a standard model [26],

$$h(\mathbf{x}) = h_1(\mathbf{x}) + h_2(\mathbf{x}),\tag{17}$$

where  $h_1(\mathbf{x})$  and  $h_2(\mathbf{x})$  are affine and deformable registrations, respectively. Affine registration accounts for large displacements including rotations, translations, and scalings. Deformable registration accounts for the finer details at a local level. We use free-form deformations (FFD) based on uniform cubic B-splines [26] to implement deformable registration. Therefore  $h_2(\mathbf{x})$  can be written as the 3D tensor product of 1D cubic B-splines,

$$h_2(\mathbf{x}) = \sum_{l=0}^{3} \sum_{m=0}^{3} \sum_{n=0}^{3} B_l(u) B_m(v) B_n(w) c_{i+l,j+m,k+n},$$
(18)

where c denotes the control points and i, j, and k represent the index of the control point.

The B-spline transformation model has three desirable properties for the present application. First, estimated deformation field is easily regularized by controlling the control point separation [26]. We use this property to balance accuracy versus smoothness of the resulting deformation field. Second, B-splines are separable in multiple dimensions, providing computational efficiency. We refer the reader to [25] and [27] for implementation details. Finally, optimization is inherently local, since changing the location of a single control point modifies only a local neighborhood of control point [28]. This permits additional computational efficiency since regions that have converged do not need further updating.

# D. Optimization

The energy functional is minimized using a gradient descent method [29] defined by

$$\boldsymbol{\mu}_{k+1} = \boldsymbol{\mu}_k - a_k g(\boldsymbol{\mu}_k), \tag{19}$$

where  $g(\mu_k)$  is the derivative of the cost function evaluated at the current position  $\mu_k$ . From our weighted joint distribution, we first re-compute the marginal probability distribution,  $\hat{p}_{I_1}(i_1)$ . In calculating the gradient of the modified MI, we follow the work [25], [27]. We first derive a Taylor expansion of the modified MI as given by

$$\mathcal{M}^{\mathcal{C}}(\boldsymbol{\mu}) = \mathcal{M}^{\mathcal{C}}(\boldsymbol{\nu}) + \sum_{i} \frac{\partial \mathcal{M}^{\mathcal{C}}(\boldsymbol{\nu})}{\partial \mu_{i}} (\mu_{i} - \nu_{i}) + \frac{1}{2} \sum_{i,j} \frac{\partial^{2} \mathcal{M}^{\mathcal{C}}(\boldsymbol{\nu})}{\partial \mu_{i} \partial \mu_{j}} \partial \mu_{j} (\mu_{i} - \nu_{i}) (\mu_{j} - \nu_{j}) + \cdots$$
(20)

We simplify the above equarray by ignoring all terms above second-order as in [25]. Here the quadratic model yields the optimal solution when the Parzen window is a B-spline of degree  $m \ge 3$  [25].

We then calculate the gradient of the cost function, which is necessary for its efficient minimization. The gradient of the modified MI with respect to the parameter  $\mu$  is given by

$$\nabla \mathcal{M}^{\mathcal{C}} = \left[\frac{\partial \mathcal{M}^{\mathcal{C}}}{\partial \mu_1}, \frac{\partial \mathcal{M}^{\mathcal{C}}}{\partial \mu_2}, \cdots\right].$$
 (21)

A single component of the gradient is found by

$$\frac{\partial \mathcal{M}^{\mathcal{C}}}{\partial \mu} = -\iint \frac{\partial p_{u}^{\mathcal{C}}(i_{1}, i_{2})}{\partial \mu} \log \frac{p_{u}^{\mathcal{C}}(i_{1}, i_{2})}{\hat{p}_{I_{1}}(i_{1})} di_{1} di_{2}.$$
 (22)

Additionally, a multi-resolution scheme is used to represent coarse-to-fine details of both volumes for fast and robust registration [30].

# IV. EXPERIMENTS AND RESULTS

In this section, we present results of experiments on synthetic data, the brain data, and the tongue data, which together demonstrate the performance of the proposed method. Software was implemented using C++ with the Insight Segmentation and Registration Toolkit (ITK) [31], an open source library. For comparison, we used the MI registration algorithm described in Mattes et al. [25], which was also implemented in ITK. All experiments were performed on an Intel i7 CPU with a clock speed of 1.74 GHz and 8 GB memory.

## A. Synthetic Data

In this experiment, we compared the performance of the proposed method to the conventional MI method on a pair of synthetic images at different noise levels, similar to the work by Fan et al. [32]. The synthetic images are shown in Figs. 5(a) and 5(b). The images in Figs. 5(c) and 5(d) are generated to test robustness of the methods against noise and spatial variation. In Fig. 5(c) the object is noisy and spatially variant; the background is spatially variant. The pattern is reversed in Fig. 5(d).

Fig. 6 shows the cost values of MI for the proposed method as a function of horizontal and vertical shifts ranging from -20 to 20 pixels. When Figs. 5(a) and 5(b) were used as template and target images, respectively, both methods found correct alignments as illustrated in Figs. 6(a) and 6(d). However, the MI failed to register as shown in Fig. 6(b) when



Fig. 5. Synthetic images: (a) and (b) are original images and backgrounds. (c) and (d) show spatial gradient or gradient plus noise imposed on the figure or on the background (Image size is  $210 \times 300$ ).



Fig. 6. Comparison between MI (top row) and proposed method (bottom row) on synthetic datasets. (a) and (d) represent the cost value plots between Figs. 5(a) and 5(b) as a function of horizontal and vertical shifts. (b) and (e) show the cost value plots between Figs. 5(c) and 5(d). (e) and (f) show the joint histograms plotted after performing registration with initial misalignments of 5 and 5 pixels in x and y axes.

Figs. 5(c) and 5(d) were used as template and target images, respectively. This is shown more clearly from Figs. 6(c) and 6(f) that the proposed method provided less dispersion of joint histogram compared to the MI after registration of Figs. 5(c) and 5(d) with initial misalignments of 5 and 5 pixels in x and y axes, respectively. The proposed method has good performance on both the noisy images as well as on the noiseless pair. It is observed from these examples that the proposed method provided robust and accurate results in the presence of noise and spatial variation. Although the intensity values inside the object are worsened due to the noise and spatial variance, the proposed method exploited the information from spatially salient regions, enabling the proposed method to find correct alignment and locating the peak value in the center as illustrated in Fig. 6(e). In these experiments, the number of histograms were set to 50 and 80% of the sample size in order to calculate MI. In this case,

the 2D Harris operator was used to extract corner and edge information by setting k = 0.05 and T = 800.

Additional experiments were carried out to test robustness against different noise levels (see Figs. 6(c) and 6(d)) and different spatial variations inside the object and background (see Fig. 7), respectively. To generate the noise, we used Gaussian noise with standard deviation varied from 10% to 50% of the pixel value. Multiple registrations were performed with different initial positions ranging from -6 to 6 pixels in steps of 3 pixels in x and y directions (25 combinations in total per one noise level). A translation spatial transformation was used in these experiments and the root-mean-square (RMS) errors between corrected displacements and initial positions were used as a metric to measure the accuracy of the registration. Registration was considered to be a success if the RMS error was less than 3 pixels. Table I and Table II list the percentage of successful registrations for each method with



Fig. 7. Synthetic images with different spatial variations: (a) low spatial variation, (b) intermediate spatial variation, and (c) high spatial variation (Image size is  $210 \times 300$ ).

TABLE I Success Rate for Both Methods on Synthetic Data (Noise)

	Noise Level	MI	Proposed method
ĺ	0%	90%	100%
	10%	72%	100%
1	20%	40%	100%
	30%	32%	100%
	40%	28%	92%
	50%	32%	92%

TABLE II Success Rate for Both Methods on Synthetic Data (Spatial Variation)

Spatial Variation	MI	Proposed method
Low	90%	90%
Intermediate	24%	68%
High	16%	52%

different noise levels; Table III and Table IV show statistics of RMS error after registration for each method. The performance of the registration using the MI deteriorated as the noise or spatial variation were added whereas the proposed method demonstrated more robust and superior performance in these experiments.

# B. In Vivo Brain Data

1) *RIRE Data:* We used the publicly available brain data from the Retrospective Image Registration Evaluation (RIRE) project [34] to objectively evaluate the performance of the proposed method. The RIRE project provides ground truth transformation to evaluate rigid registration. In our experiments, we focus on the registration between (rectified) MR-T2

TABLE III RMS Error (Pixel) After Registration for Both Methods on Synthetic Data (Noise)

Noise level	MI		Proposed method		
Noise level	Mean±SD	Max	Mean±SD	Max	
0%	$2.28 \pm 1.23$	6.45	$0.06 {\pm} 0.01$	0.06	
10%	$2.48 \pm 1.48$	7.34	$0.07 \pm 0.01$	0.07	
20%	$3.42 \pm 1.79$	6.64	$0.20 \pm 0.01$	0.21	
30%	$4.07 \pm 1.70$	6.65	$0.24{\pm}0.01$	0.24	
40%	$3.87 \pm 1.45$	6.61	$0.69 \pm 1.32$	5.54	
50%	$3.73 \pm 1.53$	6.49	$0.79 \pm 1.58$	6.47	

TABLE IV RMS Error (Pixel) After Registration for Both Methods on Synthetic Data (Spatial Variation)

Spatial Variation	MI		Proposed method		
Spatial variation	Mean±SD	Max	Mean±SD	Max	
Low	$2.16 \pm 0.46$	3.25	$1.40{\pm}1.52$	5.90	
Intermediate	$3.42 \pm 0.63$	4.82	$1.98 \pm 1.95$	8.10	
High	$3.69 \pm 1.42$	9.88	$3.31 \pm 2.94$	11.45	

and PET datasets similar to [33]. The RIRE project provides one training data and seven testing data, five of which only have PET images. Thus we use five datasets for PET-MR registration (i.e., patient 001, patient 002, patient 005, patient 006, and patient 007). The voxel size is  $1.25 \times 1.25 \times 4 \text{ mm}^3$ for MR, and  $2.59 \times 2.59 \times 8 \text{ mm}^3$  for PET images.

2) Evaluation: In order to compare different statistical similarity measures, we used the results reported in [33] as they used the same data. Our method was compared with several entropy-based measures, including MI, normalized MI (NMI), entropy correlation coefficient (ECC), cumulative residual entropy correlation coefficient (CRECC), their modified overlap invariant measures (MMI, MECC, MCRECC) [35], and a learning based method [33]. Please note that we did not reimplement these algorithms but used the results in [33]. The accuracy of the registration was evaluated using target registration error (TRE) [36] provided by the RIRE project website. To obtain TRE, the physical coordinates of the transformed points were uploaded to the RIRE website, thus automatically computing and posting TRE on the website.<sup>1</sup> In our experiments, we set the number of histograms to 50, and used the entire volume as the sample size in order to calculate the modified MI. In addition, the 3D Harris operator was used to label the voxels by setting k = 0.01 and T = 50,000,000. These parameters were set empirically, but the same parameters were used throughout experiments using the RIRE data. The TRE results are shown in Table V. Our proposed method outperformed the other methods in terms of mean TRE. However, the worst case TRE of the learning based method provided a better result compared to the proposed method.

# C. In Vivo Tongue Data

By using hMRI when the tongue is at rest, one can visualize details of this muscle anatomy (see Fig. 1(c)). The motion of

<sup>&</sup>lt;sup>1</sup>http://www.insight-journal.org/rire/

TABLE V PET-MR REGISTRATION ERRORS USING THE RIRE DATA

TRE (mm)	MI	MMI	NMI	ECC	MECC	CRECC	MCRECC	Learning-based Method [33]	Proposed Method
Mean	8.19	3.00	3.20	3.10	2.96	3.34	3.29	2.60	2.44
Median	5.16	2.37	2.64	2.54	2.40	3.01	2.95	2.52	1.81
Max	37.18	7.71	7.57	7.65	7.50	7.53	7.47	4.81	5.81



Fig. 8. Anatomical landmarks used for validation shown in one example case with images from hMRI (a) and images from a cine-MRI (b). We have obtained independent sets of anatomical landmarks from two expert observers. Red crosses indicate the positions of landmarks, including tongue tip, lower lip, and posterior pharyns.

the tongue during speech can be fast and quite complex, however, and unlike the heart, it is not periodic by default. These facts make tongue motion imaging during speech quite challenging. As a consequence, cine magnetic resonance images (cine-MRIs), are captured at movie frame rates but with low spatial resolution (see Fig. 1(d)). Since hMRI captures the structure of the tongue and cine-MRI captures the motion of the tongue, these two MRI "modalities" offer complementary information about the tongue. Our approach allows to further enhance their utility by registering the hMRI data, captured in a resting position, to the cine-MRI data.

1) Subjects and Task: Nine normal native American English speakers were subjects in this experiment. The speech task was "a geese". The word was chosen because the motion is complex. The tongue moves from a neutral tongue shape in "a" to a high back tongue position for "g" which is followed by forward motion into "ee" and tongue body lowering for "s". In addition, this word does not involve jaw opening, which would assist tongue motion. Therefore the functional load, creation of different vocal tract configurations, is placed entirely on the tongue.

2) Recording and Procedure: Both types of MRI datasets hMRI and cine-MRI—were recorded in the same session using a head and neck coil. Cine-MRI datasets were collected with a 6 mm slice thickness and had an in-plane resolution



Fig. 9. Synthetic simulation using tongue data. (a) high-resolution MRI, (b) artificially deformed downsampled hMRI, and (c) the ground truth deformation. The maximum displacement of the ground truth deformation is 15 mm in this case.

of 1.875 mm/pixel. 14 axial slices were acquired. hMRI datasets were 3 mm thick with an in-plane resolution of 0.94 mm/pixel. 24 axial slices were acquired. The subjects were required to remain still from 1.5 to 3 minutes for each plane. The datasets were aligned such that one cine slice contained two hMRI slices. University of Maryland MRI facilities have an MRI trigger system that uses acoustic cues to synchronize speech utterance repetitions with MRI acquisition. The protocol for synchronized auditory cueing is based on the method of Masaki and colleagues [37]. Cine-MRI datasets were collected in multiple planes, while the subject repeated speech tasks ("a geese") to the beat of the auditory rhythm cue. To collect the datasets, the subject repeated the speech task 4 times per slice. A 15-minute training protocol, with feedback from the experimenter, was developed using the subjects. Due to the training, excellent cine images are obtained for naive subjects and patients even with long repetition sets. Recording time can take up to 1 hour and 15 minutes.

3) Evaluation on Simulated Tongue Data: We first validated the accuracy of the proposed method on simulated tongue MR images. Low-resolution MR images were generated by downsampling the hMRI by a factor of 2. Low-resolution MR images were then deformed artificially using randomly generated transformations by setting maximum displacements from 5 to 15 mm. In order to ensure regularity, the transformations were smoothed by a Gaussian kernel with standard deviation  $\sigma = 2$ . Fig. 9 shows an example of the pair of hMRI and a deformed low-resolution MR image, and the ground truth deformation, respectively.

Registration of hMRI (template) with the low-resolution MR image (target) was then performed using MI-based registration and the proposed method. The obtained deformations were compared with the ground truth deformations using the RMS error. We tested the two methods on three subjects and the results are shown in Table VI. It is clearly observed from the results that the proposed approach offers more accurate registration than the MI-based method in all three cases.

TABLE VI RMS Error (mm) Between Ground Truth and Recovered Deformations Using Synthetic Tongue Simulation

Max	S	ub 1	S	Sub 2	S	ub 3
IVIAN	MI	Proposed	MI	Proposed	MI	Proposed
15	1.167	0.959	1.275	1.230	1.361	1.190
14	1.134	0.939	1.223	1.175	1.380	1.199
13	1.102	0.904	1.194	1.122	1.285	1.112
12	1.080	0.869	1.153	1.086	1.257	1.061
11	1.035	0.826	1.089	1.013	1.189	0.986
10	1.027	0.785	1.066	0.971	1.154	0.934
9	1.003	0.745	1.038	0.956	1.127	0.901
8	0.984	0.719	0.994	0.884	1.118	0.852
7	0.968	0.682	0.980	0.876	1.058	0.800
6	0.947	0.651	0.963	0.829	1.036	0.770
5	0.952	0.635	0.940	0.802	0.997	0.731

4) Evaluation on Real Data: Registration was performed on the two static volumes: (1) the first time frame of the axial cine-MRI that was acquired during speech task of "a geese" and (2) the axial hMRI volume that was acquired at rest. We used first time frame of cine-MRI and hMRI for our registration since first frame of cine-MRI is neutral (schwa position) and therefore the position is close to the resting tongue position found in the hMRI data. The registration methods used affine registration as an initialization, followed by the deformable registration using the proposed and MI-based method using FFD. In our experiments, we set the number of histograms to 30, and used 80% of the entire volume as the sample size. We used control point spacings of 16 mm in each axis. For the 3D Harris operator, we set k = 0.001 and T = 50,000,000. The parameters were chosen empirically to provide the best registration performance. The method stops when the movement is less than 0.001 mm or iteration reaches the predefined iteration number 200 in both methods.

To evaluate the accuracy and robustness of the proposed method, we performed two experiments on the nine pairs of 3D axial MRI volumes described above. The first experiment assessed the accuracy of the registration method using TRE. Two expert observers independently selected three corresponding anatomical landmarks from each volume including tongue tip, lower lip, and posterior pharynx as illustrated in Fig. 8. Table VII lists the mean and standard deviation of TRE and inter-observer variability using both methods. In all cases, the original data misalignments were larger than 3.5 voxels. Affine registration and further deformable registration using MI and proposed method reduced the mean TRE to 2.7 and 2.1 voxels, respectively (p < 0.05). The TRE results show that the proposed method provided accurate results compared to the traditional MI-based method. In addition, the TRE results obtained from the proposed method were not much different from the observer variability (p = NS). Fig. 10 shows one result of the first experiment. It is apparent in the figure that the proposed method has better alignment.

The second experiment further demonstrated the performance of the registration method. Three different levels of intensity non-uniformity (bias) were generated including small (20%), medium (40%), and large (60%) bias fields (see Fig. 11). In these experiments, we also used TRE to measure the performance of the methods. As shown in Table VIII, the results of the proposed method were superior to the MI-based registration and were also robust against the bias fields.

## V. DISCUSSION

We used a 3D Harris operator to characterize and label the tissue into three disjoint regions and the local structure matrix was used in the calculation of a modified MI image similarity criterion. It is possible to use other image descriptors to label the tissue. For example, scale invariant feature transform (SIFT) [38] or maximally stable extremal regions (MSER) [39] are likely to give similar pixel labeling results. However, the local structure matrix defined at each voxel offers information about the local geometry, which we were able to exploit in a new weighted MI image similarity computation. As demonstrated in the synthetic experiments, the capture range may not be wide, compared to the MI. However, in practice, we used affine registration as an initialization to further perform deformable registration and therefore the capture range did not affect the results greatly.

Multimodal image registration has been explored in great detail (see [3], [4], [11] and references therein); however, only a few methods have been proposed for tongue images. Yang and Stone [40] proposed to reconstruct 3D tongue motion by aligning temporal data from ultrasound images. Li et al. [41] performed tongue motion averaging to provide best representation of speech motion from several repetitions. Singh et al. [42] proposed to register multiple swallows for generating high temporal resolution in MRI videos. Aron et al. [43] investigated registration of multimodal data including ultrasound, and stereovision data within MRI.

The tongue has three orthogonal muscle fiber directions and extensive fiber inter-digitation and it has no bones or joints. Therefore MRI is an excellent tool because it images soft tissue very well; however, there are challenges. hMRI, collected when the tongue is at rest, can visualize the muscle anatomy in great detail, as shown in Fig. 1(b). Cine-MRI captures tongue motion at movie frame rates, providing good temporal resolution, but with low spatial resolution, as shown in Fig. 1(d). Moreover, these two modalities are different enough that their registration is more akin to multimodal registration than unimodal registration. Our 3D registration approach successfully registered the two datasets. A second challenge is that the motion of the tongue during speech can be fast and quite complex unlike the heart and it is not periodic by default. These facts make high quality tongue image analysis during speech quite challenging. Nonetheless, our method provided registered complementary information: temporal and spatial.

The current study registered cine-MRI to hMRI. This is useful for comparing multiple subjects in the same spatial coordinate space, such as, high-resolution atlas space [44]. In the future, registration from hMRI to cine-MRI is planned. The goal is to incorporate high-quality muscle definition into cine-MRI to better interpret tongue muscle motion from cine-MRI data.



Fig. 10. One example of the results: (a) template image (hMRI) (b) volume labeling of the template image using the Harris operator, (c) a resulting image using MI-based registration, (d) a resulting image using the proposed method and (e) the target image (cine-MRI). The red arrows show that (d) and (e) are better aligned than (c) and (e) in terms of vocal tract edge.



Fig. 11. Different non-uniformity fields: (a) original image, (b) small bias field (20%), (c) medium bias field (40%), (d) large bias field (60%).

 TABLE VII

 Registration Errors and Observer Variability (Voxel)

TRE (voxel)	Before Registration	Affine Registration	MI	Proposed method	Observer Variability
Tongue Tip	$6.2 \pm 3.7$	3.8±1.4	$3.6{\pm}1.8$	$2.5 \pm 1.2$	$1.8{\pm}1.3$
Lower Lip	$3.9{\pm}1.8$	$2.8{\pm}1.1$	$2.6 \pm 1.4$	$1.8 \pm 1.2$	$2.7{\pm}1.7$
Posterior pharynx	$3.9{\pm}5.8$	$1.9{\pm}0.9$	$1.5 {\pm} 0.7$	$1.5 {\pm} 0.9$	$1.4{\pm}1.3$
Average	4.7±4.0	3.1±2.8	2.7±2.6	2.1±1.2	2.0±1.4

 TABLE VIII

 REGISTRATION ERRORS IN DIFFERENT NON-UNIFORMITY FIELDS (VOXEL)

TRE (voxel)	Affine Registration	MI	Proposed method
Small bias field (20%)	$3.8\pm1.6$	$3.5\pm 2.6$	$2.3\pm1.2$
Medium bias field (40%)	$3.7\pm1.3$	$3.6\pm 2.6$	$2.4\pm1.2$
Large bias field (60%)	$3.8\pm1.5$	$3.8\pm 2.5$	$2.7\pm1.5$

To validate the proposed method, experiments with both synthetic and in vivo human datasets including the tongue and the brain were performed. Experiments with synthetic datasets allowed evaluation of the accuracy and robustness of the proposed method to different noise levels. Experiments with human tongue datasets allowed the evaluation of both accuracy and robustness in images with different magnitudes of intensity non-uniformity. Experiments with brain datasets allowed to comparisons of several entropybased measures used in [33] and [35] and allowed to objectively compare the similarity measure itself. This is because the evaluation focused on the rigid transformation and thus the TRE results were independent of the choice of the transformation models or different regularization methods. It was observed that the proposed method provided superior performance even when noise level or the intensity non-uniformity became stronger. Also the proposed method outperformed entropy-based similarity measures such as MMI, ECC, MECC, CRECC, MCRECC [35], and the learning based method [33].

In our method, the choice of T determines how much anatomical detail is included in the registration, which will be different from one application to another. In our case, anatomical details are captured in Type 1 and 2 regions and Type 3 captures a homogeneous region. In the registration process, Type 1 and 2 regions are combined; therefore we chose T based on the Type 3 homogeneous region. Once we set T, we used the same parameters within the same experiments as the anatomical details are similarly presented in the images. The choices of the associated weights of Type 1 and 2 and the parameters for the Harris operator are subject to modification. Although in this work we used spatially meaningful regions (i.e., Type 1 and 2) with the same weights, this can be improved by a mechanism to incorporate importance of each region in an adaptive manner such as the work by Yi et al. [16].

The choice of the width of the Parzen window in estimating MI could influence the registration performance. While a large width will over-smooth the density estimation and mask the structure of the data, a small width will yield a density estimation that is noisy. In our work, we used MI metric available in ITK, which used a third order B-spline function as the Parzen window [31]. The use of B-spline function satisfies the constraint for the partition of unity, while remaining positive, thus being an admissible Parzen window [25], [27]. In addition, in the ITK implementation, in calculating the probability density function, the image intensity values are linearly scaled between zero and one. Thus, in order to handle image data with varying magnitude and range, a fixed B-spline kernel bandwidth of one is used [31].

A challenge for the method to be applied in large population studies and routine clinical practice is the computational cost. It takes on average 1–2 hours for the in vivo tongue experiments on an Intel i7 CPU with a clock speed of 1.74 GHz. This could sped up by either implementing the method with GPU or using parallel computing. In addition, more sophisticated optimization schemes such as gradient descent with backtracking line search could be employed to improve the convergence speed.

Validation of any registration algorithm is a challenging task. Compared to brain registration, both tongue image registration and its validation are inherently more difficult due to the movement of tongue. Therefore, selecting anatomical landmarks is of great importance and, at the same time, a challenging task even for humans, in assessing the accuracy of the registration method. This is because there is no true gold standard other than visual judgment, which is marred by inter-observer variability. The chosen landmarks are soft tissue points that abut the pharyngeal airway. They were chosen because they are visible on both datasets; we did not choose other landmarks due to differences in spatial resolution of muscles, slice thickness, and T2-weighting. For example, the hinge point of the soft palate is not used as a landmark because the hinge point of the soft palate is not comparable in the two datasets and the high-resolution dataset captures quiet breathing and the soft palate is open. The cine dataset captures speech and the soft palate is closed. Therefore the hinge is not the same tissue point in the two datasets. The chosen landmarks are tongue tip, lower lip, and posterior pharynx. In the present study, the inter-observer variability was high especially in the lower lip but the overall TRE was comparable to inter-observer variability (p = NS).

The tongue moves during speech, therefore, different anatomical features including the tongue surface and velum move on cine time frames. As a result, there could be geometrical ambiguities in finding the correct features to match between hMRI and cine-MRI images. Thus, we aligned the hMRI data to the first time frame of the cine image because the tongue is in a relatively neutral position due to the "uh" speech sound being made. The position of the tongue and surrounding structures should be close to the resting tongue position found in hMRI data.

To our knowledge, this is the first report addressing registration of hMRI and cine-MRI in tongue images. These images are multimodal in a certain sense and a method to provide routine coregistration could yield a unique resource in the scientific research of speech science and speech-related disorders. The proposed work holds promise to bridge the two modalities, thereby enriching the information for further tongue image and motion analyses.

## VI. CONCLUSION

In this work, we presented a novel multimodal image registration algorithm. In order to address limitations of the MI, we utilized structural information computed from the 3D Harris operator to encode spatial and geometric cues into the computation of MI. The proposed method was validated extensively on synthetic data, tongue and brain data to demonstrate the benefit of its novel features.

## ACKNOWLEDGMENT

The authors would like to thank Katie Dietrich-Burns for serving as an independent observer.

#### REFERENCES

- J. V. Hajnal, D. L. G. Hill, and D. J. Hawkes, *Medical Image Registra*tion. Boca Raton, FL, USA: CRC Press, 2001.
- [2] J. Woo et al., "Geometric feature-based multimodal image registration of contrast-enhanced cardiac CT with gated myocardial perfusion SPECT," *Med. Phys.*, vol. 36, no. 12, pp. 5467–5479, 2009.
- [3] F. Maes, D. Vandermeulen, and P. Suetens, "Medical image registration using mutual information," *Proc. IEEE*, vol. 91, no. 10, pp. 1699–1722, Oct. 2003.
- [4] B. Zitová and J. Flusser, "Image registration methods: A survey," Image Vis. Comput., vol. 21, no. 11, pp. 977–1000, 2003.
- [5] R. A. McLaughlin, J. Hipwell, D. J. Hawkes, J. Alison Noble, J. V. Byrne, and T. Cox, "A comparison of 2D-3D intensitybased registration and feature-based registration for neurointerventions," in *Proc. 5th Int. Conf. Med. Image Comput. Comput.-Assist. Intervent. (MICCAI)*, 2002, pp. 517–524.
- [6] P. Viola and W. M. Wells, III, "Alignment by maximization of mutual information," Int. J. Comput. Vis., vol. 24, no. 2, pp. 137–154, 1997.
- [7] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal, "Automated multimodality medical image registration using information theory," in *Proc. 14th Int. Conf. Inf. Process. Med. Imag.*, Y. Bizais, C. Barillot, and R. Di Paola, Eds., pp. 263–274, Computational Imaging and Vision 3, Kluwer Academ. Publ., Jun. 1995.
- [8] X. Huang, J. Ren, G. Guiraudon, D. Boughner, and T. M. Peters, "Rapid dynamic image registration of the beating heart for diagnosis and surgical navigation," *IEEE Trans. Med. Imag.*, vol. 28, no. 11, pp. 1802–1814, Nov. 2009.
- [9] D. Gering *et al.*, "An integrated visualization system for surgical planning and guidance using image fusion and an open MR," *J. Magn. Reson. Imag.*, vol. 13, no. 6, pp. 967–975, 2001.
- [10] M. L. Kessler, "Image registration and data fusion in radiation therapy," *Brit. J. Radiol.*, vol. 79, no. 1, pp. 99–108, 2006.
- [11] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever, "Mutualinformation-based registration of medical images: A survey," *IEEE Trans. Med. Imag.*, vol. 22, no. 8, pp. 986–1004, Aug. 2003.

- [12] X. Zhuang, D. J. Hawkes, and S. Ourselin, "Unifying encoding of spatial information in mutual information for nonrigid registration," in *Information Processing in Medical Imaging*. New York, NY, USA: Springer-Verlag, 2009, pp. 491–502.
- [13] J. P. W. Pluim, J. B. Antoine Maintz, and M. A. Viergever, "Image registration by maximization of combined mutual information and gradient information," in *Medical Image Computing and Computer-Assisted Intervention.* New York, NY, USA: Springer-Verlag, 2000, pp. 452–461.
- [14] D. Rueckert, M. J. Clarkson, D. L. G. Hill, and D. J. Hawkes, "Nonrigid registration using higher-order mutual information," *Proc. SPIE*, vol. 3979, p. 438, Mar. 2000.
- [15] D. B. Russakoff, C. Tomasi, T. Rohlfing, and C. R. Maurer, Jr., "Image similarity using mutual information of regions," in *Proc. 8th Eur. Conf. Comput. Vis.*, 2004, pp. 596–607.
- [16] Z. Yi and S. Soatto, "Nonrigid registration combining global and local statistics," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 2200–2207.
- [17] D. Loeckx, P. Slagmolen, F. Maes, D. Vandermeulen, and P. Suetens, "Nonrigid image registration using conditional mutual information," *IEEE Trans. Med. Imag.*, vol. 29, no. 1, pp. 19–29, Jan. 2010.
- [18] X. Zhuang, S. Arridge, D. J. Hawkes, and S. Ourselin, "A nonrigid registration framework using spatially encoded mutual information and free-form deformations," *IEEE Trans. Med. Imag.*, vol. 30, no. 10, pp. 1819–1828, Oct. 2011.
- [19] A. Myronenko and X. Song, "Intensity-based image registration by minimizing residual complexity," *IEEE Trans. Med. Imag.*, vol. 29, no. 11, pp. 1882–1891, Nov. 2010.
- [20] J. Woo, M. Stone, and J. L. Prince, "Deformable registration of highresolution and cine MR tongue images," in *Medical Image Computing* and Computer-Assisted Intervention. Berlin, Germany: Springer-Verlag, 2011, pp. 556–563.
- [21] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. New York, NY, USA: Wiley, Nov. 2001.
- [22] C. Harris and M. Stephens, "A combined corner and edge detector," in Proc. Alvey Vis. Conf., Manchester, U.K., 1988, pp. 147–151.
- [23] H. Luan, F. Qi, Z. Xue, L. Chen, and D. Shen, "Multimodality image registration by maximization of quantitative-qualitative measure of mutual information," *Pattern Recognit.*, vol. 41, no. 1, pp. 285–298, 2008.
- [24] M. Donoser, M. Urschler, M. Hirzer, and H. Bischof, "Saliency driven total variation segmentation," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, Sep./Oct. 2010, pp. 817–824.
- [25] D. Mattes, D. R. Haynor, H. Vesselle, T. K. Lewellen, and W. Eubank, "PET-CT image registration in the chest using free-form deformations," *IEEE Trans. Med. Imag.*, vol. 22, no. 1, pp. 120–128, Jan. 2003.
- [26] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, "Nonrigid registration using free-form deformations: Application to breast MR images," *IEEE Trans. Med. Imag.*, vol. 18, no. 8, pp. 712–721, Aug. 1999.
- [27] P. Thévenaz and M. Unser, "Optimization of mutual information for multiresolution image registration," *IEEE Trans. Image Process.*, vol. 9, no. 12, pp. 2083–2099, Dec. 2000.
- [28] F. Wang and B. C. Vemuri, "Non-rigid multi-modal image registration using cross-cumulative residual entropy," *Int. J. Comput. Vis.*, vol. 74, no. 2, pp. 201–215, 2007.
- [29] S. Klein, M. Staring, and J. P. W. Pluim, "Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines," *IEEE Trans. Image Process.*, vol. 16, no. 12, pp. 2879–2890, Dec. 2007.
- [30] F. Maes, D. Vandermeulen, and P. Suetens, "Comparative evaluation of multiresolution optimization strategies for multimodality image registration by maximization of mutual information," *Med. Image Anal.*, vol. 3, no. 4, pp. 373–386, 1999.
- [31] L. Ibanez, W. Schroeder, L. Ng, and J. Cates, *The ITK Software Guide: The Insight Segmentation and Registration Toolkit*. New York, NY, USA: Kitware, 2003.
- [32] X. Fan, H. Rhody, and E. Saber, "A spatial-feature-enhanced MMI algorithm for multimodal airborne image registration," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 6, pp. 2580–2589, Jun. 2010.
- [33] D. Lee, M. Hofmann, F. Steinke, Y. Altun, N. D. Cahill, and B. Scholkopf, "Learning similarity measure for multi-modal 3D image registration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 186–193.
- [34] J. West *et al.*, "Comparison and evaluation of retrospective intermodality brain image registration techniques," *J. Comput. Assist. Tomogr.*, vol. 21, no. 4, pp. 554–568, 1997.

- [35] N. D. Cahill, J. A. Schnabel, J. A. Noble, and D. J. Hawkes, "Revisiting overlap invariance in medical image alignment," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2008, pp. 1–8.
- [36] J. M. Fitzpatrick, J. B. West, and C. R. Maurer, Jr., "Predicting error in rigid-body point-based registration," *IEEE Trans. Med. Imag.*, vol. 17, no. 5, pp. 694–702, Oct. 1998.
- [37] S. Masaki *et al.*, "MRI-based speech production study using a synchronized sampling method," *J. Acoust. Soc. Jpn.*, vol. 20, no. 5, pp. 375–379, 1999.
- [38] D. G. Lowe, "Object recognition from local scale-invariant features," in *Proc. 7th IEEE Int. Conf. Comput. Vis.*, vol. 2. Sep. 1999, pp. 1150–1157.
- [39] J. Matas, O. Chum, M. Urban, and T. Pajdla, "Robust wide-baseline stereo from maximally stable extremal regions," *Image Vis. Comput.*, vol. 22, no. 10, pp. 761–767, 2004.
- [40] C. Yang and M. Stone, "Dynamic programming method for temporal registration of three-dimensional tongue surface motion from multiple utterances," *Speech Commun.*, vol. 38, nos. 1–2, pp. 201–209, 2002.
- [41] M. Li, C. Kambhamettu, and M. Stone, "Tongue motion averaging from contour sequences," *Clin. Linguistics Phonetics*, vol. 19, nos. 6–7, pp. 515–528, 2005.
- [42] M. Singh, R. Thompson, A. Basu, J. Rieger, and M. Mandal, "Image based temporal registration of MRI data for medical visualization," in *Proc. IEEE Int. Conf. Image Process.*, Oct. 2006, pp. 1169–1172.
- [43] M. Aron, A. Toutios, M.-O. Berger, E. Kerrien, B. Wrobel-Dautcourt, and Y. Laprie, "Registration of multimodal data for estimating the parameters of an articulatory model," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Apr. 2009, pp. 4489–4492.
- [44] J. Woo et al., "A high-resolution atlas and statistical model of the vocal tract from structural MRI," Comput. Methods Biomech. Biomed. Eng., Imag. Vis., to be published.



Jonghye Woo (S'08-M'10) received the B.S. degree from Seoul National University, Seoul, Korea, in 2005, and the M.S. and Ph.D. degrees from the University of Southern California (USC) at Los Angeles, Los Angeles, CA, USA, in 2007 and 2009, respectively, all in electrical engineering. He is currently a Faculty Member with the Department of Radiology, Massachusetts General Hospital, Boston, MA, USA, and the Harvard Medical School, Boston, after his post-doctoral training with Cedars-Sinai Medical Center, Los Angeles, CA, USA, the Uni-

versity of Maryland, Baltimore, MD, USA, and Johns Hopkins University, Baltimore, MD, USA. His current research interests are in computer vision and machine learning with an application to medical imaging. He was a recipient of the USC Viterbi School of Engineering Best Dissertation Award in 2010 and the NIH Pathway to Independence Award in 2013. He is a member of the International Society for Optics and Photonics, the Medical Image Computing and Computer Assisted Intervention, the Acoustical Society of America, and Sigma Xi.



**Maureen Stone** is currently a Professor with the School of Dentistry, University of Maryland at Baltimore, Baltimore, MD, USA, with a joint appointment with the Department of Neural and Pain Sciences and the Department of Orthodontics. She is the Director of the Vocal Tract Visualization Laboratory. She studies tongue biomechanics and motor control using several imaging techniques, in particular, ultrasound and magnetic resonance imaging (MRI). She is studying the effects of glossectomy surgery on tongue motion using MRI. The

goal is to determine how motor control strategies have adapted and which ones are successful in producing high-quality speech. Four techniques are used to discover the effects of subject-based and surgery-based factors on post-surgical 3D tongue anatomy, motion, and speech. These are tagged MRI, cine MRI, high-resolution MRI, and diffusion tensor MRI. She has written numerous articles on multi-instrumental approaches to studying vocal tract function. She is also a fellow of the Acoustical Society of America.



Jerry L. Prince (S'78–M'83–SM'96–F'04) received the B.S. degree from the University of Connecticut, Storrs, CT, USA, in 1979, and the S.M., E.E., and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1982, 1986, and 1988, respectively, all in electrical engineering and computer science.

He joined The Analytic Sciences Corporation, Inc., Reading, MA, USA, as the Technical Staff, where he contributed to the design of an automated vision system for synthetic aperture radar imaging,

after receiving the Ph.D. degree in geometric reconstruction in computed tomography. He joined Johns Hopkins University, Baltimore, MD, USA, as a faculty, in 1989, where he has been a Professor with the Department of Electrical and Computer Engineering since 1998, and received the William B. Kouwenhoven Chaired Professorship in 2001. He holds joint appointments with the Department of Radiology and the Department of Biomedical Engineering, and secondary appointments with the Department of Applied Mathematics and Statistics and the Department of Computer Science. He was an Associate Director for Research of the Engineering Research Center for Computer Integrated Surgical Systems and Technology from 2001 to 2005.

Over the past 25 years, he has authored numerous book chapters, peerreviewed journal and conference papers, and abstracts. He co-edited the book entitled *Measurement of Cardiac Deformations from Magnetic Resonance Imaging (MRI): Physical and Mathematical Models* with A. A. Amini, and co-authored a textbook entitled *Medical Imaging Signals and Systems* with J. M. Links, now in its second edition. His current research interests are in image processing and computer vision with an application to medical imaging. His major projects, including MRI of tissue motion, 3D brain image analysis, optical coherence tomography of the retina, and computer integrated surgery.

Dr. Prince has served as an Associate Editor of the IEEE TRANSACTIONS ON IMAGE PROCESSING and the IEEE TRANSACTIONS ON MEDICAL IMAG-ING. He is on the Editorial Board of *Medical Image Analysis* and is a member of the Sigma Xi Professional Society and the Tau Beta Pi, Eta Kappa Nu, and Phi Kappa Phi Honor Societies. He was a member of the IEEE Signal Processing Society Image and Multidimensional Signal Processing Technical Committee from 1994 to 2000. He was a recipient of several awards, including the 1993 National Science Foundation Presidential Faculty Fellows Award and the 1997 Maryland Outstanding Young Engineer Award. He became a fellow of the Medical Image Computing and Computer Assisted Intervention (MICCAI) Society in 2011, and was also a recipient of the MICCAI Society Enduring Impact Award in 2012.