

Research Article

Slice-to-Volume Nonrigid Registration of Histological Sections to MR Images of the Human Brain

Sergey Osechinskiy¹ and Frithjof Kruggel¹

¹Department of Biomedical Engineering, University of California, Irvine, CA 92697, USA

Correspondence should be addressed to Sergey Osechinskiy, sosechin@uci.edu

Received 15 June 2010; Revised 12 August 2010; Accepted 8 September 2010

Academic Editor: Feng C. Zhou

Copyright © 2011 S. Osechinskiy and F. Kruggel. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Registration of histological images to three-dimensional imaging modalities is an important step in quantitative analysis of brain structure, in architectonic mapping of the brain, and in investigation of the pathology of a brain disease. Reconstruction of histology volume from serial sections is a well-established procedure, but it does not address registration of individual slices from sparse sections, which is the aim of the slice-to-volume approach. This study presents a flexible framework for intensity-based slice-to-volume nonrigid registration algorithms with a geometric transformation deformation field parametrized by various classes of spline functions: thin-plate splines (TPS), Gaussian elastic body splines (GEBS), or cubic B-splines. Algorithms are applied to cross-modality registration of histological and magnetic resonance images of the human brain. Registration performance is evaluated across a range of optimization algorithms and intensity-based cost functions. For a particular case of histological data, best results are obtained with a TPS three-dimensional (3D) warp, a new unconstrained optimization algorithm (NEWUOA), and a correlation-coefficient-based cost function.

1. Introduction

Registration of histological sections to magnetic resonance (MR) images is an important step in quantitative analysis and cross-modality comparison of brain histology and MR data. Magnetic resonance imaging (MRI) and histological examination are two complementary modalities that considerably differ in their scope of application and scale of analysis. MRI of the brain became an indispensable method in clinical evaluation of central nervous system (CNS) diseases and in brain research. MRI is extensively used for *in vivo* characterization of intracranial lesions and tumors, for longitudinal studies of neurodegenerative effects, and so forth. Postmortem MRI of isolated formalin-fixed brain specimens is also becoming a promising method for the investigation of the pathology of a brain disease [1], where it can serve as a high-throughput screening tool detecting regions with abnormal signal contrast. Unfortunately, in the near future, the spatial resolution of MRI is not likely to improve beyond an intermediate, mesoscopic scale of 0.1 mm, contrary to

histology that has long been operating on a microscopic, cellular scale. In addition, MRI contrast mechanisms have complex origins that make exact relationships between MRI-detected abnormal regions and pathological characteristics unclear; correlates of MRI findings have to be studied and validated. Histological examination remains a gold standard for precise characterization of neuropathology. However, histology has its own share of shortcomings: it is a low-throughput, inherently two-dimensional (2D), time-consuming process, which is invasive, irreversible to the specimen, and destructive to gross morphological features. Consequently, there is a great value in bringing these two modalities together, to complement each other and bridge the scales of mesoscopic and microscopic examination. In doing so, it is important to establish the best possible spatial correspondence between the two, which can be achieved by the design of a specialized brain slicer [2] or, more generally, by the registration approach (e.g., see [3]). Registration may become a key technology bridging these two modalities. Specifically, the registration of individual

histological slices from sparse section sets directly to three-dimensional (3D) MR images (*slice-to-volume* registration) constitutes an important unsolved problem.

Apart from neuropathological investigations, histology-to-MRI registration plays an important role in quantitative analysis of brain structure and, more specifically, in architectonic mapping of the brain. Cyto- or myeloarchitectonic mapping of the brain cortex has long since relied on examination of histological specimens, but the spatial location of delineated areas was traditionally sketched only roughly, as exemplified by renowned Brodmann maps. In modern development, architectonic mapping features observer-independent quantitative analysis of histology and is brought into a standard anatomical spatial reference system by reconstruction and stereotaxic alignment of a histological volume. Schormann and Zilles [4] describe reconstruction of histology volume from serial sections, which at one point involves registration to MRI. It should be noted that the methodology involved in reconstruction of a histology volume, which has become a well-established procedure, relies on coregistration of serial histological sections, and is not directly applicable to registration of individual slices from sparse sections.

Advances in high-resolution MRI make an alternative approach to architectonic parcellation of the cortex feasible: homogeneous areas can be delineated directly from MR images, potentially *in vivo*. However, as in the case of neuropathological investigation, correlation of MRI findings to the gold standard of histological examination must be studied and validated. Our present work in histology-to-MRI registration is primarily motivated by the search for a ground truth reference to validate neocortical parcellation results obtained by Kruggel et al. [5, 6]. Previous studies have demonstrated a striking qualitative similarity of optical profiles in myelin-stained tissue with intensity profiles of high-resolution MRI in a fixated brain [5, 6]. Profiles of histological specimens stained for myelin and cell bodies were quantitatively compared to *in vivo* T1 and postmortem T2 MRI by Eickhoff et al. in [7], focusing on two architectonic regions only: a striate and an extrastriate cortex. It was found that MR intensity profiles were best fitted by a weighted sum of both myeloarchitectonic and cytoarchitectonic profiles, with a significantly larger contribution from the former. There, an approximate correspondence of histological and MRI slices was established manually, and only two image modalities were originating from the same brain specimen. Our study is seeking to establish a reference for quantitative comparison of histological and MRI intensity profiles, where correspondence of section images is optimized by a nonrigid registration procedure.

The registration of histological images is challenging because of three main issues: (1) a “deformation problem”: a brain sample is subject to complex deformations in the process of preparation and histological sectioning; in case of sparse histological slices (2) a “search problem”: an inherent one-to-many correspondence of the underlying 2D-to-3D mapping; (3) a “statistical consistency” problem: the reliance of a similarity measure on small sampling size of an input 2D image.

Nonlinear and nonuniform global and local distortions pertinent to different stages of histological sectioning procedure are described in [8]. A brain sample undergoes global shrinking during fixation in formalin, and each slice is experiencing local uncorrelated nonlinear distortions when sectioned. In addition, nonelastic deformations such as folds and rips may be introduced when a section is transferred to the object slide for scanning. Schormann et al. [9] list nine sources of deformations induced by histological processing and present a theoretical analysis of the statistics of deformations, emphasizing that local deformations impede exact registration of MRI with histological sections.

Only a limited number of papers have been published on the particular topic of histology-to-MRI registration (see [8] for a review), most of them dealing with reconstruction of histology volume from serial sections. When contiguous data for serial sections are available, the issue of a 2D-to-3D mapping may be bypassed, and the registration can proceed in two stages: first, sequential slices are 2D aligned into a 3D histology volume and second, the reconstructed histology volume is 3D registered with an MRI or a positron emission tomography (PET) image. This way, each stage can benefit from a vast body of research and algorithms developed for 2D and 3D, intra- and intermodal, rigid and nonrigid registration. For example, Ourselin et al. [8] applied a block matching algorithm to 2D align histological sections into a reconstructed volume, which was further registered to MRI data using a 3D version of the same algorithm. If available, optical images of a block face, for example, photos of the cryomacrotome cut view with external landmarks, can be used as intermediate data, facilitating the stage of 2D alignment in histology volume reconstruction. Mega et al. [10] reconstructed a cryovolume from optical data, adapted a 3D elastic warping/surface matching approach to a single 2D slice for warping of each stained section image to a corresponding optical image, and applied a 3D rigid registration to align the cryovolume with PET. Similarly, Schormann and Zilles [4] reconstructed a histology volume by 2D alignment with optical data, 3D aligned it to MRI, and applied full multigrid elastic registration with MRI to refine the histology volume reconstruction. Intermediate optical images may also serve as a reference in case histological sections are not contiguous, for example, see [11].

When histology data are sparse, and no intermediate optical data are available, coregistration with MRI is posed as a 2D-to-3D registration problem. The 2D-to-3D medical image registration has received considerable attention in the context of projective registration in intraoperative imaging modalities, for example, X-ray fluoroscopy with CT images. Projective registration is seeking to find a correspondence between reference 2D image and a projection of a volume to a plane. On the contrary, there are relatively few publications on the topic of slice-to-volume registration, which is seeking to find correspondence between reference 2D image and a cross-section of a volume by a plane or a warped surface. Different approaches to rigid slice-to-volume registration are described in [12–14]. In prostate imaging, Fei et al.

[12] registered interventional MRI (iMRI) slices to preoperative high-resolution MRI volume, using multiresolution approach with two cost functions and automatic restarting to avoid local minima. Chandler et al. [13] evaluated slice-to-volume registration in connection to the planning of MR scans, registering single thick brain MR slice, and applying information-based criterion on vector-valued probabilistic images of tissue classification. Birkfellner et al. [14] studied registration of single slices from FluoroCT, CineMR, or iMRI to 3D volumes, and used cross-correlation with repeated application of local optimization algorithms. All authors in [12–14] note that slice-to-volume registration raises several challenges in cost function selection, convergence behavior, and optimization strategy, which require special considerations. We will return to a more detailed discussion in Section 2.

Slice-to-volume registration of histology to MRI was studied by few research groups. Jacobs et al. [15] registered rodent brain histology to MRI in two steps: first by rigid alignment of a histological slice with an MR image based on a surface matching algorithm, followed by a 2D warping of a planar section of the MRI volume onto a histological image. Here, the 2D warping step used thin-plate splines with matching control points automatically selected along both sets of contours. Similarly, Gefen et al. [16] registered sections of a mouse brain to a 3D atlas by first matching them to a corresponding atlas plane with a global affine transformation, and then by warping in 2D. However, modeling the sectioning surface as a plane may not correctly reflect global and local deformations of a brain sample; generally, such a cross-section is better modeled as a warped surface. In a study of the pathology of disease, Kim et al. [3] registered postmortem human brain histology images to premortem MRI reference volume. Warped cross-sections of MRI volume were modeled by polynomial functions of second or third order, recursively derived and matched to postmortem slices. A polynomial model allows to compensate for global distortions between premortem and postmortem images, but does not account for local deformations.

We present a flexible framework for intensity-based slice-to-volume nonrigid registration algorithms with geometric transformation that combines a rigid alignment with a 3D deformation field parametrized by various classes of spline functions: thin-plate splines (TPS) [17], Gaussian elastic body splines (GEBS) [18], or cubic B-splines [19]. We apply algorithms to cross-modality registration of histological and MR images of the human brain and evaluate registration performance across a range of optimization algorithms and intensity-based cost functions.

The remainder of the paper is organized as follows. In Section 2, we present the building blocks of the algorithmic framework: geometric transformations, cost functions with similarity measures, and optimization strategies. In Section 3, we describe our computational experiments and discuss their results, comparing registration performance across spline models, similarity measures, and optimization algorithms. The last section discusses the key features, various algorithmic details, and future developments.

2. Algorithms

In our framework, an intensity-based registration algorithm consists of three main building blocks: (1) a geometric transformation with a deformation model, (2) a cost function with a similarity measure, and (3) an optimization scheme. In this section, we describe three implementations of (1) with a 3D deformation field parametrized by TPS, GEBS, or B-splines, several implementations of (2) with various similarity measures, and several optimizers and optimization schemes implementing (3).

2.1. Problem Formulation and General Notation. Slice-to-volume nonrigid registration takes as an input a 2D image (called target or floating image) and finds a corresponding warped cross-section slice (called registered image, 2D) in a 3D volume (called source or reference image). Let $I_d(i, j)$ and $V_d(i, j, k)$ denote a 2D histological image and a 3D MR discrete grey-scale image, respectively. Let $I((x, y) \in \Omega) \in \mathbb{R}$ and $V((x, y, z) \in \Phi) \in \mathbb{R}$ denote the real-valued intensity function of a continuous version of a 2D and a 3D image with domain $\Omega \subset \mathbb{R}^2$ and $\Phi \subset \mathbb{R}^3$, respectively. Continuous images are obtained from discrete images I_d and V_d by a chosen interpolation model, for example by trilinear or spline interpolation. A geometric transformation \mathcal{T} is defined as a mapping $\mathcal{T} : \Omega \mapsto \Psi, \Psi \subset \Phi$; the mapping defines the domain of a warped slice $\Psi \subset \mathbb{R}^3$. Then a registered image I_R is defined as the following intensity function: $I_R(V, \mathcal{T}) = \{i_R = V(\vec{r}), \vec{r} \in \Psi\}$. Let $S_{\text{sim}}(I_1, I_2)$ denote an intensity-based similarity measure between two images I_1, I_2 , such that closer similarity of images results in higher value. A cost function is defined as: $C(I, V, \mathcal{T}) = -S_{\text{sim}}(I, I_R) + \beta E(\mathcal{T}(\Omega))$, where $E(\mathcal{T}(\Omega))$ is the deformation energy of the geometric transformation, and β is the weight parameter at the energy penalty term. Image registration seeks the optimal transformation $\mathcal{T} = \arg \min(C(I, V, \mathcal{T}))$ that minimizes the cost function.

2.2. Geometric Transformation. The slice-to-volume geometric transformation \mathcal{T} is composed of two parts: alignment and deformation, $\mathcal{T} = \mathcal{T}_{\text{aln}} \circ \mathcal{T}_{\text{def}}$. The flat 2D slice is deformed in 3D, and the resulting warped surface is rotated, translated, and scaled within bounds of a 3D volume; the registered image is obtained at the cross-section of the volume by the warped surface.

2.2.1. Alignment. Alignment \mathcal{T}_{aln} is modeled by a 3D linear operator $\mathcal{T}_{\text{aln}}[\vec{x}] = \mathbf{A}(\vec{\alpha}, \vec{\kappa})\vec{x} + \vec{r}$, where \vec{x} denotes 3D coordinates of a transformed point, \mathbf{A} is a 3×3 rotation/scaling matrix, $\vec{\alpha}$ is a vector with three rotation angles, $\vec{\kappa}$ is a vector with three scaling factors, and \vec{r} is a translation 3D vector. With scaling factors, \mathcal{T}_{aln} is called Procrustes transformation and has nine degrees of freedom (DOF). If scaling is omitted, \mathcal{T}_{aln} corresponds to rigid transformation with six DOFs.

2.2.2. Thin-Plate Spline Deformation Field. Thin-plate spline (TPS) is a solution of an interpolation problem that minimizes a certain smoothness energy function based on local

curvature [17]. TPS interpolation function is defined as

$$T(x, y, \boldsymbol{\theta}) = a_0 + a_1x + a_2y + \sum_{(x_i, y_i) \in \boldsymbol{\theta}} c_i U(\|(x_i, y_i) - (x, y)\|), \quad (1)$$

where $U(r) = r^2 \log r$ is a radial basis function (RBF), $\boldsymbol{\theta} = \{(x_i, y_i)\}$ are control points (where RBFs are centered), $\{c_i\}$ are RBF weights, and $\{a_0, a_1, a_2\}$ are affine coefficients. Given a set of interpolation conditions $T(x_i, y_i) = t_i \{(x_i, y_i) \in \boldsymbol{\theta}\}$, there exists a closed form solution for TPS coefficients [20]. TPS interpolation minimizes the following functional, called TPS bending energy:

$$E_T = \iint_{\mathbb{R}^2} [T_{xx}^2 + 2T_{xy}^2 + T_{yy}^2] dx dy. \quad (2)$$

TPSs are typically used in a 2D nonrigid registration based on point-landmark matching [15, 21]: point-landmarks in a target image serve as control points, and the displacement (u, v) of a 2D warp is modeled by two independent TPS functions. Matched landmarks in a source image define displacements at the control points, from which TPS weights can be computed by solving a system of linear equations [17]. Our slice-to-volume registration does not use landmarks. Instead, TPSs are used in a novel way to parametrize a smooth 3D deformation of a 2D surface: control points are placed in a regular grid on the 2D image domain Ω , and a 3D warp is modeled by three independent TPS functions $u(x, y, \boldsymbol{\theta})$, $v(x, y, \boldsymbol{\theta})$, and $w(x, y, \boldsymbol{\theta})$: $\mathcal{T}_{\text{def}} : (x, y) \mapsto (x + u, y + v, w)$. The displacements (u_i, v_i, w_i) at the control points serve as DOFs and thus “steer” the deformation; TPS coefficients are computed from these displacements. The smoothness energy of such a deformation field is the sum of three individual TPS bending energies $E_{\mathcal{T}} = E_u + E_v + E_w$, which can easily be computed from TPS weights [20].

2.2.3. Gaussian Elastic Body Spline Deformation Field. Gaussian elastic body splines (GEBS) [18] model 3D deformations of elastic homogeneous isotropic media under the influence of forces that are applied at control points and are decreasing with distance as Gaussian functions. The deformation is defined as:

$$\vec{T}(\vec{x}) = \vec{x} + \sum_{\vec{p}_i \in \boldsymbol{\theta}} \mathbf{G}(\vec{x} - \vec{p}_i) \vec{c}_i, \quad (3)$$

where $\mathbf{G}(\vec{x})$ is a 3×3 matrix of basis functions, and $\boldsymbol{\theta} = \{\vec{p}_i\}$ are control points in 3D. For given interpolation/approximation conditions, the coefficients \vec{c}_i are determined from a linear system of equations (see [18, 22]). As with TPS, we define control points $\boldsymbol{\theta}$ on a regular 2D grid. The 3D displacements of control points serve as DOFs and specify an interpolating/approximating GEBS deformation. We characterize the smoothness of GEBS deformation by the same energy functional E_T (TPS bending energy), which is commonly used in registration [19, 23]. An analytic expression for E_T is not readily available; therefore, we compute it numerically, approximating partial derivatives by finite difference schemes.

2.2.4. B-Spline Free-Form Deformation Field. Free-form deformation (FFD) using 3D tensor product of cubic B-splines is described in [19] and is generalized in [24]. We define a 3D FFD of a 2D domain as

$$\vec{T}(x, y, \boldsymbol{\theta}) = (x, y, 0)' + \sum_{i=0}^{L-1} \sum_{j=0}^{M-1} \beta_3\left(\frac{x}{h} - i\right) \beta_3\left(\frac{y}{h} - j\right) \vec{c}_{i,j}, \quad (4)$$

where $\boldsymbol{\theta} = \{\vec{c}_{i,j}\}$ variables are 3D displacements of control points in a 2D $L \times M$ grid that has regular spacing h and β_3 denotes the cubic B-spline basis function [25]. Four main distinctions of a B-spline FFD are noted below: (1) in contrast to TPS or GEBS, where each control point contributes globally, FFD is locally controlled, that is, each control point affects only its local neighborhood of 16 grid tiles; (2) FFD does not require solving a linear system for computing weights from displacements; (3) for B-splines, control points must be placed in a regular grid, and must enclose the domain boundaries (“closed grid”), whereas for TPS or GEBS control points can, in principle, be arbitrarily placed inside the domain, and a grid can have open boundaries (“open grid”); consequently, for the same domain size and grid spacing, FFD will have a larger number of control points and DOFs; (4) unlike TPS and interpolating GEBS, a surface warped by FFD does not necessarily pass through the displaced control points (see properties of β_3 in [25]). We regularize B-spline FFD with the same TPS bending energy functional, which can be efficiently computed from analytical expression, because B-spline derivatives have local support.

2.3. Cost Function Similarity Measures. Our registration framework provides a flexible selection of cost function similarity measure:

$$\text{SSD: } S_{\text{SSD}}(I_1, I_2) = -\frac{1}{N} \sum_i (I_1(i) - I_2(i))^2,$$

$$\text{CC: } S_{\text{CC}}(I_1, I_2) = \frac{\sum_i (I_1(i) - \langle I_1 \rangle) (I_2(i) - \langle I_2 \rangle)}{\sqrt{\sum_i (I_1(i) - \langle I_1 \rangle)^2 \sum_i (I_2(i) - \langle I_2 \rangle)^2}},$$

$$\text{MI: } S_{\text{MI}}(I_1, I_2) = H(I_1) + H(I_2) - H(I_1, I_2),$$

$$\text{NMI: } S_{\text{NMI}}(I_1, I_2) = \frac{2S_{\text{MI}}}{H(I_1) + H(I_2)},$$

$$\text{NMI}_1: S_{\text{NMI}_1}(I_1, I_2) = \frac{H(I_1) + H(I_2)}{H(I_1, I_2)},$$

$$\text{LC: } S_{\text{LC}}(I_1, I_2)$$

$$= \frac{1}{N} \sum_j \frac{\left(\sum_{i \in n(j)} (I_1(i) - \langle I_1 \rangle_j) (I_2(i) - \langle I_2 \rangle_j) \right)^2}{\sum_{i \in n(j)} (I_1(i) - \langle I_1 \rangle_j)^2 \sum_{i \in n(j)} (I_2(i) - \langle I_2 \rangle_j)^2},$$

$$\begin{aligned} \text{RC: } S_{\text{RC}}(I_1, I_2) &= S_{\text{CC}}(\text{Rank}(I_1), \text{Rank}(I_2)), \\ \text{CR: } S_{\text{CR}}(I_1 | I_2) &= \frac{\text{var}(E(I_1 | I_2))}{\text{var}(I_1)}, \end{aligned} \quad (5)$$

where N is the number of pixels in Ω , $H(X)$ and $H(X, Y)$ correspond to the entropy and joint entropy, $E(X|Y)$ is the conditional expectation of X in terms of Y , $n(j)$ is the local neighborhood of j th voxel, and brackets $\langle \rangle$ denote averaging. The listed similarity measures are: sum of squared differences (SSD) [26], Pearson's correlation coefficient (CC) [27], mutual information (MI) and normalized MI (NMI) [28], overlap-invariant NMI (NMI₁) [29], local correlation (LC) [30], Spearman's rank correlation coefficient (RC) [31], and correlation ratio (CR) [32]. The similarity measures can be ordered with respect to assumed optimal relationship between image intensities, from more restrictive to less restrictive assumption: SSD (additive Gaussian noise), CC (globally linear relationship), LC (linear transfer functions for small patches), RC (some monotonic relationship), CR (some functional dependence), and MI-NMI (statistical dependence). The choice of a similarity measure for a particular registration application depends not only on relationship between imaging modalities, but also on other factors, like robustness with respect to outliers/image defects, computational cost, sensitivity, statistical consistency (adequate sampling size), presence of local maxima, and availability of analytical derivatives for derivative-based optimization.

2.4. Optimization. The optimal transformation is sought by numerical multidimensional nonlinear optimization. For our geometric transformations, the total number of optimization variables is $N_{\text{opt}} = N_{\text{aln}} + 3N_{\text{cp}}$, where N_{aln} is equal to nine or six for Procrustes or rigid transformation, respectively, and N_{cp} is the number of the control points in a 2D grid that defines a spline deformation model.

Our registration framework provides modular selection of optimization algorithms. The following derivative-free optimizers are included: New Unconstrained Optimization Algorithm (NEWUOA) [33], Powell's Direction Set (PDS) method [31], Nelder-Mead's Downhill Simplex [31], Genetic [34], and Differential Evolution (DE) [35] algorithm (see [36–39] for more details on some of these algorithms in medical image registration). For the SSD similarity measure and B-spline deformation field combination only, two derivative-based optimizers are implemented: Conjugate Gradient (CG) [31] and Levenberg-Marquardt-(LM-) type [31] algorithm (see [24, 37, 40] for more details on LM in registration applications).

The registration framework also allows for an evaluation of different optimization strategies: multiresolution and/or multistage optimization schemes. A multiresolution scheme [36] proceeds through several stages, from coarser to finer resolution, for example, on a Gaussian pyramid of input images. A multistage scheme partitions the optimization into several sequential stages, each stage being focused on a particular aspect of the problem such as initialization, a subset of optimization variables, a particular cost function,

or a certain optimizer. For example, a two-stage scheme can combine global optimization of rigid alignment by DE and local optimization of deformation DOFs by PDS or NEWUOA.

3. Experiments and Result Discussion

3.1. Data and Preprocessing. High-resolution MR images of an isolated left brain hemisphere fixated in 4% paraformaldehyde were acquired postmortem on a Bruker 3T Medspec 100 system equipped with a bird cage quadrature coil, using T_1 -weighted 3D MDEFT protocol (FOV $179.2 \times 89.6 \times 179.2$ mm, matrix $512 \times 256 \times 256$, voxel size $0.35 \times 0.35 \times 0.7$ mm³, scanning time 12 h). The white-grey matter contrast of the MR images is inverted due to fixation in formalin, which makes it similar to the contrast in histological images (see Figure 1). After MRI scanning, the brain sample was cut on a macrotome into ten coronal blocks, each approx. 1.5 cm thick in sagittal direction. Blocks were frozen to -80°C , and up to five sections were cut from the face of each block using cryomicrotome. The $50 \mu\text{m}$ -thick slices were washed in 30% sucrose in sodium phosphate buffer (PBS), glass-mounted, stained for myelin with Sudan Black B, and scanned on a flat-bed scanner at 2000dpi resolution ($12.7 \mu\text{m}/\text{pix}$). The MR image was preprocessed to correct for intensity inhomogeneities, and converted to the isotropic resolution of 0.35 mm^3 . Histological images were converted to an 8-bit grey-scale intensity range (0 to 255), and resampled to match MRI spatial resolution; the bright background was suppressed by setting intensities above 240 to zero. A continuous model of the discrete MR image was built with the fourth degree B-spline interpolation [25]. Registration experiments were applied to histological slices from eight distinct blocks of the brain sample. Because both histology and MR images originated from a formalin-fixated brain, there was no global shrinkage/swelling to account for, therefore a rigid transformation was used in the alignment part \mathcal{T}_{aln} (scaling was disabled). The software was developed in C++ in a Linux environment; running times are reported on a 2.4 GHz Athlon-64 CPU with 2 GB RAM.

3.2. Initialization and Choice of Parameters. This subsection explains how initial conditions and parameter values were chosen. Initial conditions specify the starting vector of optimization variables. Initial values for translation $\vec{\tau}$ and rotation $\vec{\alpha}$ were first determined approximately by a human observer, and were further refined by a sequence of fast rigid registrations sampling the diagonals of the $\vec{\tau}, \vec{\alpha}$ DOF space; they were also confirmed by global optimization with DE algorithm. The deformation field was always initialized to zero displacements of the control points, that is, no deformation.

The right choice of deformation energy penalty weight β (Section 2.1) is very important for high-quality nonrigid registration results. If β is too low, the deformation is not penalized, and a higher value of intensity similarity measure can be reached at the expense of unrealistic deformations, producing clearly visible artifacts and distortions in the




FIGURE 1: Histological section (a) and MR image (b).

registered image. If β is too large, then deformation is suppressed, and registration approaches the rigid alignment. A suitable value of β has been established experimentally: the similarity measure and bending energy of registered images were plotted as a function of β , and the value of β was picked from an interval that was preceding a sharp rise in bending energy, and where the similarity measure was flat. The value of $\beta = 0.1$ has been chosen for our experiments reported below.

Control point grid spacing influences both the spatial scale of control over deformation, and the dimensionality N_{opt} of the optimization problem. A coarser grid does not model local deformations at a smaller scale, but is easier for the optimizer. Finer grids allow fine-tuning of deformation, but may be very difficult for optimization. As an empirical trade-off between an increase in registration time and an improvement of similarity measure, the grid spacing of 15 and 20 pixels has been chosen for histological image area below and above 40,000 square pixels, respectively.

3.3. Metaoptimization. Choosing the right similarity measure for intensity-based registration can be challenging. Although some arguments in favor of a certain similarity measure can be based solely on image modalities (e.g., SSD or NMI is a typical choice in intra- or intermodal registration, resp.), the best cost function has to be determined experimentally. We test different cost functions and evaluate registration results by computing (post registration) the *cumulative similarity measure* SM as a sum of seven coefficients: CC, LC, RC, CR, NMI, NMI_1 , and Dice's coefficient (DC) on segmented cortical ribbon (intensity thresholds of 185 and 150 were chosen to segment the grey matter ribbon in histological and registered MR slices, resp.). A comparison on any single similarity measure may be biased towards the underlying assumption of functional/statistical relationship. We argue that the cumulative measure represents a "consensus vote" if, on average, all individual similarity measures are consistent, that is, they may disagree on quantity but do not contradict each other qualitatively. Such cumulative similarity measure is aimed at "unbiased" and "smoothed"

comparison of cost function performances, when the ground truth evaluation, either based on reliable landmarks or on prior knowledge, is not feasible.

For a particular slice-to-volume nonrigid registration task, it is not clear a priori what is the best: (1) optimization strategy, (2) cost function, (3) deformation model, and more practically, (4) what is the best combination of the above three? These questions are left open to experimental confirmation. Therefore, the goal of "metaoptimization" is to find the best combination of "metadimensions" (a cost function, a deformation model, and an optimization strategy) within our framework. Our framework has open building blocks that may be freely combined and tested. Since it is not practical to exhaustively test all possible combinations of implemented models and algorithms, the combination is optimized in a series of "line searches" along one metadimension at a time. For example, first, performance of optimization algorithms is compared using TPS- and CC-based cost function, next performance of cost functions with various similarity measures is compared using the best optimizer from the first step and keeping TPS, then different deformation models are compared using the best cost function and optimizer, and finally, multiresolution/multistage schemes are evaluated for the best performing combinations.

3.4. Comparison of Optimizers. The first series of experiments addresses the following question: which optimization algorithms are robust and perform well in nonrigid and/or rigid registration? The goal is to identify candidates for further use in multistage/multiresolution optimization schemes and for sampling along the remaining metadimensions (cost functions, deformation models). Rigid registration is the easiest test for an optimizer, while nonrigid registration (full resolution, single-stage) is the hardest test. Rigid and nonrigid registrations were performed on our data using six optimization algorithms: NEWUOA, PDS, Simplex, Genetic, DE, and CG. All cases used CC cost function and TPS (N/A to rigid case), except for CG, which is implemented only for SSD and B-Splines. The results of rigid and nonrigid

registrations are compared across optimization algorithms in Table 1.

The registration quality is evaluated by three criteria listed in order of priority: (1) the cumulative similarity SM, (2) the bending energy E_T (N/A to rigid case), and (3) the running time. Best results should have maximal SM with constrained bending energy and a reasonably low running time. In the easiest optimization task, rigid registration, two optimizers stand out as winners: NEWUOA (total SM 39.099) and DE (total SM 39.245), but DE runs much longer (total time 125.9 min versus 1.1 min). In the hardest test, single-stage nonrigid registration, some optimizers fail to find a plausible solution: Simplex optimization does not finish in more than two days, while Genetic and DE evolutionary algorithms fail to converge to an optimum (note large bending energy at low SM values). Two algorithms, NEWUOA and PDS, manage to find a solution that is better than rigid registration, and NEWUOA results are clearly the best in terms of all three criteria. Based on these findings, NEWUOA is identified as the top performer for our slice-to-volume nonrigid registration task.

As evident from the table, different optimizers do not find the same solution. The slice-to-volume registration problem is underdetermined; despite a deformation model regularization and a smoothness constraint, there are multiple locally optimal solutions that are close to each other. Also, the results are not uniform across input slices (see columns with block # 1–8): registration of some slices takes considerably longer time. For this reason, the table includes detailed results for each block, instead of just the summary column. The results are so different not only because of differences in the size of images and control point grids, but also because of different image content. A cost function multidimensional landscape is unique in each case of a histological image; some cases are more difficult to optimize because of many local minima due to noise and local matches.

In order to simplify the nonrigid optimization task, it was divided into stages optimizing alignment and deformation DOFs separately. Each stage in a two-stage scheme employs the NEWUOA. Better results are obtained if the second stage is also allowed to optimize the translation $\vec{\tau}$ variables (it should be intuitively expected because the position of the deformed surface in 3D is fine-tuned in a second stage). A two-stage scheme works better (total SM 46.109 in Table 2 versus 45.233 in Table 1) and faster (total time 69.2 versus 139.4 min).

We conclude that the NEWUOA optimizer provides a clear advantage in our problem for both rigid and nonrigid cases; two-stage scheme provides a further improvement.

3.5. Evaluation of Cost Functions. The second series of experiments addresses the question: which cost function's similarity measure performs best in the nonrigid registration exemplified by TPS deformation and the two-stage NEWUOA optimization scheme? Experiments were performed on our data using six cost functions with the following similarity measures: CC, CR, LC, RC, NMI, and




FIGURE 2: Density of intensity distribution (solid line: histology, dashed line: MRI).

SSD. The results of cost function evaluation are presented in Table 2.

Two cost functions, CC and SSD, stand out as top performers, with SSD being slightly better in cumulative similarity but running slower on average. Even though SSD is computationally less expensive than CC, the number of cost function evaluations is typically larger with SSD compared to CC (total sum 48940 versus 38351), which indicates that SSD is a more difficult task for the optimizer. In our data, both imaging modalities have similar trimodal histograms that are closely matched by linear histogram expansion (see Figure 2); this could explain in part why CC performs well. While in theory SSD is not optimal for different imaging modalities, in practice, SSD produces good results; this justifies using analytically simple SSD expression in derivative-based optimizers, as implemented in our CG and LM.

NMI does not yield the best results, which may seem surprising in the light of NMI's popularity in cross-modality registration. This may be in part due to the "complexity" of its optimization landscape, which exhibits many "noisy" local minima and a narrower optimum (see [14]). Traditionally, these problems are rectified by a multiresolution approach, therefore the NMI cost function is revisited in Section 3.7. Another plausible explanation refers to problems with statistical consistency of NMI (see [41]), given small sample size of the 2D image. NMI becomes less reliable and loses statistical consistency if used with images of small size, because it estimates the joint distribution of intensity values from a 2D histogram, and for small images, the number of samples may be insufficient for a robust estimate. Consequently, the number of bins in the histogram becomes an important parameter. Our implementation of NMI can have 256, 128, or 64 bins in the histogram; coarser binning has shown better performance, which agrees with the statistical consistency issue.

TABLE 1: Single-stage full resolution registration results for six optimizers (NEWUOA, Powell’s Direction Set, Simplex, Genetic, Differential Evolution, and Conjugate Gradient). Rigid and nonrigid registration results are given in successive rows. Each registration result is presented by the cumulative similarity measure SM, the bending energy E_T (nonrigid only), and the running time T in minutes. Columns 1–8 are for eight slices from eight distinct blocks of the brain sample. The last column is the sum of 1–8. Values referenced in the text are in boldface.

Optimizer	Reg	Var\Blk#	1	2	3	4	5	6	7	8	\sum_1^8
NEWUOA	Rigid	SM	5.321	5.583	5.265	4.613	4.234	4.188	4.980	4.917	39.099
		T (min)	0.1	0.1	0.1	0.2	0.1	0.1	0.1	0.1	1.1
	TPS	SM	6.031	6.129	5.345	5.764	5.330	5.244	5.708	5.682	45.233
		E_T	4.4	2.5	15.3	9.6	8.1	11.8	4.2	5.0	60.9
		T (min)	0.66	7.83	40	14.3	25	25	23	3.6	139.4
PDS	Rigid	SM	4.336	4.431	3.944	3.457	3.533	4.177	4.100	4.055	32.034
		T (min)	0.2	0.2	0.6	0.5	0.5	0.6	0.4	0.2	3.3
	TPS	SM	6.026	5.867	5.628	5.218	4.984	5.414	5.253	5.425	43.814
		E_T	4	5.6	11.1	22.0	14.6	9.4	12.2	6.4	85.4
		T (min)	15.5	164.3	486.7	363	473	490.5	334.5	82.8	2410.3
Simplex	Rigid	SM	4.705	4.389	4.223	3.358	3.552	4.134	3.800	4.432	32.592
		T (min)	0.1	0.2	0.2	0.3	0.3	0.3	0.3	0.4	2.2
	TPS	SM	5.740	—	—	—	—	—	—	5.138	10.877
		T (min)	95.3	3000+	3000+	3000+	3000+	3000+	3000+	1977	20000+
Genetic	Rigid	SM	3.926	3.492	3.104	3.370	3.372	4.146	4.099	4.397	29.907
		T (min)	3.5	8.5	12.6	15.9	17.5	19.3	17.8	6.4	101.4
	TPS	SM	3.904	3.268	2.847	3.350	2.729	3.029	3.266	3.188	25.581
		E_T	141.8	1019.9	2396.0	649.5	908.2	913.8	670.3	579.2	7278.7
		T (min)	6.7	28.3	51.1	55	64	64.5	55.7	18.3	343.6
DE	Rigid	SM	5.320	5.580	5.266	4.616	4.233	4.189	4.980	5.061	39.245
		T (min)	4.2	10.5	15.0	19.9	22.6	23.1	22.0	8.6	125.9
	TPS	SM	2.318	1.733	2.353	2.581	2.610	2.605	2.724	2.318	19.241
		T (min)	8.7	35.6	64.5	64.5	82	83	71.3	24.3	433.9
CG	Rigid	SM	3.587	4.259	4.523	3.451	3.774	4.163	4.117	3.990	31.863
		T (min)	0.1	0.2	0.2	0.9	0.3	0.1	0.2	0.2	2.2
	B-SPL	SM	4.042	4.181	4.520	3.417	3.799	4.201	4.136	4.067	32.363
		T (min)	2.7	0.2	0.3	0.0	0.2	0.0	0.1	0.4	3.9
			0.2	0.4	0.2	0.7	0.5	0.5	0.3	0.6	3.3

CR, LC, and RC cost functions return good results that, however, are not superior to CC nor SSD. In case of LC and RC, the running time is significantly longer, as expected due to their computational cost. Based on these findings, we choose the “speedy” CC cost function for further tests with NEWUOA.

3.6. Comparison of Deformation Models. The next series of experiments addresses the question: which model of the deformation field ensures the optimal performance? The three deformation models in our framework are all built with smooth splines on a regular grid of control points, but they differ in underlying “physical” properties, computational costs, and (for B-splines) optimization dimensionalities. All three deformation models have been tested on our data in the context of CC cost function and NEWUOA two-stage

optimization scheme. The results are shown in Table 3, where rigid registration results are also included for baseline reference.

The SM and E_T results show a close tie between all three models, with B-Spline FFD quality being slightly better (higher similarity and lower energy) on all eight data blocks. As for the running speed, registration with TPS model is clearly the fastest (see Section 4 for discussion).

3.7. Multiresolution Optimization. The final round of experiments addresses the following questions: can a multiresolution scheme further improve the so far identified best performing combination NEWUOA/TPS/CC? Are there any other “overlooked” combinations whose performances peak in multiresolution setting? Multiresolution schemes are evaluated for four specific combinations: (1) NEWUOA

TABLE 2: Two-stage NEWUOA registration results for six cost functions (correlation coefficient, correlation ratio, local correlation, rank correlation, Maes' normalized mutual information with 64 bins, and sum of squared differences). Each registration result is presented by the cumulative similarity measure SM, the bending energy E_T , and the running time T in minutes. Columns 1–8 are for eight slices from eight distinct blocks of the brain sample. The last column is the sum of 1–8. Values referenced in the text are in boldface.

CostFunc	Var\Blk no.	1	2	3	4	5	6	7	8	Σ_1^8
CC	SM	6.056	6.124	5.941	5.814	5.349	5.261	5.688	5.876	46.109
	E_T	3.1	2.3	2.4	9.5	6.7	12.2	3.5	4.7	44.3
	T (min)	0.7	4.5	12.1	9.0	10.4	22.0	6.9	3.5	69.2
CR	SM	6.029	6.125	5.941	5.421	5.397	5.249	5.773	5.871	45.808
	E_T	4.3	2.3	4.4	18.7	11.6	24.2	8.9	9.2	83.6
	T (min)	0.5	5.5	15.9	21.8	11.2	22.1	15.8	3.3	96.1
LC	SM	5.215	6.150	5.250	5.806	5.306	5.217	5.713	5.681	44.338
	E_T	23.6	3.9	19.7	10.5	6.6	14.8	5.1	6.3	90.5
	T (min)	3.6	16.4	69.4	43.2	56.0	60.7	47.4	10.9	307.6
RC	SM	5.958	6.064	5.379	5.806	5.389	5.221	5.687	5.71	45.215
	E_T	3.3	4.2	12.7	8.1	7.4	11.4	3.4	6.5	56.9
	T (min)	0.3	9.7	42.8	14.0	24.2	30.1	31.4	3.3	155.8
NMI	SM	5.570	5.778	5.810	5.304	4.938	4.540	5.544	5.418	42.901
	E_T	6.9	5.7	7.1	8.6	6.7	5.5	6.0	5.1	51.7
	T (min)	0.3	4.6	19.3	10.8	15.9	14.2	15.9	2.4	83.4
SSD	SM	6.075	6.079	5.925	5.882	5.392	5.309	5.778	5.712	46.152
	E_T	4.1	6.2	4.3	12.2	10.6	19.7	7.7	10.3	75.1
	T (min)	0.8	6.8	13.5	16.9	15.3	24.0S	16.6	2.6	96.5

TABLE 3: Two-stage NEWUOA registration results for three deformation models (thin plate splines, Gaussian elastic body splines, and cubic B-splines). Rigid registration is also given for the reference. Each registration result is presented by the cumulative similarity measure SM, the bending energy E_T (nonrigid only), and the running time T in minutes. Columns 1–8 are for eight slices from eight distinct blocks of the brain sample. The last column is the sum of 1–8. Values referenced in the text are in boldface.

Transform	var\Blk no.	1	2	3	4	5	6	7	8	Σ_1^8
TPS	SM	6.056	6.124	5.941	5.814	5.349	5.261	5.688	5.876	46.109
	E_T	3.1	2.3	2.4	9.5	6.7	12.2	3.5	4.7	44.3
	T (min)	0.7	4.5	12.1	9.0	10.4	22.0	6.9	3.5	69.2
GEBS	SM	6.040	6.133	5.917	5.843	5.357	5.225	5.713	5.746	45.976
	E_T	3.6	2.7	3.5	9.9	8.3	14.1	4.2	4.0	50.2
	T (min)	1.3	9.3	26.0	46.6	45.9	70.3	27.8	1.7	228.8
B-SPL	SM	6.073	6.145	5.953	5.866	5.422	5.278	5.761	5.894	46.392
	E_T	2.9	2.3	2.1	9.1	7.5	12.2	4.6	3.9	44.5
	T (min)	13.3	31.0	306.7	189.3	99.1	133.1	166.5	23.5	962.5
Rigid	SM	5.321	5.581	5.265	4.613	4.234	4.188	4.979	4.917	39.098
	T (min)	0.1	0.1	0.1	0.2	0.2	0.2	0.2	0.1	1.2

with TPS and CC, (2) NEWUOA with TPS and NMI, (3) NEWUOA with B-splines and SSD, and (4) LM with B-splines and SSD. Each multiresolution scheme consists of four stages: rigid alignment at coarse resolution using NEWUOA (even in LM case), followed by three stages of nonrigid registration at progressive levels of a Gaussian pyramid (progressive smoothing with the 2D or 3D Gaussian kernel, size = 7 and $\sigma = 2.0$), from coarse to full resolution, with rotational DOFs being disabled (tests showed that

performance was worse if rotational DOFs were enabled). The results are presented in Table 4. The scheme with NEWUOA, TPS, and CC returns slightly better similarity value compared to the two-stage case (refer to Table 2), but on average runs twice as slow. The NEWUOA/TPS/NMI scheme, which was tested with the expectation that it would rectify the NMI optimization landscape, does improve the result, yet not quite reaching the level attained by CC cost function, but on average runs almost three times slower. The

TABLE 4: Registration results for four multiresolution schemes (see text for details; NEWUOA is abbreviated by NWA). Each registration result is presented by the cumulative similarity measure SM, the bending energy E_T (nonrigid only), and the running time T in minutes. Columns 1–8 are for eight slices from eight distinct blocks of the brain sample. The last column is the sum of 1–8. Values referenced in the text are in boldface.

Multires optim	var \ Blk no.	1	2	3	4	5	6	7	8	\sum_1^8
NWA/TPS/CC	SM	6.058	6.122	5.899	5.849	5.390	5.408	5.769	5.878	46.373
	E_T	2.7	2.0	2.2	8.9	7.0	9.5	5.6	4.7	42.5
	T (min)	1.8	6.9	20.1	19.2	31.9	31.2	18.7	5.4	135.2
NWA/TPS/NMI	SM	5.992	6.113	5.863	5.400	5.031	4.755	5.622	5.471	44.248
	E_T	6.2	5.4	5.3	6.5	6.1	5.3	5.3	4.7	44.9
	T (min)	1.9	19.9	66.2	31.1	41.1	32.2	32.6	8.4	233.5
NWA/B-SPL/SSD	SM	6.099	6.147	5.980	5.904	5.456	5.491	5.817	5.856	46.750
	E_T	4.3	3.2	4.0	11.7	9.7	15.8	7.9	10.7	67.3
	T (min)	18.8	160.1	254.3	322.0	317.7	533.8	284.5	95.9	1987.1
LM/B-SPL/SSD	SM	6.095	6.145	5.948	5.902	5.201	5.468	5.814	5.783	46.355
	E_T	4.3	3.1	4.4	11.6	11.7	14.9	7.8	12.6	70.4
	T (min)	3.8	20.8	91.5	57.1	26.1	95.0	57.2	14.3	365.6

third scheme NEWUOA/B-spline/SSD, which was included as a reference for the LM-based entry, reaches highest values of SM, but does it very slowly in most cases. The number of cost function evaluations was almost ten times higher than in the first case (total sum 623772 versus 68772). This is explained by larger number of control points and DOFs required for B-spline FFD, as was already mentioned in Section 2.2.4.

The last entry in the table, the LM/B-spline/SSD scheme, performs certainly better than the single-stage full resolution LM, which failed to converge in our tests, and better than the two-stage LM (not listed in the tables; total SM 44.649, total running time 67.5 min), but its running time is much longer. While it runs faster than NEWUOA/B-spline/SSD (total time 365.6 versus 1987.1 min), it is slower than the first scheme, NEWUOA/TPS/CC (total time 135.2 min). In terms of cost function and gradient/Hessian evaluations (total sums 3446 and 898, resp.), LM scheme seems to converge in fewer iterations than NEWUOA/TPS/CC (total sum 68772), but evaluations of gradient and Hessian are much more costly.

To summarize the results for our data, the best combination of image similarity (subject to deformation constraint) and registration speed is obtained with two-stage NEWUOA optimization scheme, TPS deformation model, and CC cost function. In some cases, for example, for slices in blocks 1, 6, and 8, a minor improvement can be achieved by a multiresolution scheme, without incurring a significant extra running time.

3.8. Registered Images. Examples of eight histological slices and registered MR images are given in Figures 3 and 4. Images in the middle column correspond to rigid registration with NEWUOA and CC cost function. Images in the right column are from nonrigid registration with two-stage NEWUOA, TPS, and CC cost function. Rigid registration results look rather similar to histological images on the left—clearly a pay-off from refining the initial

conditions (see Section 3.2) and choosing a good optimizer (see Table 1), since our earlier results were not quite as good. A visual inspection shows that in all slices, nonrigidly registered images have some small regions that more closely resemble histological sections. While rigid registration, which runs much faster, may seem sufficient at the level of gross morphology or white matter regions, nonrigid registration may become essential at the small detail level, for example when matching of cortical ribbons is desired.

We evaluate where and how nonrigid registration was important with the help of difference images. Although difference images are not strictly applicable to cross-modality registration, a large difference of intensities is very likely to pin-point a registration error, if intensity distributions are reasonably close. The cross-modality difference images in Figures 5 and 6 show the absolute difference between histological and registered MR images, where intensities of the latter are scaled by the factor of 1.22 to match the dynamic range of the former. The same-modality difference images in the right column show the absolute difference between rigidly and nonrigidly registered MR images, without any intensity scaling. Please recall from Figure 2 that intensity peaks of grey matter (GM) and white matter (WM) on histological images stand apart by more than approximately 100 units. Therefore, if a difference image scores a pixel with value more than 100, such a pixel is very likely to correspond to GM/WM mismatch. We computed the percentage of mismatch area on difference images thresholded above 100. Table 5 compares rigid and nonrigid registrations in terms of this measure. For all slices, the mismatch area is smaller in the case of nonrigid registration.

4. Discussion

The proposed algorithmic framework has the following advantages. (a) It allows registration of individual, sparsely




FIGURE 3: Example of registered images (block 1–4, top to bottom) (images not to scale).

spaced histological sections. (b) It models nonplanar cross-sections by a smooth surface warped in 3D space. (c) It combines 3D aligning and warping transformation in a flexible optimization scheme. While the framework appears to have a component structure that is pretty general for image registration, it is specialized for the slice-to-volume registration because (1) it offers deformation models

particularly suitable for a 2D-to-3D transformation and (2) it has an emphasis on flexible selection of cost functions and optimization strategies, for which slice-to-volume registration problems are known to be particularly demanding [12–14]. TPS and GEBS are well-known techniques in point-landmark registration, but here they are used in a novel way for parametrization of a free form (smooth)




FIGURE 4: Example of registered images (block 5–8, top to bottom) (images not to scale).

deformation in a completely different context of intensity-based registration.

While the reconstruction of a histology volume from serial sections is a well-established procedure, it requires expertise and specialized equipment and is costly and time-consuming. Individual, sparsely spaced histological sections may be more readily available from a nonspecialized histology lab; therefore, the registration method enables a wider range of histological-to-MRI evaluation and validation experiments, that otherwise could not afford a full-scale serial sectioning procedure.

TABLE 5: Comparison of rigid and nonrigid TPS registration by mismatch area (%) computed on difference images thresholded above 100.

Reg\Blk#	1	2	3	4	5	6	7	8
Rigid	0.073	0.049	0.053	0.096	0.106	0.136	0.077	0.106
TPS	0.009	0.011	0.018	0.024	0.035	0.063	0.037	0.035

Proper histological protocols and careful handling of specimen are very important to minimize the slice deformation, and in some cases rigid registration may be sufficient,

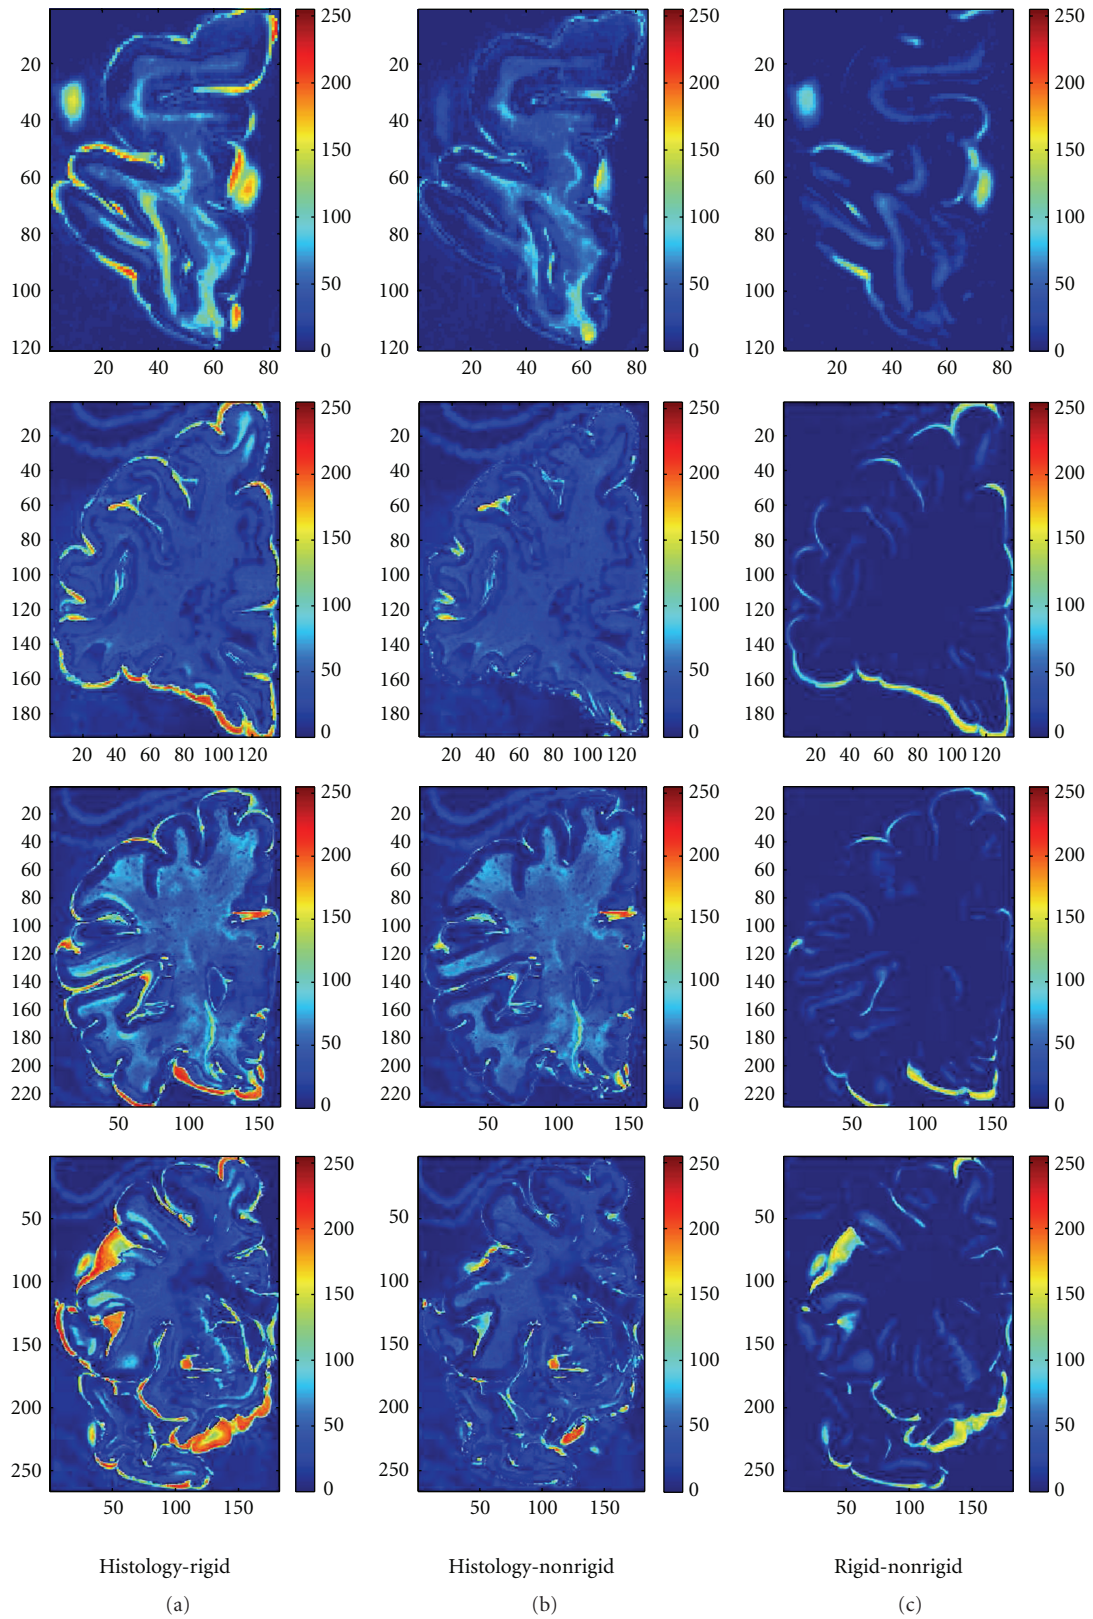


FIGURE 5: Difference images (block 1–4, top to bottom) (images not to scale).




FIGURE 6: Difference images (block 1–8, top to bottom) (images not to scale).

for example, when matching gross-morphology or regions of white matter. In most cases, however, it is desirable to recover deformations of complex local and global nature, so the modeling of nonplanar cross-sections is sought. Kim et al. [3] applied a polynomial deformation model, which captures global deformation in 3D. In their study, MRI of

the brain was performed before autopsy and fixation. Global deformations occur when the brain sample is extracted, and fixation in formalin causes global shrinkage of the specimen, therefore a global deformation model is well justified. In our framework, we also implemented a second- and a third-order polynomial models similar to those in [3], but they did not

yield any improvement over rigid registration, because in our data the postmortem MRI was performed after autopsy and fixation. Our deformation model goes a step further to account for local displacements, but it may also be easily combined with a polynomial model of global deformation, if needed.

In general, registration approaches must be tailored to a specific problem at hand: there is no single registration algorithm that suits all needs. Likewise, we cannot declare our registration problem as “solved” and “closed,” and future applications are likely to require modifications of the algorithms. In this light, flexible optimization schemes constitute an advantage of our framework, since different strategies and combinations can be tested and compared. This seems to be particularly important for slice-to-volume registration, whose performance shows greater dependence on the input images, compared to typical 3D-to-3D MRI registration.

Our finding that correlation coefficient cost function works better than other similarity measures may seem surprising and contradictory to the mainstream. After all, CC use is traditionally restricted to intramodal registration, and NMI is usually preferred in intermodal registration (although there is no proof that NMI is in any way optimal [42]). Here, we reiterate possible explanations of why CC works better in our case: (1) although registered images originate from different modalities, the correlation coefficient-based cost function may be efficiently used in cases where histograms are similar (see Figure 2); (2) the use of a NMI-based cost function with small input images suffers from statistical inconsistency; (3) CC cost function has a “smoother” optimization landscape, which is less demanding on the optimizer. The SSD cost function works well for similar reasons.

The fact that registration with TPS model is the fastest, may also appear counterintuitive. The computational cost of TPS is usually considered to be high and prohibitive for practical applications with large number of landmarks or control points [19]. However, some aspects of the computational cost may be efficiently handled by trading memory space for computation time: by iteratively reusing the decomposition of the linear system, and by tabulating RBFs at all pixel positions. This way, a speedup of factor 24 is achieved. In our experiments, TPSs are computed on grids of 100–200 control points. The computational cost may be further reduced by using approximating TPS [20], but in that case there is no closed form expression for the bending energy. An efficient computation of the TPS bending energy is one computational advantage of TPS versus GEBS, where the bending energy integral has to be computed numerically. In the case of GEBS, the basis functions can also be precomputed, therefore more complex expressions do not translate to extra running time, but we found that the TPS linear system has better numerical stability properties compared to GEBS. LU-decomposition of the linear system was used for TPS, whereas we chose to use singular value decomposition (SVD) for solving GEBS linear system (see [31]). LU-decomposition is typically faster than SVD; this also explains why registration with TPS model

is faster than GEBS. These arguments do not apply to the explanation of why B-spline model works slower, since it does not require solving any linear system, and the bending energy integral can be efficiently computed with locally supported derivatives of B-spline functions. In the case of B-spline deformation model, slower performance is mainly due to the larger dimensionality of the optimization task. It may be argued that, instead of control point displacements, the TPS or GEBS weight coefficients themselves could be used as DOFs/optimization variables: this would eliminate the need to solve a linear system for weight coefficients, and thus reduce the computational cost. Although it is equivalent to a linear transformation of the search space, nonlinear optimization problems are not invariant under such transformations. We verified experimentally with both TPS and GEBS that it would result in a much higher optimization cost and convergence issues.

Multiresolution LM optimization strategy proved to be very successful in cross-modality affine registration of 3D volumes [37], where a cost function naturally does not have any bending energy term. The combination of multiresolution approach with an LM-type optimization and cosine basis function deformation field is also known to work well in SPM toolkit nonrigid registration of smoothed MR images [40]. There, SSD-based cost function does not include bending energy penalty; instead, the deformation is regularized by Maximum a Posteriority (MAP) estimation with an ad hoc model assuming multinormal distribution of parameters. In our case, explicit bending energy term not only makes the programming of gradient and Hessian routines a tedious and time-consuming task, but also adds to the computational cost of each update iteration in LM. We recognize the fact that performance and running times of an optimization algorithm may be heavily dependent on implementation; therefore, proven state-of-the-art implementations were used in our framework whenever possible. In-house implementation of some algorithms, for example, LM and CG, might be less efficient, which could affect the running times. At the same time, the necessity of a custom implementation of these algorithms for a particular cost function and deformation model argues against them and makes the comparison fair.

The validation of registration results is problematic in the absence of a ground truth reference and reliable landmarks. Manual selection of landmarks is difficult and imprecise, especially in the case of mixed 2D and 3D images. Complementary to evaluation by a range of intensity-based similarity measures, we attempted to evaluate registration accuracy by 3D matching of automatically detected sulcal landmarks. Sulcal lines of maximal depth (*sulcal fundi* or *sulcal bottom lines*) were automatically extracted in 3D by the procedure described in [43]. A similar procedure was developed in 2D, extracting *sulcal bottom points*. The position of each bottom point was transformed in 3D according to the geometric transformation (alignment and deformation) recovered by our registration algorithm, and a 3D distance to the bottom lines was computed. For each registered slice, the maximal, the minimal, and the average distances with the standard deviation were computed. We compared these

average distances for rigid and nonrigid registrations but did not find any statistically significant difference. Since there is a clear visual distinction between rigid and nonrigid registration results, confirmed by similarity measures, one possibility is that the bottom line/bottom point landmark detection and matching method is not sensitive enough to detect it. Another possibility is that main differences are outside of sulcal fundi, which are well matched, and gyral contour or edge landmarks should be tried instead. Validation of our registration results by automatic landmark detection remains an open topic.

The described method has useful potential for various applications in mapping of brain histology to 3D imaging. Applications include (a) compiling computerized 2D-3D atlases of rodent and human brains, (b) comparing histology with MRI for a better characterization of MRI-detectable (and potentially, MRI-invisible) pathological signs, and (c) seeking for a ground truth reference to validate our neocortical parcellation results [5, 6]. This approach narrows the gap between the macroscopic scale in MRI and the microscopic scale of histological examination.

With further improvement of performance, the slice-to-volume registration framework may also find useful applications beyond histology: in areas where rigid registration is used for speed, up to this point, but where a nonrigid registration may be beneficial, for example, in planning of MR scans, or in intraoperative registration of single slices from FluoroCT, CineMR, or iMRI to 3D preoperative scans. In current implementation, nonrigid registration of a single slice takes on the order of several minutes and is likely to be too slow for intraoperative applications, but the rigid registration scheme with NEWUOA is robust and sufficiently fast, and we believe that it deserves the attention of research in the medical field.

Acknowledgments

The authors would like to thank professors Th. Arendt and M. K. Brückner (Paul-Flechsig-Institut für Hirnforschung, Leipzig) for providing the specimen and professor C. J. Wiggins (MPI für Kognitions- und Neurowissenschaften, Leipzig) for MRI acquisition, under project C15 supported by a grant from the Interdisziplinäres Zentrum für Klinische Forschung (IZKF), University of Leipzig. The authors gratefully acknowledge three anonymous reviewers for their critiques and helpful suggestions, which have led to substantial improvement of the paper.

References

- [1] L. Van den Hauwe, P. M. Parizel, J.-J. Martin, P. Cras, P. De Deyn, and A. M. A. De Schepper, "Postmortem MRI of the brain with neuropathological correlation," *Neuroradiology*, vol. 37, no. 5, pp. 343–349, 1995.
- [2] L. Bronge, N. Bogdanovic, and L.-O. Wahlund, "Postmortem MRI and histopathology of white matter changes in Alzheimer brains: a quantitative, comparative study," *Dementia and Geriatric Cognitive Disorders*, vol. 13, no. 4, pp. 205–212, 2002.
- [3] T.-S. Kim, M. Singh, W. Sungkarat, C. Zarow, and H. Chui, "Automatic registration of postmortem brain slices to MRI reference volume," *IEEE Transactions on Nuclear Science*, vol. 47, no. 4, pp. 1607–1613, 2000.
- [4] T. Schormann and K. Zilles, "Three-dimensional linear and nonlinear transformations: an integration of light microscopical and MRI data," *Human Brain Mapping*, vol. 6, no. 5-6, pp. 339–347, 1998.
- [5] F. Kruggel, M. K. Brückner, T. Arendt, C. J. Wiggins, and D. Y. von Cramon, "Analyzing the neocortical fine-structure," in *Proceedings of the Interational Conference on Information Processing in Medical Imaging (IPMI '01)*, vol. 2082, pp. 239–245, 2001.
- [6] F. Kruggel, "Techniques in analyzing the neocortical fine-structure," in *Medical Imaging Systems*, C. T. Leondes, Ed., vol. 5, pp. 255–279, World Scientific Press, 2005.
- [7] S. Eickhoff, N. B. Walters, A. Schleicher et al., "High-resolution MRI reflects myeloarchitecture and cytoarchitecture of human cerebral cortex," *Human Brain Mapping*, vol. 24, no. 3, pp. 206–215, 2005.
- [8] S. Ourselin, E. Bardinet, D. Dormont et al., "Fusion of histological sections and MR images: towards the construction of an atlas of the human basal ganglia," in *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'01)*, vol. 2208, pp. 743–751, 2001.
- [9] T. Schormann, A. Dabringhaus, and K. Zilles, "Statistics of deformations in histology and application to improved alignment with MRI," *IEEE Transactions on Medical Imaging*, vol. 14, no. 1, pp. 25–35, 1995.
- [10] M. S. Mega, S. S. Chen, P. M. Thompson et al., "Mapping histology to metabolism: coregistration of stained whole-brain sections to premortem PET in Alzheimer's disease," *NeuroImage*, vol. 5, no. 2, pp. 147–153, 1997.
- [11] E. Bardinet, S. Ourselin, D. Dormont et al., "Co-registration of histological, optical and MR data of the human brain," in *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI '02)*, vol. 2488, pp. 548–555, 2002.
- [12] B. Fei, J. L. Duerk, D. T. Boll, J. S. Lewin, and D. L. Wilson, "Slice-to-volume registration and its potential application to interventional MRI-guided radio-frequency thermal ablation of prostate cancer," *IEEE Transactions on Medical Imaging*, vol. 22, no. 4, pp. 515–525, 2003.
- [13] A. G. Chandler, T. Netsch, C. A. Cocosco, J. A. Schnabel, and D. J. Hawkes, "Slice-to-volume registration using mutual information between probabilistic image classifications," in *Progress in Biomedical Optics and Imaging—Medical Imaging*, vol. 5370 of *Proceedings of SPIE*, pp. 1120–1129, 2004.
- [14] W. Birkfellner, M. Figl, J. Kettenbach et al., "Rigid 2D/3D slice-to-volume registration and its application on fluoroscopic CT images," *Medical Physics*, vol. 34, no. 1, pp. 246–255, 2007.
- [15] M. A. Jacobs, J. P. Windham, H. Soltanian-Zadeh, D. J. Peck, and R. A. Knight, "Registration and warping of magnetic resonance images to histological sections," *Medical Physics*, vol. 26, no. 8, pp. 1568–1578, 1999.
- [16] S. Gefen, L. Bertrand, N. Kiryati, and J. Nissano, "Localization of sections within the brain via 2D to 3D image registration," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '05)*, vol. 2, pp. 733–736, 2005.

- [17] F. L. Bookstein, "Principal warps: thin-plate splines and the decomposition of deformations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, pp. 567–585, 1992.
- [18] J. Kohlrausch, K. Rohr, and H. S. Stiehl, "A new class of elastic body splines for nonrigid registration of medical images," *Journal of Mathematical Imaging and Vision*, vol. 23, no. 3, pp. 253–280, 2005.
- [19] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. Hill, M. O. Leach, and D. J. Hawkes, "Nonrigid registration using free-form deformations: application to breast mr images," *IEEE Transactions on Medical Imaging*, vol. 18, no. 8, pp. 712–721, 1999.
- [20] G. Donato and S. Belongie, "Approximate thin plate spline mappings," in *Proceedings of the European Conference on Computer Vision (ECCV'02)*, vol. 2352, pp. 21–31, 2002.
- [21] K. Rohr, H. S. Stiehl, R. Sprengel, T. M. Buzug, J. Weese, and M. H. Kuhn, "Landmark-based elastic registration using approximating thin-plate splines," *IEEE Transactions on Medical Imaging*, vol. 20, no. 6, pp. 526–534, 2001.
- [22] S. Wörz and K. Rohr, "Physics-based elastic image registration using splines and including landmark localization uncertainties," in *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI '06)*, vol. 4191, pp. 678–685, 2006.
- [23] M. Holden, "A review of geometric transformations for nonrigid body registration," *IEEE Transactions on Medical Imaging*, vol. 27, no. 1, pp. 111–128, 2008.
- [24] J. Kybic and M. Unser, "Fast parametric elastic image registration," *IEEE Transactions on Image Processing*, vol. 12, no. 11, pp. 1427–1442, 2003.
- [25] P. Thévenaz, T. Blu, and M. Unser, "Interpolation revisited," *IEEE Transactions on Medical Imaging*, vol. 19, no. 7, pp. 739–758, 2000.
- [26] M. Holden, D. L. G. Hill, E. R. E. Denton et al., "Voxel similarity measures for 3-D serial MR brain image registration," *IEEE Transactions on Medical Imaging*, vol. 19, no. 2, pp. 94–102, 2000.
- [27] D. L. G. Hill and D. J. Hawkes, "Across-modality registration using intensity-based cost functions," in *Handbook of Medical Imaging*, I. Bankman, Ed., pp. 537–553, Academic Press, 2000.
- [28] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *IEEE Transactions on Medical Imaging*, vol. 16, no. 2, pp. 187–198, 1997.
- [29] C. Studholme, D. L. G. Hill, and D. J. Hawkes, "An overlap invariant entropy measure of 3D medical image alignment," *Pattern Recognition*, vol. 32, no. 1, pp. 71–86, 1999.
- [30] T. Netsch, P. Rösch, A. Van Muiswinkel, and J. Weese, "Towards real-time multi-modality 3-D medical image registration," in *Proceedings of the 8th International Conference on Computer Vision*, vol. 1, pp. 718–725, 2001.
- [31] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C++: The Art of Scientific Computing*, Cambridge University Press, Cambridge, UK, 2002.
- [32] A. Roche, G. Malandain, X. Pennec, and N. Ayache, "The correlation ratio as a new similarity measure for multimodal image registration," in *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI '98)*, vol. 1496, pp. 1115–1124, 1998.
- [33] M. J. D. Powell, "The NEWUOA software for unconstrained optimization without derivatives," Tech. Rep. NA2004/08, Department of Applied Mathematics and Theoretical Physics, Numerical Analysis group, 2004.
- [34] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley Professional, Reading, Mass, USA, 1989.
- [35] R. Storn and K. Price, "Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces," *Journal of Global Optimization*, vol. 11, no. 4, pp. 341–359, 1997.
- [36] F. Maes, D. Vandermeulen, and P. Suetens, "Comparative evaluation of multiresolution optimization strategies for multimodality image registration by maximization of mutual information," *Medical Image Analysis*, vol. 3, no. 4, pp. 373–386, 1999.
- [37] P. Thévenaz and M. Unser, "Optimization of mutual information for multiresolution image registration," *IEEE Transactions on Image Processing*, vol. 9, no. 12, pp. 2083–2099, 2000.
- [38] N. Wiest-Daesslé, P. Yger, S. Prima, and C. Barillot, "Evaluation of a new optimisation algorithm for rigid registration of MRI data," in *Medical Imaging*, vol. 6512 of *Proceedings of SPIE*, 2007.
- [39] N. Wiest-Daesslé, S. Prima, S. P. Morrissey, and C. Barillot, "Validation of a new optimisation algorithm for registration tasks in medical imaging," in *Proceedings of the 4th IEEE International Symposium on Biomedical Imaging (ISBI '07)*, vol. 1, pp. 41–44, April 2007.
- [40] J. Ashburner and K. J. Friston, "Nonlinear spatial normalization using basis functions," *Human Brain Mapping*, vol. 7, no. 4, pp. 254–266, 1999.
- [41] A. Andronache, M. von Siebenthal, G. Székely, and P. H. Cattin, "Non-rigid registration of multi-modal images using both mutual information and cross-correlation," *Medical Image Analysis*, vol. 12, no. 1, pp. 3–15, 2008.
- [42] D. L. G. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes, "Medical image registration," *Physics in Medicine and Biology*, vol. 46, no. 3, pp. R1–R45, 2001.
- [43] G. Lohmann and F. Kruggel, "Extracting lines of maximal depth from MR images of the human brain," in *Proceedings of the International Conference on Pattern Recognition (ICPR '96)*, vol. 3, pp. 518–522, 1996.