

Medical Image Registration Using Mutual Information

FREDERIK MAES, DIRK VANDERMEULEN, AND PAUL SUETENS

Invited Paper

Analysis of multispectral or multitemporal images requires proper geometric alignment of the images to compare corresponding regions in each image volume. Retrospective three-dimensional alignment or registration of multimodal medical images based on features intrinsic to the image data itself is complicated by their different photometric properties, by the complexity of the anatomical objects in the scene and by the large variety of clinical applications in which registration is involved. While the accuracy of registration approaches based on matching of anatomical landmarks or object surfaces suffers from segmentation errors, voxel-based approaches consider all voxels in the image without the need for segmentation. The recent introduction of the criterion of maximization of mutual information, a basic concept from information theory, has proven to be a breakthrough in the field. While solutions for inpatient affine registration based on this concept are already commercially available, current research in the field focuses on interpatient nonrigid matching.

Keywords—Image registration, medical imaging, mutual information.

I. INTRODUCTION

With current medical imaging modalities, such as computed tomography (CT), magnetic resonance imaging (MRI), or positron emission tomography (PET), it is possible to routinely and minimally invasively acquire three-dimensional (3-D) images of the internal organs of the human body. These radiological images provide detailed information about anatomy and function of the imaged organs and play an essential and crucial role in the medical decision

Manuscript received July 15, 2002; revised March 17, 2003. This work was supported in part by the F.W.O.-Vlaanderen under Grants G.0340.98 and G.0258.02 and in part by the Research Fund K.U. Leuven under Grants GOA/99/05 (VHS+) and IDO/99/005. The work of F. Maes is supported by the Fund for Scientific Research—Flanders, Belgium (F.W.O.—Vlaanderen).

The authors are with the Laboratory for Medical Image Computing (Radiology-ESAT/PSI), Faculties of Medicine and Engineering, University Hospital Gasthuisberg, B-3000 Leuven, Belgium (e-mail: Frederik.Maes@uz.kuleuven.ac.be).

Digital Object Identifier 10.1109/JPROC.2003.817864

process in diagnosis, therapy planning, and assessment. The classical representation of medical images, printed on radiological film and visualized using a light box, makes image interpretation necessarily subjective and qualitative. However, most modern medical image acquisition systems generate digital images that can be processed by a computer and transferred over computer networks. Digital imaging allows to extract objective, quantitative parameters from the images by image analysis. Medical image analysis exploits the numerical representation of digital images to develop image processing techniques that facilitate computer-aided interpretation of medical images.

The continuing advancement of image acquisition technology and the resulting improvement of radiological image quality have led to an increasing clinical need and physician's demand for quantitative image interpretation in routine practice, imposing new and more challenging requirements for medical image analysis. Typical applications include for instance volumetry of organs or lesions using CT or MRI, morphometry of the brain using MRI, correlation of anatomical information from MRI with functional information from PET, or the planning of therapeutic interventions such as surgery or radiotherapy using CT images. However, accurate quantification of structural and physiological parameters from medical images is often quite difficult, due to the 3-D nature of the problem and due to limitations intrinsic to the imaging process itself, such as insufficient contrast, limited resolution, noise, inhomogeneities, or artifacts. Ambiguity in the interpretation of the images introduces inter- and intraobserver variability in the measurements that may be of the same order of magnitude as the parameter to be quantified itself. This variability affects the reproducibility of the measurements and undermines the significance of the clinical findings derived from them. There is, therefore, great need for more automated methods for computer-aided image interpretation to improve the reliability of the measurements.

A fundamental problem in medical image analysis is the integration of information from multiple images of the same

subject, acquired using the same or different imaging modalities and possibly at different time points. One essential aspect thereof is image registration, i.e., recovering the geometric relationship between corresponding points in multiple images of the same scene. While various more or less automated approaches for image registration have been proposed in the field of medical imaging and image analysis, one strategy in particular, namely maximization of mutual information (MMI), has been extremely successful at automatically computing the registration of 3-D multimodal medical images of various organs from the image content itself. Mutual information (MI) is a basic concept from information theory, that is applied in the context of image registration to measure the amount of information that one image contains about the other. The MMI registration criterion postulates that MI is maximal when the images are correctly aligned. MMI has been demonstrated to be a very general and powerful criterion, that can be applied automatically and very reliably, without prior segmentation or preprocessing, on a large variety of applications. This makes the method highly suited for routine use in clinical practice.

This paper presents a review of the MMI registration criterion and its application for 3-D medical image registration. We start with a discussion of the need and requirements for 3-D medical image registration in Section I-A. Section I-B reviews previous work on intensity-based image registration and situates the MMI criterion within the state-of-the-art at the time of its introduction in the field. The theory and implementation of the MMI registration criterion are presented in Section II. The concept of mutual information is introduced in Section II-A and the MMI registration criterion is formulated in Section II-B. This section also discusses some defective aspects of the assumptions underlying MMI, which are known to result in registration failures in some particular cases. The MMI registration algorithm as originally presented by Collignon *et al.* [17] is outlined in Section II-C. Alternative implementations, other MI related registration measures and extensions of the MMI approach are discussed in Section II-D.

Validation of the MMI registration criterion is discussed in Section III. Section III-A reviews the validation of the robustness of the MMI criterion with respect to implementation issues, such as image sampling, intensity interpolation, and optimization strategy, and with respect to image related aspects, such as image degradation and partial overlap of the registered images. The accuracy of the MMI criterion has been validated as part of the Retrospective Registration Evaluation Project (RREP) conducted by Fitzpatrick *et al.* [50] at Vanderbilt University, Nashville, TN, by comparison with external marker based registration. The results of this study, which are discussed in Section III-B, demonstrate the subvoxel accuracy of MMI for registration of CT, MR, and PET images of the brain. In Section IV, we illustrate the clinical relevance of the MMI registration criterion for 3-D affine image registration in a wide range of applications involving CT, MR, and PET images of the head, the thorax, and the abdomen, including: registration of MR time series images of the brain for the study of lesion evolution over time in mul-

tle sclerosis (MS); staging of metastatic mediastinal lymph nodes from CT and PET images in lung cancer patients; and planning of radiotherapy treatment of brain and prostate tumors from CT and MR images. All these registration applications can be handled by a single software tool using a single registration algorithm, without having to tune it to each specific application at hand.

While affine image registration using MMI is well established and used already in routine clinical practice, extension of the MMI criterion to nonrigid image registration is still an active area of research. The state-of-the-art of nonrigid image registration using MMI is reviewed in Section V.

In this paper, we focus on the application of MMI for registration of 3-D image volumes. MMI has also been applied for other registration problems, such as two-dimensional (2-D)/3-D image registration [16], [78], registration of surface models to 2-D and 3-D images [122], or registration in other contexts, such as microscopy [10] or histology [55]. For an extensive survey of MMI registration applications we refer to [83].

A. Image Registration Problem

Different imaging modalities, such as CT, MRI, and PET, are based on different physical principles and capture different and often complementary information. Many applications in clinical practice benefit from an integrated visualization and combined analysis of such multimodal images. In radiotherapy planning, for instance, computation of the radiation dose distribution around the target volume is based on CT, while the lesion to be radiated itself can often better be discriminated and defined in the corresponding MRI scan. In neurology, functional information about the brain derived from PET needs to be correlated with anatomical information from MRI. For the diagnosis of lung cancer, tumors can be easily detected in PET, but a CT scan is required for anatomical localization, etc. In many applications, it is also necessary to compare images acquired at a different time point, such as: dynamic sequences of 3-D images, for instance, for functional analysis of the heart; followup studies, for instance, for diagnosis of evolving disease processes such as MS or for assessing the outcome of therapy; or pre- and post-operatively acquired images, for instance, for validation of a surgical intervention. Analysis of a single scene, i.e., the patient, from multiple images assumes that the geometrical correspondence or registration between these images is known, such that from the coordinates of a point in one image the coordinates of the geometrically corresponding point in each of the other images can be obtained. Anatomically identical points can then be precisely located and compared in each of the images. The principle of image registration is illustrated in Fig. 1.

While in some cases, the images may be assumed to be aligned at acquisition, such as simultaneously acquired proton density (PD) and T2 weighted MR images or PET transmission and emission images acquired without the patient leaving the scanner, images acquired by different scanners or at different time points are usually acquired




Fig. 1. Principle of 3-D image registration illustrated on CT and MR brain images of the same subject. (a) Original CT slice. (b) Original MR slice. (c) Slice through the reformatted CT image volume that geometrically corresponds to image (b) after 3-D registration to the MR image. Note the rotation around the vertical axis. Points at the same position in both images are anatomically identical and can immediately be compared. Both images offer complementary information—CT shows primarily bony structures, i.e., the skull, while MRI allows better differentiation of the internal soft tissue brain structures. A lesion is clearly visible in the MR image.

independently one of another and their relative position is generally unknown, unless specific provisions were made prior to acquisition. In stereotactic brain surgery, for instance, a reference frame is rigidly attached to the skull of the patient prior to image acquisition. The frame holds a set of markers with known geometry and filled with an appropriate contrast medium, such that they are clearly visible as bright spots and easy to detect automatically in the images. From the position of all markers in each of the image slices, the transformation from image coordinates to frame coordinates, and, hence, to patient coordinates and vice versa, can be computed. If multiple images were acquired, the location of corresponding frame markers in each of the images allows to compute the registration of one image to another (see [62] and references therein). The accuracy of external frame based registration is usually very good and of the order of 1–2 mm, provided that registration errors induced by scanner miscalibration or by geometric distortion in the MR images have been minimized by proper calibration and appropriate selection of scan protocol parameters, respectively. Less invasive and more patient friendly skin markers have been developed, but these are less accurate than bone-implanted or skull-fixated markers due to the nonrigid nature of the skin. However, external marker based registration can not be applied in all applications, such as registration of CT and MR images of the abdomen for instance, and does not allow retrospective registration of images acquired without markers attached to the patient. Retrospective registration requires procedures that allow to recover the registration transformation from the image content itself.

With the traditional film-based representation of 3-D image volumes as a sequence of 2-D cross-sectional slices, the clinician mentally has to construct a picture of the 3-D geometric relationship between the different sets of slices when jointly interpreting multiple image volumes, using prior anatomical knowledge to combine the geometrical clues provided by corresponding anatomical landmarks visible in each of the images separately. It is clear that when one wants to exploit the benefits of digital imaging and medical image processing for quantitative analysis of

multimodal images, a formal computer-based registration procedure is required. The registration may be performed manually by the user tuning the registration parameters assisted by visualization software that supplies visual feedback of the quality of the registration [75], [79], [98]. While such subjective approach may be sufficient to support clinical decisions in some applications, a more objective registration measure is needed when higher accuracy is required, as in brain surgery for instance, or if clear anatomical references are missing or uncertain in one of the images, as when staging mediastinal tumors from PET and CT images, or when one wants to exploit the benefits of digital image processing for quantitative analysis, as for instance in the followup of MS lesion load over time from MRI. The retrospective registration of 3-D multimodal images is, therefore, a fundamental task in medical image processing. For any image registration tool to be useful and successful in clinical practice, it needs to be sufficiently accurate and reliable in all cases and the effort and time required by the user should be minimal. These requirements can best be met by automating the registration process as much as possible.

Retrospective registration of 3-D medical images for clinical applications is not a trivial problem, due to the sometimes large differences in patient positioning in the scanner, the different image content and resolution of the images, the possible presence of modality specific image degradations and artifacts and of pathology induced distortions. Moreover, patient related changes over time may have occurred between both acquisitions, for instance because of a surgical intervention. Another difficulty is the large variety of different applications in which image registration is important. These do not only involve radiological images of the head and the brain, but also of other organs, such as the thorax, the abdomen, or the pelvic region. Each of these applications poses specific problems the registration method should be able to cope with, such as for instance breathing induced motion artifacts in case of the thorax or differences in bladder filling in case of the pelvis. In many applications, local nonrigid tissue deformations are negligible or irrelevant and the geometric relationship between the images to be registered can be modeled by a rigid or affine linear transformation, composed of a 3-D translation and rotation and possibly also 3-D scaling and skew. The registration problem then consists of determining the 6 or 12 parameters of the rigid or affine geometrical transformation that correctly aligns both images. In other applications, it may be needed to correct for local nonrigid image distortions, for instance to compensate for breathing induced deformations in the thorax, or to quantify local morphological differences from similar images of different subjects, for instance, for building geometrical models of the brain. In these cases, a more general nonrigid deformation model, with typically much more degrees of freedom, needs to be used.

B. Voxel Based Strategies for Image Registration

For an extensive overview and classification of medical image registration techniques, we refer to Maintz [62], extending earlier overviews of [7], [70], [117]. Registration

procedures can be classified according to their registration paradigm [62], i.e., the image features and the correspondence criterion used to compute the registration. Registration features are either extrinsic, i.e., relating to external objects or markers introduced in the imaged space specifically for registration purposes, or intrinsic, i.e., based on the image information generated by the object under study itself. Extrinsic registration can not be applied retrospectively. Intrinsic methods can be classified as either point based, surface based or voxel based. *Point based* methods rely on manually indicated anatomical landmarks or automatically computed geometrically salient points. Although very versatile, these methods are labor-intensive if user interaction is required and their accuracy relies on the accurate localization of a sufficient number of corresponding landmarks in all modalities. *Surface based* registration requires delineation of corresponding surfaces in each of the images separately. However, surface segmentation algorithms are generally highly data and application dependent and difficult to automate, and surfaces are not easily identified in functional modalities such as PET. *Voxel based* registration methods optimize a functional measuring the similarity of all geometrically corresponding voxel pairs for some feature. Their main advantage is that feature calculation is straightforward or even absent when only grey-values are used, such that accuracy is not limited by segmentation errors as in point or surface based methods. By extending the feature space from points to surfaces to all voxels in the image, knowledge of point-to-point correspondences is traded for a larger set of samples that is accounted for by the similarity measure.

For intramodality registration, multiple voxel based methods have been proposed that optimize some global measure of the absolute difference between corresponding voxel intensities within overlapping parts or in a region of interest [13], [32], [90], [121]. These criteria all rely on the assumption that the intensities of the two images are linearly correlated, which is generally not satisfied for intermodality registration. Cross-correlation of feature images derived from the original image data has been investigated using geometrical features such as edges [64] and ridges [116] or using especially designed intensity transformations [115]. However, feature extraction may introduce new geometrical errors, requires extra calculation time, and may be modality or application specific, which makes automation more difficult. Moreover, correlation of sparse features like edges or ridges may have a very peaked optimum at the registration solution, but at the same time be rather insensitive to misregistration at larger distances, as all nonedge or nonridge voxels correlate equally well. A multiresolution optimization strategy is, therefore, required, which is not necessarily a disadvantage, as it can be computationally attractive.

In the approach of Woods *et al.* [129] and Hill *et al.* [40], [42], misregistration is measured by the dispersion of the scatter-plot or 2-D joint histogram of image intensities of corresponding voxel pairs, which is assumed to be minimal in the registered position. However, the dispersion measures they propose are largely heuristic. Hill's criterion

requires segmentation of the images or delineation of specific histogram regions to make the method work [103], while Woods' criterion is based on additional assumptions concerning the relationship between the grey-values in the different modalities, which reduces its applicability to some very specific multimodality combinations (PET/MR). Collignon *et al.* [18], inspired by the work of Hill and Woods, recognized the information theoretic nature of the image registration problem and proposed the much more general notion of entropy of the grey-level scatter-plot as a new matching criterion. However, this measure is sensitive to partial overlap of the images, as it does not account for the fact that the information content of each of the images separately within their region of overlap may not be constant during registration.

Finally, two different groups, Collignon and Maes *et al.* [17], [59] at KU Leuven, Belgium, and Viola and Wells *et al.* [122], [124] at the Massachusetts Institute of Technology, Cambridge, almost simultaneously but independently of each other, introduced MMI of image intensities as a new registration criterion. MI, or relative entropy, is a basic concept from information theory, which can be considered a nonlinear generalization of cross-correlation. MI measures the statistical dependence between two random variables or the amount of information that one variable contains about the other [20]. The MMI registration criterion postulates that the MI of the image intensity values of corresponding voxel pairs is maximal if the images are geometrically aligned. The information theoretical relationship of MI with the registration criteria of Hill and Woods is examined in [59]. The mathematical elegance, algorithmic simplicity and spectacular robustness of the MMI registration criterion immediately attracted a lot of interest from the research community. Related early work in this area includes the work by Studholme *et al.* [103], Pluim [82], and Meyer *et al.* [51].¹

II. THEORY OF IMAGE REGISTRATION USING MMI

A. Mutual Information

Two discrete random variables A and B with marginal probability distributions $p_A(a)$ and $p_B(b)$ and joint probability distribution $p_{AB}(a, b)$ are statistically independent if $p_{AB}(a, b) = p_A(a) \cdot p_B(b)$, while they are maximally

¹Historical footnote: Collignon and Maes formulated the MMI registration criterion early 1995. It was presented orally by Collignon at the CVRMED'95 conference in Nice, France, in April of that year [18], and first published in the proceedings of the IPMI'95 conference in Brest, France, in June [17]. Almost simultaneously, Viola and Wells presented their own formulation of MMI at ICCV'95 in Cambridge, MA, also in June [122]. Hawkes, Hill, and Studholme, at UMDS, London, U.K., already working on histogram dispersion measures for image registration [42], got informed about MMI by personal communication with Collignon and Maes during mutual visits between both groups in January and March, 1995. At IPMI'95, Collignon's enthusiasm about the new registration measure was supported by Studholme's comparison of MMI with other voxel-based registration measures, which was presented at that same conference [103]. The group of Viergever, Maintz, and Pluim at ISI, Utrecht, the Netherlands, started working on MMI in 1996 using software provided by Maes [82]. Meyer attended Collignon's presentation at CVRMED'95. His group was the first to apply MMI for nonrigid matching [51].

dependent if they are related by a one-to-one mapping T : $p_A(a) = p_B(T(a)) = p_{AB}(a, T(a))$. The mutual information $I(A, B)$ of A and B measures the degree of dependence of A and B as the distance between the joint distribution $p_{AB}(a, b)$ and the distribution associated to the case of complete independence $p_A(a) \cdot p_B(b)$, by means of the Kullback-Leibler measure [20], [114], i.e.,

$$I(A, B) = \sum_{a,b} p_{AB}(a, b) \log \frac{p_{AB}(a, b)}{p_A(a) \cdot p_B(b)}. \quad (1)$$

Mutual information is related to the information theoretic notion of entropy by the following equations:

$$I(A, B) = H(A) + H(B) - H(A, B) \quad (2)$$

$$= H(A) - H(A|B) \quad (3)$$

$$= H(B) - H(B|A) \quad (4)$$

with $H(A)$ and $H(B)$ being the entropy of A and B respectively, $H(A, B)$ their joint entropy, and $H(A|B)$ and $H(B|A)$ the conditional entropy of A given B and of B given A , respectively. $H(A)$, $H(A|B)$ and $H(A, B)$ are defined as

$$H(A) = - \sum_a p_A(a) \log p_A(a) \quad (5)$$

$$H(A|B) = - \sum_{a,b} p_{A,B}(a, b) \log p_{A|B=b}(a) \quad (6)$$

$$H(A, B) = - \sum_{a,b} p_{A,B}(a, b) \log p_{A,B}(a, b) \quad (7)$$

with $p_{A|B=b}(a)$ the conditional probability of A given $B = b$. The entropy $H(A)$ is known to be a measure of the amount of uncertainty about the random variable A , while $H(A|B)$ is the amount of uncertainty left in A when knowing B . Hence, from (3), $I(A, B)$ is the reduction in the uncertainty of the random variable A by the knowledge of another random variable B , or, equivalently, the amount of information that B contains about A . If A and B are independent, $p_{AB}(a, b) = p_A(a) \cdot p_B(b)$ and $I(A, B) = 0$, while if A and B are one-to-one related, $I(A, B) = H(A) = H(B)$. It can be shown [20] that mutual information is nonnegative ($I(A, B) \geq 0$) for any two random variables A and B . Some properties of mutual information are summarized in Table 1 (see [114] for their proof).

B. Mutual Information Registration Criterion

Multimodal images of the same scene represent measurements of different properties of the objects in that scene. Although the image intensities corresponding to the same object may be very different between different modalities, in general they are not independent observations as the underlying physical reality, i.e., the objects or tissues, are the same. The intensity values in different images of the same scene at image positions that correspond to the same location in 3-D physical space, i.e., the same volume element of the same

Table 1
Some Properties of Mutual Information

Independence:	$I(A, B) = 0 \Leftrightarrow p_{AB}(a, b) = p_A(a) \cdot p_B(b)$
Symmetry:	$I(A, B) = I(B, A)$
Self information:	$I(A, A) = H(A)$
Lower bound:	$I(A, B) \geq 0$
Upper bound:	$I(A, B) \leq \min(H(A), H(B))$ $\leq (H(A) + H(B))/2$ $\leq \max(H(A), H(B))$ $\leq H(A, B)$ $\leq H(A) + H(B)$
Data processing:	$I(A, B) \geq I(A, T(B))$

object, are not independent quantities, but are statistically related measurements. Knowledge of the outcome of one measurement provides some information about the underlying physical reality from which it was obtained and, therefore, reduces the uncertainty about the outcome expected from other measurements of that same reality. This is intuitively clear for modalities that capture information which is linked directly to anatomy, such as X-ray attenuation in CT or water content in MR, but also holds for modalities that represent different kinds of information, such as anatomy in MR and function in PET.

If \mathcal{A} and \mathcal{B} are two images that are geometrically related by the registration transformation \mathbf{T}_α with parameters α such that voxels \mathbf{p} in \mathcal{A} with intensity a physically correspond to voxels $\mathbf{T}_\alpha(\mathbf{p})$ in \mathcal{B} with intensity b , the statistical dependence between a and b or the information that one value contains about the other is measured by the mutual information $I(A, B)$ of the variables $A = \{a\}$ and $B = \{b\}$

$$a = \mathcal{A}(\mathbf{p})$$

$$b = \mathcal{B}(\mathbf{T}_\alpha(\mathbf{p}))$$

$$I(A, B) = \sum_{a,b} p_{AB}(a, b) \log_2 \frac{p_{AB}(a, b)}{p_A(a) \cdot p_B(b)}$$

with $p_{AB}(a, b)$, $p_A(a)$ and $p_B(b)$ the joint and marginal distributions of the pair (a, b) and of a and b , respectively. Estimates for these distributions can be obtained by simple normalization of the joint and marginal histograms of the overlapping parts of both images. The relationship $p_{AB}(a, b)$ between a and b and, hence, their mutual information $I(A, B)$ depends on \mathbf{T}_α , i.e., on the registration of the images. The mutual information registration criterion postulates that the images are geometrically aligned by the transformation \mathbf{T}_{α^*} for which $I(A, B)$ is maximal

$$\alpha^* = \arg \max_{\alpha} I(A, B).$$

If both marginal distributions $p_A(a)$ and $p_B(b)$ can be considered to be independent of the registration parameters α , the MI criterion reduces to minimizing the joint entropy

$H_{AB}(A, B)$ [18]. If either $p_A(a)$ or $p_B(b)$ is independent of α , which is the case if one of the images is always completely contained in the other, the MI criterion reduces to minimizing the conditional entropy $H(A|B)$ or $H(B|A)$. However, if both images only partially overlap, which is very likely during optimization, the volume of overlap will change when α is varied and both marginal distributions $p_A(a)$ and $p_B(b)$ and, therefore, also their entropies $H(A)$ and $H(B)$ will in general depend on α . The MI criterion takes this into account explicitly, as becomes clear in (2), which can be interpreted as follows [122]: “Maximizing mutual information will tend to find as much as possible of the complexity that is in the separate datasets (maximizing the first two terms) so that at the same time they explain each other well (minimizing the last term).”

For $I(A, B)$ to be useful as a registration criterion and well-behaved with respect to optimization, $I(A, B)$ should vary smoothly as a function of misregistration $|\alpha - \alpha^*|$. This requires $p_A(a)$, $p_B(b)$ and $p_{AB}(a, b)$ to change smoothly when α is varied, which will be the case if the image intensity values are spatially correlated [59]. Although the formulation of the MI criterion suggests that spatial dependence of image intensity values is not taken into account, such dependence is in fact essential for the criterion to be well-behaved around the registration solution.

The MMI registration criterion is illustrated in Fig. 2 for the CT and MR brain images of Fig. 1. Fig. 2 shows the 2-D histogram of the image intensity values in a nonregistered and in the registered position. If the images are properly aligned, the high-intensity values in the histogram of the CT image originating from the bone of the skull are most likely to correspond to low-intensity values in the histogram of the MR image, resulting in a peak in their 2-D joint histogram. The uncertainty about the MR voxel intensity is, thus, largely reduced if the corresponding CT voxel is known to be of high intensity. This correspondence is lost in case of misregistration, resulting in a much more dispersed histogram. The mutual information between both images is, therefore, larger in the registered than in a nonregistered position. Fig. 3 shows traces of MI of the images of Fig. 1 for translation and rotation around the registered position. These illustrate the robustness and reliability of the MI criterion: within a large range of -25 to $+25$ mm translation and -25 to $+25^\circ$ rotation around the registered position there is only a single strong optimum, which coincides with the correct registration solution.

Mutual information does not rely on the intensity values directly to measure correspondence between different images, but on their relative occurrence in each of the images separately and co-occurrence in both images combined. As such it is insensitive to intensity permutations or one-to-one intensity transformations and is capable of handling positive and negative intensity correlations simultaneously. Unlike other voxel-based registration criteria, based on for instance intensity differences or intensity correlation, the MI criterion does not make limiting assumptions about the nature of the relationship between the image intensities of corresponding voxels in the different modalities, which is highly




Fig. 2. Joint histogram of the volume of overlap of the CT and MR images of Fig. 1: (a) before registration, $I(CT, MR) = 0.52$; (b) after registration, $I(CT, MR) = 0.86$. Misregistration was about 15 mm translation and 10° rotation.

data dependent, and does not impose constraints on the image content of the modalities involved. This explains the success of MMI for multimodal image registration in a wide range of applications involving various modality combinations, some of which are illustrated in Section IV.

While originally introduced for intermodality image registration, MMI has also been used successfully for intramodality image registration, for instance of serial MR images [43], fMRI time-series images [27], or SPECT perfusion images [89]. In some studies, unimodal similarity measures, relying on minimizing intensity differences or maximizing intensity correlation, or measures specifically designed for particular modality combinations were found to perform equally well or better than MMI [4], [78], [110]. However, the underlying assumptions of these methods, especially the assumption that the intensities of corresponding voxels in unimodal images to be registered are identical, may be violated in practice, for instance due to imaging artifacts such as noise or intensity inhomogeneity, nonrigid tissue deformations not compensated for by affine image registration, or local contrast changes, for instance as a result of perfusion or brain activation. A statistical multimodal registration criterion such as MMI may,

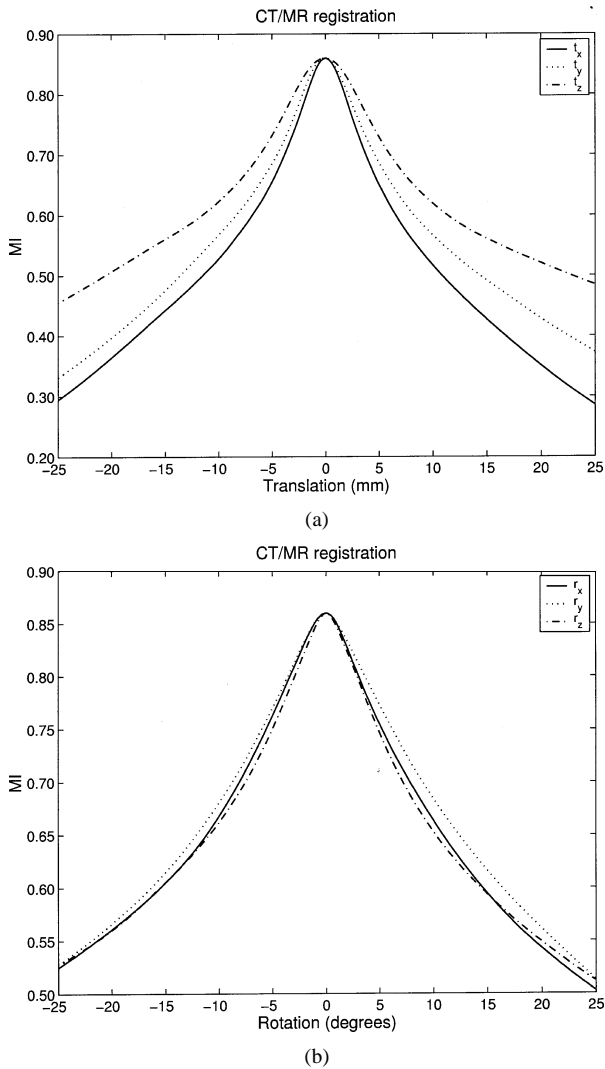


Fig. 3. Traces of MI of the CT and MR images of Fig. 1 for (a) translation and (b) rotation around the registered position.

therefore, be preferred for unimodal registration as well [43], [89].

Nevertheless, there are cases in which MMI fails as a registration criterion. Such failures occur due to insufficient mutual information in the images, ambiguity about the intensity relationship between both images if this is not spatially invariant, or inability to reliably estimate MI if the number of image samples is small.

The fundamental assumption of MMI based registration that image intensities in both images are related to corresponding objects that should be aligned by registration, may be invalid if the information in both images is very different. A typical example is the co-registration of anatomical information from CT and MR with functional information from PET. MMI-based registration of CT and MR images of the brain to the corresponding PET brain images has been demonstrated to be feasible with subvoxel accuracy [126], because in the brain the rigid body assumption is usually well satisfied and the PET images contain sufficient anatomical clues for the registration to be successful. However, such clues are much less apparent

in PET emission images of other body regions, such as the abdomen or the thorax, which are blurred by breathing induced motion averaging and which typically only show hot spots of functional activity against a fairly homogeneous background as illustrated in Fig. 13. Consequently, correct alignment of CT and PET emission images as in Fig. 13 may not coincide with a global or even local maximum of MI and robust registration of thorax CT and PET emission images directly using MMI is not feasible. Instead, PET transmission images, aligned with the emission images and showing similar anatomical content as the CT images but at much lower resolution, have been used as intermediary as described in Section IV-D [120].

Since MI is computed from the joint intensity probability of both images that is estimated by pooling contributions from everywhere in the image domain, the MMI criterion implicitly assumes that the statistical relationship between corresponding voxel intensities is identical over the whole area of overlap. However, the photometric relationship between two multimodal images of the same scene may not be spatially invariant if one of the images suffers from severe intensity inhomogeneity. In particular, MMI is not well adapted for registration involving images that show large shading artifacts, a typical example being MR images acquired using surface coils [99], [107].

The MMI registration criterion assumes that the joint probability distribution of corresponding voxel intensities can be estimated reliably at and near the registration solution. In practice, this requires the volume of overlap at registration to contain a sufficiently large number of voxels. For low-resolution images or if the region of overlap is small, the statistical relationship between both images needs to be derived from a small number of samples, which is not robust. In these cases, the computed MI may show multiple local optima around the correct registration solution or the registered position may not coincide with a local maximum of MI [87], [106].

C. Registration Algorithm and Implementation Issues

We first describe the MMI registration algorithm as formulated by Collignon and Maes *et al.* [17], [59]. Alternative formulations, implementations, and extensions are discussed in Section II-D.

The MMI registration criterion does not require any preprocessing or segmentation of the images. With each of the image volumes is associated a 3-D world coordinate frame in millimeter units with its origin at the center of the volume and its axes directed left/right, anterior/posterior, and inferior/superior relative to the patient. The orientation of the image can be retrieved automatically from the information in the image header recorded by the scanner. Voxel indices in each image are translated into coordinates in millimeter by taking the pixel size and the interslice spacing of each image into account. Anisotropy in voxel size and differences in voxel size between both images are, thus, accounted for without resampling the images first. One of the images to be registered is selected to be the *floating* image \mathcal{F} from which samples $s \in S$ are taken and transformed by the geometric transformation \mathbf{T}_α with parameters α into the *reference* image \mathcal{R} .




Fig. 4. Graphical illustration of NN, TRI, PV, and PI interpolation in 2-D. NN and TRI interpolation find the reference image intensity value at position $T_{\alpha} s$ and update the corresponding joint histogram entry, while PV and PI interpolation distributes the contribution of this sample over multiple histogram entries using weights that are smooth functions of the registration parameters α .

S may include all voxels in \mathcal{F} or a subset thereof to increase speed performance.

The joint image intensity histogram $h_{\alpha}(f, r)$ of the volume of overlap $s \in S_{\alpha} \subset S$ of \mathcal{F} and \mathcal{R} is constructed by binning the image intensity pairs $(f(s), r(\mathbf{T}_{\alpha}(s)))$ for all $s \in S_{\alpha}$. In order to do this efficiently, the floating and the reference image intensities are first linearly rescaled to the range $[0, n_{\mathcal{F}} - 1]$ and $[0, n_{\mathcal{R}} - 1]$ respectively, $n_{\mathcal{F}} \times n_{\mathcal{R}}$ being the total number of bins in the joint histogram, with typically $n_{\mathcal{F}} = n_{\mathcal{R}} = 256$. In general, $\mathbf{T}_{\alpha}(s)$ will not coincide with a grid point of \mathcal{R} and interpolation of the reference image is needed to obtain the image intensity value $r(\mathbf{T}_{\alpha}(s))$. Several interpolation schemes are graphically depicted in Fig. 4 [59]. Nearest neighbor (NN) interpolation of \mathcal{R} is generally insufficient to guarantee subvoxel accuracy, as it is insensitive to translations up to 1 voxel. Other interpolation methods, such as trilinear (TRI) or higher order interpolation, may introduce new intensity values which are originally not present in the reference image, leading to unpredictable changes in the marginal distribution $p_{\mathcal{R}, \alpha}(r)$ of the reference image for small variations of α . To avoid this problem, an alternative approach, namely trilinear partial volume distribution (PV) interpolation, was proposed in [17] to update the joint histogram for each voxel pair $(s, T_{\alpha} s)$. Instead of interpolating new intensity values in \mathcal{R} ,

the contribution of the image intensity $f(s)$ of the sample s of \mathcal{F} to the joint histogram is distributed over the intensity values of all eight NNs of $T_{\alpha} s$ on the grid of \mathcal{R} , using the same weights as for trilinear interpolation. Each entry in the joint histogram is then the sum of smoothly varying fractions of 1, such that the histogram changes smoothly as α is varied. PV interpolation results in a continuous and a.e. differentiable registration criterion. A variant thereof, trilinear partial intensity distribution (PI) interpolation, was proposed in [58]: the image intensity value $r(\mathbf{T}_{\alpha}(s))$ is computed first by trilinear interpolation and the contribution of the sample s to the joint histogram is distributed over the bins $(f(s), \lfloor r \rfloor)$ and $(f(s), \lceil r \rceil)$ with weights $\lceil r \rceil - r$ and $r - \lfloor r \rfloor$, respectively, with $\lfloor r \rfloor$ the largest integer that is not larger than r and $\lceil r \rceil = \lfloor r \rfloor + 1$.

Estimations for the marginal and joint image intensity distributions $p_{\mathcal{F}, \alpha}(f)$, $p_{\mathcal{R}, \alpha}(r)$ and $p_{\mathcal{FR}, \alpha}(f, r)$ are obtained by normalization of $h_{\alpha}(f, r)$ and the MI registration criterion is evaluated using

$$I(\alpha) = \sum_{f, r} p_{\mathcal{FR}, \alpha}(f, r) \log_2 \frac{p_{\mathcal{FR}, \alpha}(f, r)}{p_{\mathcal{F}, \alpha}(f) p_{\mathcal{R}, \alpha}(r)} \quad (8)$$

with

$$p_{\mathcal{FR}, \alpha}(f, r) = \frac{h_{\alpha}(f, r)}{\sum_{f, r} h_{\alpha}(f, r)} \quad (9)$$

$$p_{\mathcal{F}, \alpha}(f) = \sum_r p_{\mathcal{FR}, \alpha}(f, r) \quad (10)$$

$$p_{\mathcal{R}, \alpha}(r) = \sum_f p_{\mathcal{FR}, \alpha}(f, r). \quad (11)$$

The optimal registration parameters α^* are found by maximization of $I(\alpha)$. In [17], [59], Powell's multidimensional direction set method [88] is used to maximize $I(\alpha)$ (or minimize $-I(\alpha)$). The images are initially positioned such that their centers coincide and that the corresponding scan axes of both images are aligned and have the same orientation. Powell's direction matrix is initialized with unit vectors in each of the parameter directions. In [59], it is noted that an appropriate choice for the order in which the parameters are considered by Powell's optimization method needs to be specified, as this may influence optimization robustness. For brain images, for instance, the horizontal translation and the rotation around the vertical axis are more constrained by the shape of the head than the pitching rotation around the left to right horizontal axis. Therefore, it is suggested in [59] to first align the images in the transversal plane by optimizing the in-plane translation and rotation parameters first, in order to facilitate the optimization of the out-of-plane parameters. As the optimization proceeds, the Powell algorithm may introduce other optimization directions and change the order in which these are considered [88].

D. Alternative Implementations and Extensions

Various authors presented alternative optimization schemes for maximization of $I(\alpha)$. Studholme *et al.* [102],

[104] use an heuristic search procedure in which the parameters are one by one iteratively changed by small amounts such as to maximize MI. Meyer *et al.* [73] use a simplex search method to optimize the parameters of an affine or thin plate spline warped registration. Maes *et al.* [60] obtained analytic expressions for the gradient of MI w.r.t. the affine registration parameters, using partial volume distribution interpolation, and compared multiple gradient and not gradient based optimization strategies in terms of precision and speed. Thévenaz and Unser [110] presented a gradient based optimization strategy using analytic expressions for the gradient of MI computed using cubic spline based image interpolation and Parzen windowing of the joint intensity histogram. Jenkinson and Smith [47] proposed a scheme for global optimization of MI that combines a fast local optimization method with an initial search phase, demonstrating improved robustness of their scheme compared to local search strategies in the context of affine intersubject brain image registration.

Other schemes can be used to estimate the image intensity distributions in (8), for instance by using Parzen windowing [25] on a set of samples taken from the overlapping part of both images. This approach was used by Viola *et al.* [122]. In this approach, the image intensity probability densities are estimated as a mixture of Gaussian distributions from a small number of samples by weighting each sample's contribution to the histogram with a Gaussian window function G . A first set \mathcal{A} of N_A samples is used to estimate the densities $p_{\mathcal{FR}}^A(f, r)$, $p_{\mathcal{F}}^A(f)$, and $p_{\mathcal{R}}^A(r)$ by Parzen windowing as

$$p_{\mathcal{FR}}^A(f, r) = \frac{1}{N_A} \sum_{s_i \in \mathcal{A}} G(f - f(s_i), r - r(s_i))$$

and a second set \mathcal{B} of N_B samples is used to estimate mutual information as the log likelihood

$$I = \frac{1}{N_B} \cdot \sum_{s_i \in \mathcal{B}} \log_2 \frac{p_{\mathcal{FR}}^A(f(s_i), r(s_i))}{p_{\mathcal{F}}^A(f(s_i)) \cdot p_{\mathcal{R}}^A(r(s_i))}$$

with typically $N_A = N_B = 50$. By replacing each of the densities by a sum of Gaussians, analytic expressions are derived for stochastic approximations of mutual information and its gradient. Steepest gradient descent is then used to find the optimal registration transformation by maximization of I . While stochastic sampling of the images increases speed performance, the fact that only a limited number of samples is used, is likely to introduce local optima in the registration criterion and reduce registration robustness. Wells *et al.* [124] reported a 90% success rate of convergence to near the correct solution for a single pair of CT and MR images of the brain in 111 registration trials starting from a randomized initial position using translational and rotational offsets for each axis uniformly selected in the range $[-25, 25]$ mm and $[-10, +10^\circ]$ around the registered position. Parzen windowing was also used by Thévenaz and Unser [110] to estimate the joint probability distribution, but using B-splines as weighting functions instead.

Other information-theoretic registration measures can be derived from the mutual information criterion presented above, such as the entropy correlation coefficient $ECC(A, B)$ [59] or the normalized mutual information NMI [106]:

$$ECC(A, B) = 2 \cdot \frac{I(A, B)}{H(A) + H(B)} \quad (12)$$

$$NMI(A, B) = \frac{H(A) + H(B)}{H(A, B)} \quad (13)$$

with $0 \leq ECC \leq 1$ and $1 \leq NMI \leq 2$. ECC and NMI are one-to-one related by $ECC = 2 \cdot (1 - 1/NMI)$, such that maximization of each of these measures is equivalent. While Maes *et al.* [59] reported that they could not establish a clear preference for ECC versus MI for registration of full brain CT, MR, and PET images, Studholme *et al.* [106] found that maximization of NMI is superior to MMI itself in case the region of overlap of both images is relatively small at the correct registration solution, as MMI may be biased toward registration solutions with larger total amount of information $H(A) + H(B)$ within the region of overlap.

Several authors have proposed adaptations of the MI measure to include spatial information that is contained in each of the images separately, in order to increase registration robustness and accuracy in cases where the joint intensity histogram by itself is insufficient.

Butz and Thiran [8] applied MMI not to the original intensities, but to an edginess measure such as the gradient magnitude, indicative for the location of object surfaces. However, the fact that a lot of intensity information is discarded and due to the sparseness of edge features, the edginess-based MI shows more and more pronounced local maxima, necessitating the use of a genetic optimization strategy to find the global optimum.

In the approach of Pluim *et al.* [84], spatial information is incorporated in the registration criterion by multiplying MI with a gradient term that measures the coincidence of voxels in both images with strong gradients of similar orientation, which is considered to be indicative of alignment of tissue transitions in both images. The combined measure was evaluated for rigid-body registration of CT, MR, and PET brain images and was found to yield a smoother function than mutual information or its normalized version, to be less affected by interpolation artifacts and to perform better for low-resolution images.

Studholme *et al.* [105] experimented with incorporating additional information channels in the mutual information criterion. In some difficult applications, such as registration of pelvic MR and PET images, Studholme observes that registration on intensities alone may not provide a distinct optimum due to a lack of differentiation between spatially unconnected regions in one modality which seem to be connected in the other. Incorporation of region labeling information is then useful to express region connectivity or higher level anatomical information. In general, the mutual information registration criterion can be straightforwardly extended to the case in which two already registered images A and B

have to be registered to a third one C by expressing the mutual information between the pair (A, B) and C as

$$\begin{aligned} I(A, B; C) &= H(A, B) - H(A, B|C) \\ &= (H(A) + H(B|A)) - (H(A|C) + H(B|A, C)) \\ &= (H(A) - H(A|C)) + (H(B|A) - H(B|A, C)) \\ &= I(A; C) + I(B; C|A) \end{aligned}$$

or as the sum of the information that A contains about C and the additional information that B contains about C that is not already contained in A .

Rueckert *et al.* [94] considered the cooccurrence of neighboring voxel intensities in each of the images to be registered and defined second-order mutual information to assess the correspondence of intensities of adjacent voxel pairs in both images, explicitly taking the dependence of neighboring voxel intensities into account. The method was evaluated for nonrigid registration of brain MR images with simulated and real intensity inhomogeneity, showing that higher-order mutual information is much less sensitive to shading artifacts than traditional mutual information.

Mutual information $I(A, B)$ is only one example of the more general f -information measures of dependence $f(P||P_1 \times P_2)$ [114] with P the set of joint probability distributions $P(A, B)$ and $P_1 \times P_2$ the set of joint probability distributions $P(A) \cdot P(B)$ assuming A and B to be independent. f -information is derived from the concept of f -divergence, which is defined as

$$f(P||Q) = \sum_i q_i \cdot f\left(\frac{p_i}{q_i}\right)$$

with $P = \{p_1, p_2, \dots\}$ and $Q = \{q_1, q_2, \dots\}$, with an appropriate choice for f and with suitable definitions when $q_i = 0$. One example of f -divergence is

$$I_\alpha(P||Q) = \frac{1}{\alpha(\alpha-1)} \left[\sum_i \frac{p_i^\alpha}{q_i^{\alpha-1}} - 1 \right]$$

with corresponding f -information

$$I_\alpha(P||P_1 \times P_2) = \frac{1}{\alpha(\alpha-1)} \left[\sum_{i,j} \frac{p_{ij}^\alpha}{(p_i \cdot p_j)^{\alpha-1}} - 1 \right]$$

with $p_{ij} = P(i, j)$ and $p_i = \sum_j p_{ij}$ and $p_j = \sum_i p_{ij}$. I_α equals MI for $\alpha \rightarrow 1$. Several f -information measures, including I_α for various α , were investigated by Pluim *et al.* [83], [86] and compared to MI in terms of accuracy of their optimum, robustness with respect to initial positioning and overall registration performance for rigid-body registration of CT, MR, and PET brain images with marker-based gold standard registration (taken from the RREP study [50], [126], see Section III-B). In some cases, these measures were found to outperform MI in terms of accuracy when initiating the optimization with the reference solution itself, but not in overall registration performance when starting the optimization from an initial position further off from the reference.

Inspired by the MMI registration criterion, Buzug *et al.* [9] and Bro-Nielsen [5] experimented with alternative histogram based registration measures motivated by convexity arguments and by the literature on grey-level cooccurrence matrices respectively, but their approaches lack the mathematical elegance of the mutual information registration criterion.

III. VALIDATION

A. Robustness

The robustness of the MMI registration algorithm was investigated by various authors with respect to implementation issues such as (sub)sampling [59], [87], [104], optimization [60], [110], and especially interpolation [59], [83], [85], [110]. Thurfjell *et al.* [113] and Zhu and Cochoff [130] adopted the approach of [59] and studied the influence of several implementation options, such as interpolation method, number of histogram bins, optimization approach and multiresolution strategy in the context of MR and SPECT brain image registration, tuning the algorithm toward optimal performance in terms of registration accuracy, robustness and speed.

Different interpolation schemes used for histogram binning show different behavior of the MI registration measure in the neighborhood around the optimum, such that the sensitivity of the registration solution to the initial positioning of the images and to the optimization strategy used depends on the choice of the interpolation method [59]. While NN and TRI interpolation usually show many weak local optima near the correct registration solution, PV interpolation yields a smooth MI registration criterion with a large basin of attraction around the correct optimum [59]. However, PV interpolation may introduce artifacts in the registration criterion if the images to be registered have identical voxel grids or if their voxel sizes are multiples of each other in one or more dimensions [58]. Indeed, if the images are aligned such that all or many of the voxels in both images coincide exactly, many of the PV weights are zero and the histogram dispersion introduced by the PV distribution scheme is less at grid-aligning registration positions. Since PV histogram dispersion increases again when moving away from such a position in any direction, PV interpolation may yield a local maximum of MI at these positions (especially with low resolution images with low spatial intensity correlation such that neighboring voxels in the reference image are more likely to contribute to different histogram bins). If the true registration solution is sufficiently far away from a grid-aligning position (for instance, if it involves significant rotation of one image relative to the other) and has a broad attraction range, these local maxima are unlikely to have an impact on the optimization process and it usually suffices to start the optimization from an initial position in which the grids are not aligned (for instance, by specifying a nonzero initial rotation around one or more axes). However, if the true registration differs very little from a grid-aligning position, the optimization is likely to be attracted by the local maxima introduced by the PV interpolation scheme itself, such that these PV interpolation

artifacts are very likely to deteriorate registration robustness. Especially for recovering subvoxel small displacements of images with identical grids (e.g., motion correction of fMRI time series images) TRI interpolation is to be preferred [58].

The impact of PV interpolation induced artifacts on mutual information based registration of clinical images was investigated extensively by Pluim *et al.* [85], who also showed the occurrence of artifacts for TRI interpolation. While PV interpolation yields local maxima of MI at grid-aligning positions, TRI interpolation yields local minima at these positions. This is explained in [85] by the fact that TRI interpolation blurs the reference image, making noise and other small scale structures disappear and therefore reducing histogram dispersion. At grid-aligning positions, however, no blurring is applied and the joint entropy can, therefore, be higher (and, hence, MI smaller) for grid-aligning transformations than for nongrid-aligning transformations. An example illustrating PV and TRI interpolation artifacts is shown in Fig. 5. Pluim *et al.* [85] also showed that a small resampling of one of the images, such that the voxel sizes of both images are no longer equal and grid alignment therefore cannot occur, largely removes the interpolation artifacts and vastly increases the smoothness of the registration function. The impact thereof on registration accuracy was evaluated for registration of CT and MR brain images with identical interslice distance (taken from the RREP study [50], [126], see Section III-B), but no significant difference in accuracy for registration of CT to the original and to the resampled MR images was found.

It has been shown that for high-resolution images subsampling of the floating image can be applied without deteriorating optimization robustness of the MMI registration criterion [60], [87]. Important speedups can, thus, be realized by using a multiresolution optimization strategy, starting with a coarsely sampled image for efficiency and increasing the resolution as the optimization proceeds for accuracy [104], [110]. Multiple multiresolution strategies, involving different subsampling factors and number of resolution levels and using various optimization methods, were compared in [60], demonstrating possible speedups up to a factor of three over the original method of [59] with similar precision and without loss of robustness. Pluim *et al.* [87] found that the use of a multiresolution optimization strategy for MMI registration of low-resolution images, with and without resolution-dependent image blurring at each resolution level, merely increased computation time without improving registration robustness.

Studholme [103], [104] investigated the influence of the initial positioning of the images on the registration result, comparing his implementation of the MMI criterion with other voxel based similarity measures, such as Collignon's entropy measure [18], Hill's third-order moment [42], cross-correlation and Wood's variance of intensity ratios measure [129]. Multistart experiments reported in [104] for a single CT and MR dataset of the brain showed that the standard deviation in the registration solutions was smaller than 0.1 mm for translation and 0.15° for rotation and that clustering

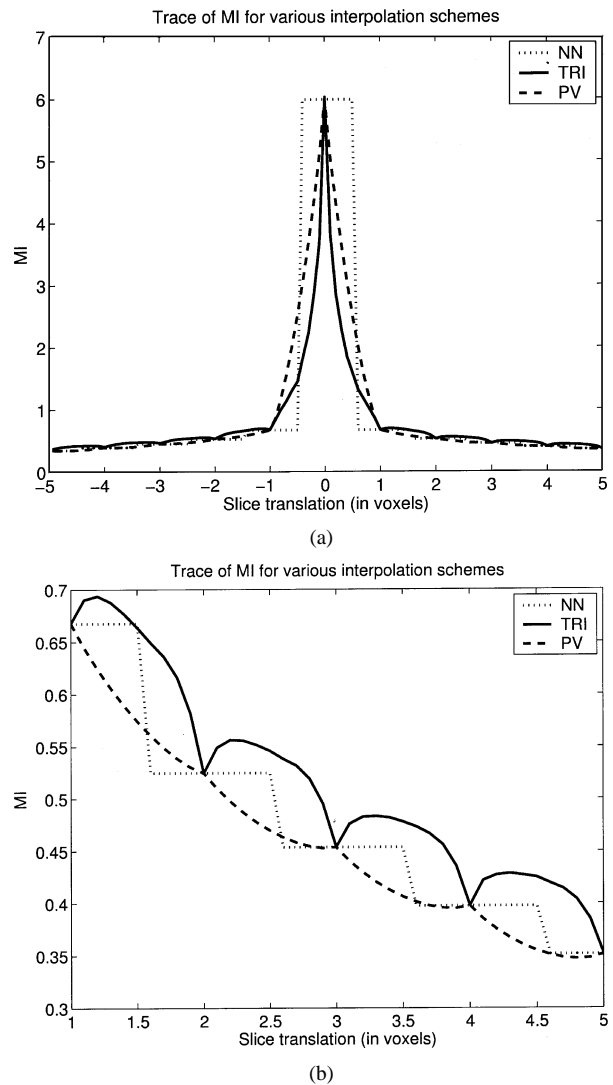


Fig. 5. Traces of MI between an MR brain image and itself translated in the slice direction. (a) Range from -5 to $+5$ voxels around the registered position. (b) Detail of (a). Note the local minima and maxima for TRI and PV interpolation respectively at grid-aligning positions at multiples of one voxel displacements. NN interpolation is insensitive to translations up to one voxel, which explains the staircase pattern.

of the results was significantly better for the MMI criterion than for the other criteria.

Studholme [104] also verified the impact of partial overlap on registration robustness and found that clustering of registration results obtained in his multistart experiments was best if the skull base was included in the registration and worse if only the upper half of the CT dataset was used. In [59], it was reported that registration of a single dataset of high resolution CT and MR images of the brain using either the full images spanning the entire head or using only part of the datasets has only subvoxel influence on the registration result.

The robustness of the MMI criterion with respect to image degradations, such as noise, intensity inhomogeneity or geometric distortion, was illustrated in [59] by inspecting the behavior of the MI criterion for registration of an original MR image and its artificially degraded version. Different degradations of the original image intensities were created

Table 2

Quantitative Validation of Registration Accuracy of Various Registration Algorithms by Comparison With External Marker Based Registration for Registration of CT, MR, and PET Images of the Brain [126]. Registration Features are: Point Landmark Based (PB), Surface Based (SB), and Voxel Based (VB). The Worst Case Median Error (in millimeters) for CT to MR and PET to MR Registration Over All Six MR Scans Available per Patient as Reported in and Three in [126] and on the RREP Website [50] (n.a = no results available). The Various Approaches Using Mutual Information Differ in Implementation Issues and in the use of Mutual Information Versus Normalized Mutual Information as Defined in (13).

Group	Method	CT/MR	PET/MR
Barillot [54]	SB chamfer matching	2.48	5.22
Collignon [17],[59]	VB mutual information	1.50	3.64
van den Elsen [115]	VB intensity correlation	2.01	n.a.
Harkness [77]	SB head-hat distance	4.24	4.16
Hemler [37]	SB head-hat distance	4.70	n.a.
Hill [104]	VB mutual information	1.94	3.25
Maintz [64],[66],[65]	VB edgeness correlation	5.05	5.33
Malandain [67],[68],[69]	SB potential field	5.42	4.85
Noz [61]	PB elastic warp	7.80	5.17
Pelizzari [77]	SB head-hat distance	2.79	3.33
Robb [48],[49]	SB chamfer matching	5.90	3.85
Woods [129]	VB intensity variance	n.a.	3.36
Ashburner [2]	VB tissue alignment	n.a.	4.20
Chen	VB mutual information [59]	2.07	3.87
Ding [24]	PB template matching	2.47	4.16
Doll	VB mutual information [59]	2.67	6.37
Hsu [46]	SB edge features	2.51	4.82
Li	VB mutual information [59]	1.71	4.74
Lloret [57]	VB creaseness correlation	3.41	n.a.
Ng	VB mutual information [124]	1.71	4.33
Ren	VB normalized mutual information [104]	1.96	4.42
Rodriguez [91]	VB mutual information	2.16	3.95
Rohlfing	VB normalized mutual information [104]	1.80	3.27
Rouet [93]	SB head-hat distance	4.52	n.a.
Tanacs	VB mutual information [59]	1.92	3.87
Thevenaz [110]	VB mutual information	2.04	5.03
Wyman	VB mutual information [124]	1.68	12.01
Zhu [130]	VB mutual information	2.26	3.38

by superimposing zero-mean Gaussian noise, by multiplication with a 2-D radially symmetric quadratic bias field and by simulating geometric distortion along one coordinate axis using the magnetic field inhomogeneity model of [74]. In these simulations, the MMI registration solution was not affected by noise or (mild) intensity inhomogeneity: MI decreased when the degradation was more prominent, but in all experiments the location of its maximum was identical to the case when no degradation was applied. In the simulations involving geometric distortion, the registration solution shifted in the same direction and over a distance proportional to the average distortion applied.

B. Accuracy

The accuracy of the MMI registration algorithm was validated within the framework of the RREP conducted by Fitz-

patrick *et al.* at Vanderbilt University [50] and reported in West *et al.* [125]–[127]. The primary objective of this study was “to perform a blinded evaluation of a group of retrospective image registration techniques using as a gold standard a prospective, marker-based registration method.” Image volumes of three modalities (CT, MR, and PET) were obtained from nine patients undergoing neurosurgery. Up to six different MR images were available for each patient: T1, T2, and proton density weighted images and geometric distortion corrected versions thereof. A stereotactic frame and additional bone-implanted fiducial markers, designed to be visible in CT, MR, and PET, were rigidly mounted to the skull of each patient. The centroid of each marker was detected in each of the images using the fiducial localization technique described by Wang *et al.* [123]. The CT and PET images were registered to the MR images by calculating the rigid

body transformation that minimizes the mean square distance between corresponding marker positions in the two images. The rigid body assumption is well satisfied inside the skull in 3-D scans of the head if scanner miscalibration, geometric distortion of the images and patient related changes, for instance due to interscanning interventions, can be neglected. These marker-based transformations were used as the gold standard for evaluation of the retrospective methods. The accuracy of this reference was estimated by numerical estimations to be 0.4 mm for CT to MR and 1.7 mm for PET to MR registration [72], [126]. To ensure blindness of the study, the frame and the fiducial markers were removed from the images by manual editing, approximately reconstructing the image background in manually outlined regions containing these structures. The gold standard results were not communicated to the participants in the study until after their results had been submitted. The transformation differences between the reference and the submitted transformations were evaluated at 10 anatomically relevant volumes of interest (VOI) in the brain. Since results for different VOIs were all very similar, results were pooled over all VOIs. The median and maximal error over all sites and all patients were reported in [126] for each method that participated in the study. These results are summarized in Table 2, together with results by some other groups on the same data sets that were submitted after the initial publication [126] and that are publicly available on the RREP website [50]. Although the number of registration experiments performed in this study was too small to draw statistically significant conclusions regarding relative performance of different methods, it is clear that some methods performed better than others. Two groups that participated in the initial RREP study [126] (Hill and Collignon in Table 2) submitted registration results computed with (different implementations of) the MMI registration criterion. The RREP evaluation shows that the results submitted by both groups differ almost always by less than a voxel from the external marker based registration solution, both for CT/MR and for PET/MR registration, which provides strong evidence for the subvoxel accuracy of the MMI approach.

While the CT and MR images in the original RREP study as reported in [126] had an interslice distance of 4 mm, a second series of nine patient data sets with higher resolution CT and MR images and their gold standard reference registration solution has been made available by Vanderbilt University [50]. Each set includes one CT image (3-mm slice distance) and up to four MR images (T1, T2, and PD weighted, 3-mm slice distance; and T1 weighted, 1.25-1.66-mm slice distance). We measured the difference between the reference registration solution and the MMI registration result obtained with the algorithm of Maes *et al.* [59] at eight points located near and well distributed over the brain surface to obtain an upper bound for the MMI registration error within the brain. Mean errors over all cases ranged from 0.5 to 1.7 mm for the high resolution MR images and from 1 to 3.1 mm for the lower resolution MR images [58]. All errors are thus smaller than or at most as large as one CT voxel.

Fig. 6 shows a visual comparison between the RREP marker-based reference registration transformation and the




Fig. 6. Visual comparison between the marker-based reference registration transformation (top) and the MMI registration solution (middle) for CT to MR registration, showing the same orthogonal MR slices and similar iso-intensity contours extracted from the corresponding reformatted CT slices after registration with either of both transformations. Overlay of both sets of contours (bottom) and comparison of the contours on the axial and sagittal slices seem to indicate that the accuracy of the MMI solution is superior to that of the reference transformation itself. However, this ignores the possibility of geometric distortion of the images that may affect the registration result (see text for discussion). This image was generated with the multipurpose imaging tool (MPITool) developed by Uwe Pietrzyk and colleagues at the Max-Planck-Institut für neurologische Forschung in Köln [79]–[81].

MMI registration result for one set of CT and PD weighted MR images (RREP patient 101). The figure shows three orthogonal slices from the MR image with the outline of the skull overlaid, extracted by intensity thresholding from the geometrically corresponding slices from the registered and reformatted CT image. In the first set of MR images in Fig. 6, the CT image was registered to the MR image using the reference registration transformation, while the contours on the second set of MR images in Fig. 6 were obtained with the MMI registration solution. In both cases the same CT intensity threshold was used to define the contours. The third set of images in Fig. 6 shows both contours overlaid to mark the difference between both. The difference between the RREP reference and the MMI rigid-body registration transformation in this case was found to be small in the right/left and inferior/superior direction (0.08 and 0.34 mm, respectively), but much larger in the anterior/posterior direction (2.92 mm or about 3.5 pixels). When visually comparing the contours on the axial and on the sagittal slices obtained with both transformations, it seems that the reference registration transformation places the skull in the CT image too much anterior relative to the skin visible in the MR image, aligning the outer surface of the skull in CT with the outer surface of the skin in MR. The MMI registration solution on the other hand seems to correctly align both the inner and the outer CT skull surface to the MR brain and

skin surface, respectively, and, at first sight, seems more accurate than the marker-based reference transformation in this case. A similar misregistration of 2.3 mm on average in the same direction is observed for CT/MR registration using MMI involving the other PD images in the study (RREP patients 101-104).

However, a possible explanation for the observed discrepancy between both registration solutions may be the presence of geometric distortion in the MR image, caused for instance by magnetic field inhomogeneity, gradient nonlinearities, magnetic susceptibility, or chemical shift artifacts [74], [109]. Such distortion is most apparent in the direction of the readout gradient, which coincides with the anterior/posterior axis for the RREP images [71]. Different sources of geometric distortion may have a different impact on both the extrinsic, marker-based, and intrinsic, image-feature-based registration [22], [38], [71], [108], each being more sensitive to some particular distortion sources and rather insensitive to others. The accuracy of marker-based registration may be affected for instance by displacement of the marker position as observed in the images due to magnetic susceptibility artifact at the air/skin interface [22]. Voxel-based registration, on the other hand, may be biased for instance by displacement of the skin relative to the skull in the direction of the readout-gradient due to chemical shift between fat and water at the skin/skull interface. (While chemical shift artifacts in MR imaging can be effectively suppressed by a variety of measures [74], [108], they can not be neglected in the RREP images due to the small readout gradient used. From the imaging parameters mentioned in [126] and using [74, (18)], the expected displacement between fat and water in these images is 3.5 mm, which is of the same order as the MMI registration error). Visual inspection of registration differences as in Fig. 6 is, therefore, not a reliable approach for evaluating absolute registration accuracy.

The effect of geometric distortion correction on registration accuracy [38], [71] was investigated in the initial RREP study [126] by also providing MR images that were corrected for geometric distortion using the image rectification technique of Chang and Fitzpatrick [12]. A significant decrease in registration error was reported only for Collignon's MMI registration of CT and T2-weighted MR images. Maes [58] observed a clear tendency toward lower registration errors and a significant increase of MI at registration for registration of CT and PET to the corrected instead of the original MR images, indicating that the applied geometric distortion correction indeed globally improves the geometric similarity of both images and, hence, the registration accuracy. In applications where accurate registration is crucial, such as for instance the use of MR images for target identification and trajectory planning in stereotactic neurosurgery [33], optimization of the MR imaging protocols is needed to minimize confounding effects of geometric distortion artifacts on the registration result [74].

IV. APPLICATIONS

Because of its reliability and generality and because of its full automation, image registration by MMI has large poten-




Fig. 7. Registered and reformatted CT and MR images of the brain of a patient showing a lesion. After registration and reformatting, the images can be jointly visualized and use of a linked cursor allows to delineate contours in one image and transfer these to the other for comparison. (a) CT. (b) MR. (c) Blending of the skull from CT with the brain tissues from MR demonstrates the high quality of the registration result.

tial for routine use in clinical practice in a variety of applications, involving various organs and imaging modalities. We illustrate some of these applications here.

A. Within-Subject Registration of CT, MR, and PET Images of the Brain

The MMI criterion has been demonstrated to consistently and robustly match CT to MR and PET to MR images with subvoxel accuracy [126]. We have applied the method successfully in a large number of experiments (more than 500) involving low- and high-resolution MR images (5- to 1-mm slice thickness) with multiple orientations (axial, coronal, and sagittal) and acquired with various imaging protocols (T1 and T2 weighted, proton density, gadolinium enhanced, diffusion weighted), high resolution MR angiography (MRA) images, conventional and spiral CT images, PET and SPECT images. Clinical applications of MMI based image registration we have experience with include: planning of stereotactic surgery from CT and MR images; diagnosis of brain tumors from MR and PET images; assessing MS lesion evolution over time from MR time series images; brain morphometry from multiprotocol MR images; and motion correction in the analysis of functional MRI (fMRI) time series and MR perfusion imaging. Some selected results are shown in Figs. 7-9.

B. Template-to-Patient Registration of MR Brain Images

Correlation of functional MRI data obtained from different individuals can be achieved by registration of the corresponding anatomical MR images with a fixed template image, for instance using the Statistical Parameter Mapping (SPM) software package developed by Friston *et al.* [28],




Fig. 8. Registered and reformatted T2-weighted MR images of the brain of an MS patient at four different time points over a period of one year. Registration of such MR time series images allows to assess the evolution of typical MS white matter lesions over time.




Fig. 9. Registered and reformatted MR images of the brain of an MS patient at a single time point. (a) T1 weighted. (b) T2 weighted. (c) Proton density. (d) Fluid attenuated inversion recovery (FLAIR) image. Availability of such coregistered multispectral MR brain images facilitates the segmentation of white and gray matter tissue and of the white matter lesions by (semi-)automated intensity-based classification algorithms.

[29]. The template provided by SPM has been normalized to Talairach space, such that after registration of each of the study images to the template, functional measurements from different individuals can be compared using Talairach coordinates. SPM also provides a statistical atlas of the distribution of white matter, gray matter, and CSF within the brain that is co-registered with the template. Template-to-pa-




Fig. 10. Registration of PET and MR images of the brain of a patient undergoing stereotactic surgery. The MR image was acquired with the patient wearing the stereotactic reference frame with external markers attached. Intensity-based registration of the previously acquired PET image with the MR image using MMI allows to transfer the PET image into the stereotactic space and to plan the surgical intervention based on MR and PET information combined.




Fig. 11. Alignment of a template brain image to a patient MR brain image by affine image registration: (a) axial; (b) coronal; and (c) sagittal slice of the template overlaid on the registered patient image.

tient image registration allows to map the information of the atlas onto the image. Each voxel in the patient image can, thus, be assigned an *a priori* probability for tissue type that can be used to guide intensity-based tissue classification [118].

Since the MMI registration criterion is multimodal and highly robust with respect to contrast differences, the method is well suited to co-register MR images acquired using different protocols to a single template, without need for tuning the algorithm each time or without having to specify different templates. A typical result is shown in Fig. 11. The template has 2-mm cubic voxels and size $91 \times 109 \times 91$. The patient MR image has 1-mm cubic voxels and dimensions $256 \times 256 \times 160$. The registration is computed from the template to the patient image using a full 12-parameter affine transformation. Fig. 11 shows overlays of an axial, coronal, and sagittal slice of the template on the registered MR image. Because of the variability of the brain topology, the rigid body assumption is not satisfied in this case, but the MMI criterion succeeds at finding such an affine transformation that on average very well matches corresponding structures in both images. Visual inspection of the registration results on a large number of cases demonstrated the excellent performance of the MMI registration criterion for affine template-to-subject registration in the context of atlas-based brain tissue segmentation [119].

C. Registration of CT and MR Images of the Prostate

Target delineation and dose distribution computation in the planning of radiotherapy treatment of prostate cancer is usually performed on CT images. However, there is some indication that the prostate, and especially its apex, can be more reliably and accurately delineated on MR images, especially when both transversally and coronally oriented MR cross-sections are considered (see [23] and the references therein). Registration of both scans allows to map the boundary contours of the objects of interest delineated in the MR images onto the corresponding locations in the CT image, such that the relevant information of both scans (object discrimination in MR, tissue density in CT) can be effectively combined during planning.

In an initial study involving ten patients [23], spiral CT images (512×512 matrix, 50 slices, 0.7-mm pixel size, 5-mm slice thickness) were acquired in radiation position, while MR images (512×512 , 22 slices, 0.7-mm pixel size, 5-mm slice thickness) were acquired using a standard abdominal imaging protocol. Care was taken to have all the image acquisition done with empty bladder and rectum. The time between both acquisitions was kept at a minimum, such that intestinal and bladder filling differences between both scans were small. The images were aligned with a rigid-body registration transformation, which was computed completely automatically using MMI without user intervention. After registration, the images were reformatted and visualized together and the registration result was visually inspected by assessing the alignment of corresponding rigid structures in both images, such as the pelvic bone. In all cases, the registration result was judged to be excellent by radiologists. The study [23] concluded that the additional use of axial and coronal MR scans, in designing the treatment plan for localized prostate carcinoma, improves substantially the localization accuracy of the prostatic apex and of the anterior aspect of the rectum, resulting in a better coverage of the prostate and a potential to reduce the volume of the rectum irradiated to a high dose.

A typical registration result is illustrated in Fig. 12(a). Despite the presence of large local nonrigid deformations of soft tissues, such as skin, fat and muscles, induced by the different table shape of both scanners, the MMI criterion in this case succeeded at aligning corresponding rigid structures in the images without being confused by structures that can not be aligned by a rigid-body transformation. Nevertheless, in order to minimize the possible impact of such large tissue distortions on registration robustness when applying the method in clinical routine, it is recommended to avoid large differences in patient positioning, for instance by placing a board on the CT table to mimic the flat shape of the MR table [Fig. 12(b)].

D. Registration of PET and CT Images of the Thorax

Joint interpretation of PET and CT images of the thorax aids the detection and staging of metastatic mediastinal lymph nodes (MLN) in patients with lung cancer [120]. PET is more accurate for detecting MLN than CT thorax,




Fig. 12. Registered and reformatted CT and MR images of the prostate used for radiotherapy planning. The CT image (left) is needed for dose distribution simulation, while the target volume and organs at risk can be more accurately delineated in the corresponding MR image (right). The central part of either image has been shown in overlay on top of the other for visual inspection. (a) This example illustrates the robustness of the mutual information criterion to local nonrigid deformations, in this case induced by the different table shape of both scanners. The rigid bony structures are matched correctly, despite the presence of large deformations of the soft tissue structures. (b) In clinical practice, a board is placed on the CT table to mimic the MR table shape, avoiding large local tissue distortions between both scans.

but localization of lesions, essential for adequate staging, is difficult due to the little anatomic references on PET images. Registration of the PET and CT images allows to effectively combine the information contained in both modalities, i.e., to detect the metastatic nodes in PET and locate them anatomically in CT.

Spiral CT images of the thorax are acquired using a standard imaging protocol with a matrix of 512×512 pixels, typically 50 slices and 5-mm slice thickness, while the PET images have a 128×128 matrix, typically 30 slices and 6-mm slice thickness. The patients are positioned similarly for both acquisitions with the arms beside the body. PET transmission images are acquired prior to the FDG emission images in the same session without the patient leaving the scanner table. The transmission images contain more anatomical information than the emission images and are therefore more suited for registration with CT. The PET transmission and CT images are registered completely automatically with the MMI registration algorithm without any preprocessing of the images. After registration, the PET transmission images are reformatted to the CT image grid to make use of the full resolution of the CT. The quality of the registration is visually inspected by radiologists using the carina, the trachea and the lung contours as anatomical landmarks. The same reformatting is applied to the PET emission images, assuming that the emission and transmission images are aligned by acquisition. The registered CT and PET emission images are then used for localization of MLN. A typical result is shown in Fig. 13.

To investigate the benefit of PET/CT registration for lung cancer staging, a study has been conducted at our institu-




Fig. 13. Registration of PET and CT images of the thorax for mediastinum cancer staging. (a) CT image. (b) PET transmission image. (c) PET emission image. Top: axial cross-section. Middle: coronal cross-section. Bottom: sagittal cross-section. The PET transmission image was matched to the CT image, assuming alignment at acquisition of the PET transmission image with the corresponding emission image. The PET images have been reformatted to the CT grid to make use of the higher resolution of the CT image. Hot spots detected in the PET emission image can easily be located and anatomically identified in the CT image using a linked cursor.

tion involving 50 patients [120]. Registration of the mediastinum was considered very good and clinically relevant in all cases. The efficiency of the automated and highly reliable MMI registration algorithm has now changed the whole diagnostic procedure. Currently, at our institution, PET-CT registration is used in clinical routine to decide on lung cancer treatment for two patients a week. Weekly clinical sessions are organized in which the available image information, after registration and after post-processing to indicate the relevant structures, is studied jointly by pneumologists, thorax surgeons, nuclear medicine experts, and radiologists.

E. Registration of Pre- and Post-Operative CT Images

Complicated bone surgery procedures can be planned and simulated using pre-operatively acquired CT images. The outcome of the intervention can be evaluated by comparing post-operatively acquired CT images with the pre-operative planning, after proper registration of pre- and post-operative images. One such application is illustrated in Fig. 14. In this study, the value of spiral CT imaging was investigated in planning and evaluation of chin bone grafts in alveolar cleft patients. Pre- and post-operative CT images were acquired two weeks before and nine months after the operation. Registration of pre- and post-operative images allowed to define the cleft region in the pre-operative image and localize the same area in the post-operative image, such that the condition of the transplant can be followed over time.

The procedure was applied to eight patients, all children between eight and ten years old. Spiral CT images (512×512 , 40-70 slices, 0.2-mm pixel size, 1-mm slice thickness) were acquired using standard imaging protocols




Fig. 14. Matching of pre- and post-operative CT images of the lower jaw of a child with an innate alveolar cleft in the jawbone after bone graft transplantation. (a) Pre-operative image. (b) Registered and reformatted post-operative image. Contours delineated on the co-registered images using a linked cursor allow to evaluate registration accuracy. Although morphological changes due to bone growth are clearly present, these do not impair the algorithm from finding the globally best rigid transform between both images.

pre- and post-operatively. The images were registered completely automatically and without pre-processing using the MMI registration criterion. After registration, the pre-operative image was reformatted to match the post-operative image and the result was visually inspected by a radiologist. A typical result is shown in Fig. 14. Contours delineated in one image and transferred to the other allow to inspect the registration accuracy. Note that in the nine month period between both acquisitions morphological changes have occurred due to bone growth, teeth movement, and surgery. However, these nonrigid movements do not impair the registration algorithm to find the rigid transform that globally very well aligns both images.

V. NONRIGID IMAGE MATCHING USING MMI

Nonrigid registration or matching of 3-D images involves finding a 3-D field of 3-D deformation vectors that maps each point in one image onto the corresponding point in the other image. Nonrigid image to image registration warps one image toward a second one such that all objects in the warped image precisely coincide with the corresponding objects in the target image. In contrast with affine registration, which only allows for translation, rotation, scaling or skew of one image relative to another in order to establish global alignment of both images all over the image domain as illustrated in Section IV, nonrigid registration allows displacement of individual voxels such that local, regional distortions between both images can be corrected for up to voxel scale.

Various approaches for multimodal nonrigid image registration using MMI have been proposed that differ in their regularization of the deformation field and in the way the variation of MI with changes in the deformation parameters is computed.

A. Regularization of the Deformation Field

While the affine registration transformation is defined by not more than 12 parameters or degrees of freedom, non-rigid registration typically requires much more degrees of freedom in order to make the deformation field sufficiently flexible to correctly match both images everywhere in the

image domain. In its most general form, nonrigid image registration allows each voxel to be displaced independently from its neighbors with as much as 3 degrees of freedom per voxel if no additional constraints are specified. For nonrigid matching of 2 high resolution $256 \times 256 \times 128$ images of the brain for instance with more than 8 million voxels each, more than 25 million parameters need to be determined to define the deformation field when allowing unconstrained, free-form deformation of all voxels individually. However, 3-D displacement of all voxels independently is undesirable, as it is likely to result in a deformation field that is not consistent with deformations of real materials, for instance implying folding of the coordinate system causing adjacent structures to “cross over” one another [15]. Moreover, the nonrigid image registration problem is ill-posed because the images most often do not contain the necessary information to uniquely define the nonrigid registration solution in every voxel in the image domain, for instance inside regions with homogeneous intensity or along smooth object boundaries. Regularization of the registration problem is, therefore, required to constrain the solution space to include only deformation fields that are physically acceptable and to smoothly propagate or extrapolate the registration solution from sites with salient registration features (e.g., object boundaries) toward regions where registration clues are absent or ambiguous (e.g., object interior). The deformation field to be recovered is typically assumed to be a homeomorphism, i.e., a continuous, one-to-one and topology preserving mapping [15]. The assumption that there exists a one-to-one mapping between the images to be registered may be violated in case of important variability in topology between corresponding objects, for instance in the sulcal and gyral patterns of the cortex [19].

One approach for regularization of the nonrigid registration solution is the representation of the deformation field as a weighted sum of smooth basis functions. The basis functions can have global support, i.e., spanning the entire image domain such as thin plate spline [73], polynomial [45], Gaussian [53] or trigonometric functions [1], or local support, i.e., being nonzero in a subregion only, such as B-splines [96], [101] or localized radial basis functions [92]. When using basis functions such as splines, the number and distribution of spline control points determines the elasticity of the deformation and the number of degrees of freedom. This can be much smaller than the number of voxels, especially when using a multiresolution strategy with a small number of sparsely distributed control points at coarse resolution levels and a larger number of more densely distributed control points at finer resolution levels. Some strategies have also been presented to adaptively refine the control point grid depending on the local misregistration in the image [92], [97]. Smoothness of the deformation field is intrinsic to the parameterization at scales smaller than the control point spacing and may be imposed at larger scales by penalizing high spline curvature [96]. Spline-based approaches can correct for gross shape differences, but a dense grid of control points is required to characterize the deformation at voxel level detail, implying high computa-

tional complexity, especially when using basis functions with global support.

Free-form nonrigid registration approaches, using a nonparameterized expression for the deformation field, are in general more flexible than representations using basis functions, but need appropriate smoothness constraints for spatial regularization of the resulting vector field. Such constraints can be imposed by including additional cost terms in the similarity metric that penalize nonsmooth deformations [36], [52] or by modeling the deforming image as a linear elastic [3], [31] or viscous fluid [14] material. Elasticity constraints are suitable when displacements can be assumed to be small, while a viscous fluid model is more appropriate when large deformations need to be recovered. Tissues deform under the influence of local image-derived forces, defined such that the resulting displacements optimize the registration criterion. Smoothness of the deformation field is imposed by the elasticity or viscosity terms of the Lagrangian partial differential equation (PDE) that governs tissue motion. Numerical schemes for solving this equation iteratively at voxel resolution over the entire 3-D image domain are computationally demanding, especially with fluid models allowing large deformations [14]. Much faster, simplified schemes have been presented that approximate the PDE regularization by Gaussian filtering of the force field [111], [112], [6], [128]. Free-form nonrigid registration methods using MMI have been proposed using elastic [39] and viscous fluid [21] models.

Rather than using a purely mathematical approach, an appropriate deformation model can also be constructed by statistical analysis or be inspired by physics-based [101] or biomechanical arguments. Regional morphological variability in the image and correlations between deformations in different image regions can be learned from a set of similar images that are all registered to the same template by statistical analysis of the parameters of their deformation fields [95]. The deformation field of new, similar cases to be registered can then in principle be constrained by the modes of variation observed in the training set, provided that the training set is sufficiently representative. Biomechanical models on the other hand take material properties of the various tissues into account to impose tissue specific deformation constraints. These methods are often implemented using a finite element model defined on the voxel grid [34] or on a volumetric meshing of the objects of interest [26], after appropriate segmentation of the various tissue classes. We are not aware of any works using MMI for nonrigid image registration in combination with a biomechanical deformation model.

B. Registration Criterion

Affine registration parameters have a global effect on registration quality all over the image domain and can therefore be estimated from registration evidence gathered from the entire region of overlap of the images to be registered. In contrast, a change in one of the nonrigid registration parameters typically changes the deformation field within a subre-

gion only, the size of which depends on the parameterization of the deformation field, such that the impact on registration quality of each such local deformation parameter can only be observed within its region of support. This may be problematic for a statistical measure such as MMI: especially when small scale, highly localized deformations need to be recovered, the number of samples within each parameter's region of support may be too small to reliably estimate its impact on the joint intensity probability and hence on the registration criterion.

For each realization of the deformation field, MI can be computed globally over the entire image domain using a single, discrete joint histogram that is constructed by mapping all voxels of the warped image into the target image and binning of corresponding voxel intensities after appropriate interpolation of the target image. This approach is typically used with spline-based representations of the deformation field. The effect on the registration criterion of a change in any of the deformation parameters can be evaluated by recomputing the mapping and updating the joint histogram accordingly, which can be performed efficiently if each parameter has only local support and only a subregion of the image has to be considered each time [96]. With a nonparametric representation of the deformation field, a block matching strategy can be used whereby the image is subdivided in a number of regions that are each iteratively displaced individually such as to maximize MI globally [30], [52]. The obtained displacements located at the region centers are regularized and propagated to all voxels in the image, for instance by convolution with a Gaussian kernel. When evaluating MI globally over the entire image domain, a local update of the deformation field affects the registration everywhere in the image through its impact on the joint intensity histogram. Hence, in contrast with registration criteria that are evaluated in each voxel separately and independently of all others, such as minimizing the sum of squared differences of corresponding voxel intensities [14], [111], [112], registration evidence at one site is propagated to other sites because of the statistical nature of the MMI criterion.

Instead of computing MI globally over the entire image domain for each change in the deformation parameters, some approaches have been presented that compute and optimize MI locally over subregions of the image domain only, re-estimating the joint histogram for each region separately [35], [36], [63], [56], [76], [53]. A multiresolution strategy is usually adopted whereby initially a small number of larger regions is considered for coarse scale matching whose size is decreased and number increased at later stages of the optimization for matching at finer scales. Regularization of the deformation field can be achieved by imposing local smoothness [36], [53], [76] or elasticity [35] constraints or local consistency criteria [56], while spline interpolation [56], [76] or Gaussian convolution [53] can be used to propagate the deformation field over the entire image domain. However, local computation of MI within small image regions is ill-conditioned due to the small number of samples available from which to estimate the local joint image intensity distribution.

Relatively large regions are required for the local MI estimate to be reliable, such that small scale deformations are difficult to recover. Some measures have been proposed to overcome this problem, for instance by taking also the global intensity distribution into account [63], [56].

Free-form deformation approaches based on physical models require an appropriate definition for the force field at each voxel in the warped image that drives the deformation such that MI is maximized. Such force field can be derived from the gradient of MI with respect to each voxel's deformation vector. This gradient can be computed empirically, by displacing a region of interest of appropriate size around each voxel [30], [52], or analytically, using Parzen windowing to construct a joint intensity histogram that is differentiable with respect to the deformation field [39], [21].

C. Applications

Several authors have demonstrated the use of nonrigid image registration using MMI for unimodal and multimodal intrasubject image registration, to quantify or compensate for tissue motion or deformation or to correct for imaging distortion. Meyer *et al.* [73] demonstrated the accuracy and clinical versatility of mutual information for automatic multimodality image registration using affine and thin-plate spline warped geometric deformations, including thoracic PET/CT and abdominal SPECT/CT image registration. Rueckert *et al.* [96] evaluated their free-form deformation registration algorithm based on B-splines for nonrigid motion correction of 3-D MR breast images. Holden *et al.* [44] used the algorithm of [96] to quantify small changes in brain ventricular volume in serial MR images of a patient and a control group. The same method was applied by Hill *et al.* [41] for nonrigid registration of pre- and post-resection interventional MR brain images to quantify intraoperative brain deformation. D'Agostino *et al.* [21] applied their viscous fluid based nonrigid MMI approach for correction of breathing-induced deformations in MR lung ventilation images. Studholme *et al.* [101] used a method similar to [96] to correct for geometric distortion of functional MR images acquired using echo planar imaging (EPI).

Nonrigid intersubject registration using MMI has been applied for the construction of atlas templates by Rueckert *et al.* [95] and for the detection and quantification of morphological differences between subjects by Studholme *et al.* [100]. Atlas-based segmentation using MMI driven nonrigid atlas to patient registration has been applied by Castellano-Smith *et al.* [11] for brain segmentation and by D'Agostino *et al.* [21] for brain tissue classification (Fig. 15). Kjems *et al.* [52] compared various registration measures and alignment schemes, including nonrigid registration using MMI, for interindividual registration of PET activation studies.

VI. CONCLUSION

Image registration by maximization of mutual information considers all voxels in the images to be registered to estimate




Fig. 15. Illustration of nonrigid atlas to patient matching for brain tissue segmentation using the method of Van Leemput *et al.* [119]. The segmentation method uses an atlas with prior probability maps of the distribution of the various tissue classes to initialize an iterative intensity-based pixel classification scheme, assuming a Gaussian intensity distribution for each class. (a) T1-weighted MR image of a patient showing enlarged ventricles. (b) CSF segmentation obtained with affine atlas registration. (c) CSF segmentation obtained with nonrigid atlas matching. Affine atlas matching fails to properly initialize the classification and the intensity parameters of each tissue class due to the poor overlap between the ventricles in the patient and in the atlas image. These morphological differences are compensated for by nonrigid atlas matching, yielding a much improved segmentation result.

the statistical dependence between corresponding voxel intensities, which is assumed to be maximal when the images are correctly aligned. The MMI criterion is histogram based rather than intensity based and does not impose limiting assumptions on the specific nature of the relationship between corresponding voxel intensities, making it applicable to a large variety of multimodality image combinations of various body regions. Since its introduction in 1995 by Collignon *et al.* [17] and by Viola and Wells [122] in the medical imaging research community, MMI has attracted large interest within this field and has sparked a flurry of publications on multimodal medical image registration and its clinical applications. Its mathematical elegance and algorithmic simplicity made it easy for other groups to quickly adopt and implement the new approach. The success of MMI for multimodal image registration can be explained by the fact that it got rid of the need for image segmentation or preprocessing as required with previous registration algorithms and that it allows for completely automated registration without need for user interaction, making the method very well suited for application in clinical practice. In the mean time, MMI has become a standard in the field and is the method of choice for multimodal image registration in a wide range of applications. The validity of the method was demonstrated for rigid body registration of CT, MR, and PET brain images by comparison with external marker based registration. New applications, involving affine registration of multimodal images of other organs such as the thorax or the pelvis, have successfully been dealt with using MMI. Apart from intra-subject registration, MMI has also been found to be well suited for affine registration of MR images of the brain of different subjects in the context of atlas-based image analysis. Several approaches have been presented for extension of the MMI criterion to nonrigid image matching in the context of image rectification, shape normalization, motion estimation or tissue deformation correction, which is still an active area of research.

REFERENCES

- [1] J. Ashburner and K. Friston, "Nonlinear spatial normalization using basis functions," *Human Brain Mapping*, vol. 7, no. 4, pp. 254–266, 1999.
- [2] J. Ashburner and K. J. Friston, "Multimodal image coregistration and partitioning—A unified framework," *NeuroImage*, vol. 6, no. 3, pp. 209–217, 1997.
- [3] R. K. Bajcsy and S. Kovacic, "Multiresolution elastic matching," *Comput. Vis., Graph. Image Process.*, vol. 46, pp. 1–21, 1989.
- [4] B. H. Brinkmann, T. J. O'Brien, S. Aharon, M. K. O'Connor, B. P. Mullan, D. P. Hanson, and R. A. Robb, "Quantitative and clinical analysis of SPECT image registration for epilepsy studies," *J. Nucl. Med.*, vol. 40, no. 5, pp. 1098–1105, 1999.
- [5] M. Bro-Nielsen, "Rigid registration of CT, MR and cryosection images using a GLCM framework," in *Proc. 1st Joint Conf. Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery (CVRMED-MRCAS'97)*, vol. 1205, Lecture Notes in Computer Science, J. Troccaz, E. Grimson, and R. Mösges, Eds., Grenoble, France, Mar. 1997, pp. 171–180.
- [6] M. Bro-Nielsen and C. Gramkow, "Fast fluid registration of medical images," in *Proc. Visualization in Biomedical Computing (VBC'96)*, vol. 1131, Lecture Notes in Computer Science, R. Kikinis and K. H. Hoehne, Eds., Hamburg, Germany, Sept. 1996, pp. 267–276.
- [7] L. G. Brown, "A survey of image registration techniques," *ACM Comput. Surv.*, vol. 24, no. 4, pp. 325–376, Dec. 1992.
- [8] T. Butz and J.-P. Thiran, "Affine registration with feature space mutual information," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'00)*, vol. 1935, Lecture Notes in Computer Science, S. L. Delp, A. M. DiGioia, and B. Jaramaz, Eds., Berlin, Germany, Oct. 2000, pp. 549–556.
- [9] T. M. Buzug, J. Weese, C. Fassnacht, and C. Lorenz, "Image registration: Convex weighting functions for histogram-based similarity measures," in *Proc. 1st Joint Conference Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery (CVRMED-MRCAS'97)*, vol. 1205, Lecture Notes in Computer Science, J. Troccaz, E. Grimson, and R. Mösges, Eds., Grenoble, France, Mar. 1997, pp. 203–212.
- [10] M. Čapek and I. Krekule, "Alignment of adjacent picture frames captured by a CLSM," *IEEE Trans. Inform. Technol. Biomed.*, vol. 3, pp. 119–124, June 1999.
- [11] A. D. Castellano-Smith, T. Hartkens, J. A. Schnabel, D. R. Hose, H. Liu, W. A. Hall, C. L. Truwit, D. J. Hawkes, and D. L. G. Hill, "Constructing patient specific models for correcting intraoperative brain deformation," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI '01)*, vol. 2208, Lecture Notes in Computer Science, W. J. Niessen and M. A. Viergever, Eds., Utrecht, The Netherlands, Oct. 2001, pp. 1091–1098.
- [12] H. Chang and J. M. Fitzpatrick, "A technique for accurate magnetic resonance imaging in the presence of field inhomogeneities," *IEEE Trans. Med. Imag.*, vol. 11, pp. 319–329, Sept. 1992.
- [13] J. Y. Chiang and B. J. Sullivan, "Coincident bit counting – A new criterion for image registration," *IEEE Trans. Med. Imag.*, vol. 12, pp. 30–38, Mar. 1993.
- [14] G. E. Christensen, R. D. Rabbitt, and M. I. Miller, "Deformable templates using large deformation kinetics," *IEEE Trans. Image Processing*, vol. 5, pp. 1435–1447, Oct. 1996.
- [15] G. E. Christensen, R. D. Rabbitt, M. I. Miller, S. C. Joshi, U. Grenander, T. A. Coogan, and D. C. Van Essen, "Topological properties of smooth anatomic maps," in *Proc. Information Processing in Medical Imaging (IPMI'95)*, vol. 3, Computational Imaging and Vision, Y. Bizais, C. Barillot, and R. Di Paola, Eds., Ile de Berder, France, June 1995, pp. 101–112.
- [16] M. J. Clarkson, D. Rueckert, A. P. King, P. J. Edwards, D. L. G. Hill, and D. J. Hawkes, "Registration of video images to tomographic images by optimising mutual information using texture mapping," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'99)*, vol. 1679, Lecture Notes in Computer Science, C. Taylor and A. Colchester, Eds., Cambridge, U.K., Sept. 1999, pp. 579–588.
- [17] 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. Information Processing in Medical Imaging (IPMI'95)*, vol. 3, Computational Imaging and Vision, Y. Bizais, C. Barillot, and R. Di Paola, Eds., Ile de Berder, France, June 1995, pp. 263–274.

- [18] A. Collignon, D. Vandermeulen, P. Suetens, and G. Marchal, "3-D multimodality medical image registration using feature space clustering," in *Proc. First Int. Conf. Computer Vision, Virtual Reality and Robotics in Medicine (CVRMED'95)*, vol. 905, Lecture Notes in Computer Science, N. Ayache, Ed., Nice, France, Apr. 1995, pp. 195–204.
- [19] D. L. Collins, A. P. Zijdenbos, W. F. C. Barré, and A. C. Evans, "ANIMAL+INSECT: Improved cortical structure segmentation," in *Proc. Information Processing in Medical Imaging (IPMI'99)*, vol. 1613, Lecture Notes in Computer Science, A. Kuba, M. Samal, and A. Todd-Pokropek, Eds., Viségrad, Hungary, June 1999, pp. 210–223.
- [20] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [21] E. D'Agostino, F. Maes, D. Vandermeulen, and P. Suetens, "A viscous fluid model for multimodal nonrigid image registration using mutual information," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'02)*, vol. 2489, Lecture Notes in Computer Science, Tokyo, Japan, Sept. 2002, pp. 541–548.
- [22] D. Dean, J. Kamath, J. L. Duerk, and E. Ganz, "Validation of object-induced MR distortion correction for frameless stereotactic neurosurgery," *IEEE Trans. Med. Imag.*, vol. 17, pp. 810–816, May 1998.
- [23] M. Debois, R. Oyen, F. Maes, G. Verswijvel, G. Gatti, H. Bosmans, M. Feron, E. Bellon, G. Kutcher, H. V. Poppel, and L. Vanuytsel, "The contribution of magnetic resonance imaging to the three-dimensional treatment planning of localized prostate cancer," *Int. J. Radiat. Oncol. Biol. Phys.*, vol. 45, no. 4, pp. 857–865, Nov. 1999.
- [24] L. Ding, A. Goshtasby, and M. Satter, "Volume image registration by template matching," *Imag. Vis. Comput.*, vol. 19, no. 12, pp. 821–832, 2001.
- [25] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*. New York: Wiley, 1973.
- [26] M. Ferrant, A. Nabavi, B. Macq, F. A. Jolesz, R. Kikinis, and S. K. Warfield, "Registration of 3-D intraoperative MR images of the brain using a finite element biomechanical model," *IEEE Trans. Med. Imag.*, vol. 20, pp. 1384–1397, Dec. 2001.
- [27] L. Freire, A. Roche, and J.-F. Mangin, "What is the best similarity measure for motion correction in fMRI a time series?," *IEEE Trans. Med. Imag.*, vol. 21, pp. 470–484, May 2002.
- [28] K. J. Friston, J. Ashburner, C. D. Frith, J.-B. Poline, J. D. Heather, and R. S. Frackowiak, "Spatial registration and normalization of images," *Human Brain Mapping*, vol. 2, pp. 165–189, 1995.
- [29] K. J. Friston, J. Ashburner, A. Holmes, J.-B. Poline, C. Büchel, C. J. Price, R. Turner, A. Howseman, G. E. Rees, J. D. Greene, and O. Josephs. (1996) "SPM short course". [Online]. Available: <http://www.fil.ion.ucl.ac.uk>
- [30] T. Gaens, F. Maes, D. Vandermeulen, and P. Suetens, "Non-rigid multimodal image registration using mutual information," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'98)*, vol. 1496, Lecture Notes in Computer Science, W. M. Wells, A. Colchester, and S. Delp, Eds., Cambridge, MA, Oct. 1998, pp. 1099–1106.
- [31] J. C. Gee and R. K. Bajcsy, "Elastic matching: Continuum mechanical and probabilistic analysis," in *Brain Warping*, A. W. Toga, Ed. San Diego, CA: Academic, 1999.
- [32] P. Gerlot-Chiron and Y. Bizais, "Registration of multimodality medical images using region overlap criterion," *CVGIP: Graph. Models Image Process.*, vol. 54, no. 5, pp. 396–406, Sept. 1992.
- [33] P. L. Gildenberg and R. R. Tasker, *Textbook of Stereotactic and Functional Neurosurgery*. New York: McGraw-Hill, 1998.
- [34] A. Hagemann, K. Rohr, H. S. Stiel, U. Spetzger, and J. M. Gilsbach, "Biomechanical modeling of the human head for physically based, nonrigid image registration," *IEEE Trans. Med. Imag.*, vol. 18, pp. 875–884, Oct. 1999.
- [35] N. Hata, T. Dohi, S. Warfield, W. Wells III, R. Kikinis, and F. A. Jolesz, "Multimodality deformable registration of pre-and intraoperative images for MRI-guided brain surgery," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'98)*, vol. 1496, Lecture Notes in Computer Science, W. M. Wells, A. Colchester, and S. Delp, Eds., Cambridge, MA, Oct. 1998, pp. 1067–1074.
- [36] P. Hellier and C. Barillot, "Multimodal nonrigid warping for correction of distortions in functional MRI," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'00)*, vol. 1935, Lecture Notes in Computer Science, S. L. Delp, A. M. Di-Gioia, and B. Jaramaz, Eds., Berlin, Oct. 2000, pp. 512–520.
- [37] P. F. Hemler, T. S. Sumanaweera, P. A. van den Elsen, S. Napel, and J. R. Adler, "A versatile system for multimodality image fusion," *J. Image Guided Surg.*, vol. 1, pp. 35–45, 1995.
- [38] P. F. Hemler, P. A. van den Elsen, T. S. Sumanaweera, S. Napel, J. Drace, and J. R. Adler, "A quantitative comparison of residual error for three different multimodality registration techniques," in *Proc. 14th Int. Conf. Information Processing in Medical Imaging (IPMI'95)*, vol. 3, Computational Imaging and Vision, Y. Bizais, C. Barillot, and R. Di Paola, Eds., Ile de Berder, France, June 1995, pp. 251–262.
- [39] G. Hermosillo, C. Chef d'Hotel, and O. Faugeras, "Variational methods for multimodal image matching," *Int. J. Comput. Vis.*, vol. 50, no. 3, pp. 329–343, 2002.
- [40] D. L. G. Hill, D. J. Hawkes, N. A. Harrison, and C. F. Ruff, "A strategy for automated multimodality image registration incorporating anatomical knowledge and imager characteristics," in *Proc. 13th Int. Conf. Information Processing in Medical Imaging (IPMI'93)*, vol. 687, Lecture Notes in Computer Science, H. H. Barrett and A. F. Gmitro, Eds., Flagstaff, AZ, June 1993, pp. 182–196.
- [41] D. L. G. Hill, C. R. Maurer Jr., A. J. Martin, S. Sabanathan, W. A. Hall, D. J. Hawkes, D. Rueckert, and C. L. Truwit, "Assessment of intraoperative brain deformation using interventional MR imaging," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'99)*, vol. 1679, Lecture Notes in Computer Science, C. Taylor and A. Colchester, Eds., Cambridge, U.K., Sept. 1999, pp. 910–919.
- [42] D. L. G. Hill, C. Studholme, and D. J. Hawkes, "Voxel similarity measures for automated image registration," in *Proc. Visualization in Biomedical Computing (VBC'94)*, vol. 2359, Proc. SPIE, 1994, pp. 205–216.
- [43] M. Holden, D. L. G. Hill, E. R. E. Denton, J. M. Jarosz, T. C. S. Cox, T. Rohlfing, J. Goodey, and D. J. Hawkes, "Voxel similarity measures for 3-D serial MR brain image registration," *IEEE Trans. Med. Imag.*, vol. 19, pp. 94–102, July 2000.
- [44] M. Holden, J. A. Schnabel, and D. L. G. Hill, "Quantifying small changes in brain ventricular volume using nonrigid registration," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'01)*, vol. 2208, Lecture Notes in Computer Science, W. J. Niessen and M. A. Viergever, Eds., Utrecht, The Netherlands, Oct. 2001, pp. 49–56.
- [45] M. A. Horsfield, "Mapping eddy current induced fields for the correction of diffusion-weighted echo planar images," *Magn. Reson. Imag.*, vol. 17, no. 9, pp. 1335–1345, 1999.
- [46] L.-Y. Hsu and M. H. Loew, "Fully automatic 3D feature-based registration of multimodality medical images," *Image Vis. Comput.*, vol. 19, no. 1–2, pp. 75–85, 2001.
- [47] M. Jenkinson and S. Smith, "A global optimization method for robust affine registration of brain images," *Med. Image Anal.*, vol. 5, pp. 143–156, 2001.
- [48] H. Jiang, K. S. Holton, and R. A. Robb, "Image registration of multimodality 3-D medical images by chamfer matching," in *Proc. SPIE, Medical Imaging: Image Processing*, 1992, vol. 1660, pp. 356–366.
- [49] H. Jiang, R. A. Robb, and K. S. Holton, "A new approach to 3-D registration of multimodality medical images by surface matching," in *Proc. SPIE, Proc. Visualization in Biomedical Computing (VBC'92)*, vol. 1808, , 1992, pp. 196–213.
- [50] J. M. Fitzpatrick, "Retrospective Image Registration Evaluation, National Institutes of Health," Vanderbilt Univ., Nashville, TN, 1 R01 CA89323, 1994.
- [51] B. Kim, J. L. Boes, and C. R. Meyer, "Mutual information for automated multimodal image warping," *NeuroImage*, vol. 3, no. 3, pp. 158–158, June 1996.
- [52] U. Kjems, S. C. Strother, J. Anderson, I. Law, and L. K. Hansen, "Enhancing the multivariate signal of [¹⁵O] water PET a studies with a new nonlinear neuroanatomical registration algorithm," *IEEE Trans. Med. Imag.*, vol. 18, pp. 306–319, Apr. 1999.
- [53] Y. H. Lau, M. Braun, and B. F. Hutton, "Non-rigid image registration using a median-filtered coarse-to-fine displacement field and a symmetric correlation ratio," *Phys. Med. Biol.*, vol. 46, no. 4, pp. 1297–1319, 2001.
- [54] D. Lemoine, D. Liegeard, E. Lussot, and C. Barillot, "Multimodal registration system for the fusion of MRI, CT, MEG, and 3-D or stereotactic angiographic data," in *Proc. SPIE Medical Imaging*, vol. 2164, Proc. SPIE, 1994, pp. 46–56.
- [55] B. Likar and F. Pernuš, "Registration of serial transverse sections of muscle fibers," *Cytometry*, vol. 37, no. 2, pp. 93–106, 1999.
- [56] —, "A hierarchical approach to elastic registration based on mutual information," *Image Vision Comput.*, vol. 19, no. 1–2, pp. 33–44, 2001.

- [57] A. M. López, D. Lloret, J. Serrat, and J. J. Villanueva, "Multilocal creaseness based on the level-set extrinsic curvature," *Comput. Vision Image Understand.*, vol. 77, no. 2, pp. 111–144, 2000.
- [58] F. Maes, "Segmentation and registration of multimodal medical images: From theory, implementation and validation to a useful tool in clinical practice," Ph.D. dissertation, Dept. Elect. Eng. (ESAT/PSI), KU Leuven, Leuven, Belgium, 1998.
- [59] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *IEEE Trans. Med. Imag.*, vol. 16, pp. 187–198, Apr. 1997.
- [60] 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.
- [61] G. Q. Maguire Jr. et al., "Graphics applied to medical image registration," *IEEE Comput. Graph. Applicat.*, vol. 11, pp. 20–28, Mar. 1991.
- [62] J. B. A. Maintz, "Retrospective registration of tomographic brain images," Ph.D. dissertation, Univ. Utrecht, Utrecht, The Netherlands, 1996.
- [63] J. B. A. Maintz, E. H. W. Meijering, and M. A. Viergever, "General multimodal elastic registration based on mutual information," *Proc. SPIE, Medical Imaging: Image Processing*, vol. 3338, pp. 144–154, 1998.
- [64] J. B. A. Maintz, P. A. van den Elsen, and M. A. Viergever, "Comparison of feature-based matching of CT and MR brain images," in *Proc. 1st Int. Conf. Computer Vision, Virtual Reality, Robotics in Medicine*, vol. 905, Lecture Notes in Computer Science, N. Ayache, Ed., Nice, France, Apr. 1995, pp. 219–228.
- [65] —, "Comparison of edge-based and ridge-based registration of CT and MR brain images," *Med. Image Anal.*, vol. 1, pp. 151–161, 1996.
- [66] —, "Evaluation of ridge seeking operators for multimodality medical image matching," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 18, pp. 353–365, Apr. 1996.
- [67] G. Malandain, S. Fernández-Vidal, and J. M. Rocchisani, "Improving registration of 3-D medical images using a mechanical based method," in *Proc. 3rd European Conf. Computer Vision (ECCV '94)*, 1994, pp. 131–136.
- [68] —, "Rigid registration of 3-D objects by motion analysis," in *Proc. 12th Int. Conf. Pattern Recognition*, 1994, pp. 579–581.
- [69] —, "Physically based rigid registration of 3-D free form objects: Applications to medical imaging," INRIA, Sophia Antipolis Cedex, France, 1995.
- [70] C. R. Maurer and J. M. Fitzpatrick, "A review of medical image registration," in *Interactive Image-Guided Neurosurgery*, R. J. Maciunas, Ed. Stuttgart, Germany: Thieme (distributed for Amer. Assoc. Neurolog. Surg.), 1993, pp. 17–44.
- [71] C. R. Maurer Jr., G. B. Aboutanos, B. M. Dawant, S. Gadamsetty, R. A. Margolin, R. J. Maciunas, and J. M. Fitzpatrick, "Effect of geometrical distortion correction in MR on image registration accuracy," *J. Comput. Assist. Tomogr.*, vol. 20, pp. 666–679, 1996.
- [72] C. R. Maurer Jr., J. M. Fitzpatrick, M. Y. Wang, R. L. Galloway, R. J. Maciunas, and G. S. Allen, "Registration of head volume images using implantable fiducial markers," *IEEE Trans. Med. Imag.*, vol. 16, pp. 447–462, Aug. 1997.
- [73] C. Meyer, J. L. Boes, B. Kim, P. H. Bland, R. L. Wahl, K. R. Zasadny, P. V. Kison, K. Koral, and K. A. Frey, "Demonstration of accuracy and clinical versatility of mutual information for automatic multimodality image fusion using affine and thin plate spline warped geometric deformations," *Med. Image Anal.*, vol. 1, no. 3, pp. 195–206, 1997.
- [74] J. Michiels, P. Pelgrims, H. Bosmans, D. Vandermeulen, J. Gybels, G. Marchal, and P. Suetens, "On the problem of geometric distortion in magnetic resonance images for stereotactic neurosurgery," *Magn. Reson. Imag.*, vol. 12, no. 5, pp. 749–765, 1994.
- [75] E. D. Morris, G. J. Muswick, E. S. Ellert, R. N. Steagall, P. F. Goyer, and W. E. Semple, "Computer-aided techniques for aligning interleaved sets of nonidentical medical images," *Proc. SPIE Medical Imaging: Image Processing*, vol. 1898, pp. 146–157, Feb. 1993.
- [76] M. Otte, "Elastic registration of fMRI data using Bézier-spline transformations," *IEEE Trans. Med. Imag.*, vol. 20, pp. 193–206, Mar. 2001.
- [77] C. A. Pelizzari, G. T. Y. Chen, D. R. Spelbring, R. R. Weichselbaum, and C.-T. Chen, "Accurate three-dimensional registration of CT, PET, and/or MR images of the brain," *J. Comput. Assist. Tomogr.*, vol. 13, pp. 20–26, 1989.
- [78] G. P. Penney, J. Weese, J. A. Little, P. Desmedt, D. L. G. Hill, and D. J. Hawkes, "A comparison of similarity measures for use in 2D-3D medical image registration," *IEEE Trans. Med. Imag.*, vol. 17, pp. 586–595, Apr. 1999.
- [79] U. Pietrzyk, K. Herholz, G. Fink, A. Jacobs, R. Mielke, I. Slansky, M. Würker, and W. D. Heiss, "An interactive technique for three-dimensional image registration: Validation for PET, SPECT, MRI and CT brain studies," *J. Nucl. Med.*, vol. 35, pp. 2011–2018, 1994.
- [80] U. Pietrzyk, K. Herholz, and W. D. Heiss, "Three-dimensional alignment of functional and morphological tomograms," *J. Comput. Assist. Tomogr.*, vol. 14, no. 1, pp. 51–59, 1990.
- [81] U. Pietrzyk, K. Herholz, A. Schuster, H.-M. v. Stockhausen, H. Lucht, and W.-D. Heiss, "Clinical applications of registration and fusion of multimodality brain images from PET, SPECT, CT and MRI," *Eur. J. Radiol.* 1996, vol. 21, pp. 174–182, 1996.
- [82] J. P. W. Pluim, "Multi-modality matching using mutual information," M.S. thesis, Dept. Comput. Sci., Univ. Groningen, Groningen, The Netherlands, 1996.
- [83] —, "Mutual information based registration of medical images," Ph.D. dissertation, Utrecht Univ., Utrecht, The Netherlands, 2001.
- [84] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever, "Image registration by maximization of combined mutual information and gradient information," *IEEE Trans. Med. Imag.*, vol. 19, pp. 809–814, Aug. 2000.
- [85] —, "Interpolation artefacts in mutual information-based image registration," *Comput. Vis. Image Understand.*, vol. 77, no. 2, pp. 211–232, 2000.
- [86] —, "f-information measures in medical image registration," in *Proc. SPIE, Medical Imaging: Image Processing*, 2001, vol. 4322, pp. 579–587.
- [87] —, "Mutual information matching in multiresolution contexts," *Image Vis. Comput.*, vol. 19, no. 1–2, pp. 45–52, 2001.
- [88] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C*, Second ed. Cambridge, U.K.: Cambridge Univ. Press, 1992.
- [89] P. E. Radau, P. J. Slomka, P. Julin, L. Svensson, and L.-O. Wahlund, "Automated segmentation and registration technique for hmpao-spect imaging of alzheimer's patients," *Proc. SPIE, Medical Imaging: Image Processing*, vol. 3979, pp. 372–384, 2000.
- [90] T. Radcliffe, R. Rajapakshe, and S. Shalev, "Pseudocorrelation: A fast, robust, absolute, grey-level image alignment algorithm," *Med. Phys.*, vol. 21, no. 6, pp. 761–769, June 1994.
- [91] C. E. Rodriguez-Carranza and M. H. Loew, "Weighted and deterministic entropy measure for image registration using mutual information," *Proc. SPIE, Medical Imaging: Image Processing*, vol. 3338, pp. 155–166, 1998.
- [92] G. K. Rohde, A. Aldroubi, and B. M. Dawant, "Adaptive free-form deformation for inter-patient medical image registration," *Proc. SPIE, Medical Imaging: Image Processing*, vol. 4322, pp. 1578–1587, 2001.
- [93] J.-M. Rouet, J.-J. Jacq, and C. Roux, "Genetic algorithms for a robust 3-d mr-ct registration," *IEEE Trans. Inform. Technol. Biomed.*, vol. 4, pp. 126–136, June 2000.
- [94] D. Rueckert, M. J. Clarkson, D. L. G. Hill, and D. J. Hawkes, "Non-rigid registration using higher order mutual information," *Proc. SPIE, Medical Imaging: Image Processing*, vol. 3979, pp. 438–447, 2000.
- [95] D. Rueckert, A. F. Frangi, and J. A. Schnabel, "Automatic construction of 3D statistical deformation models using nonrigid registration," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI '01)*, vol. 2208, Lecture Notes in Computer Science, W. J. Niessen and M. A. Viergever, Eds., Utrecht, The Netherlands, Oct. 2001, pp. 77–94.
- [96] D. Rueckert, L. I. Sonoda, C. Hayes, D. 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, pp. 712–721, Aug. 1999.
- [97] J. A. Schnabel, D. Rueckert, M. Quist, J. M. Blackall, A. D. Castellano-Smith, T. Hartkens, G. P. Penney, W. A. Hall, H. Liu, C. L. Truweit, F. A. Gerritsen, D. L. G. Hill, and D. J. Hawkes, "A generic framework for nonrigid registration based on nonuniform multilevel free-form deformations," in *Proc. SpMedical Image Computing and Computer-Assisted Intervention (MICCAI '01)*, vol. 2208, Lecture Notes in Computer Science, W. J. Niessen and M. A. Viergever, Eds., Utrecht, The Netherlands, Oct. 2001, pp. 573–581.

- [98] M. Soltys, D. V. Beard, V. Carrasco, S. Mukherji, and J. Rosenman, "FUSION: A tool for registration and visualization of multiple modality 3-D medical data," *Proc. SPIE, Medical Imaging 1995: Image Processing*, vol. 2434, pp. 74–80, Feb. 1995.
- [99] C. Studholme, "Measures of 3-D medical image alignment," Ph.D. dissertation, Univ. London, United Medical and Dental Schools of Guy's and St Thomas's Hospitals, Division of Radiological Sciences, Computational Imaging Science Group, London, U.K., 1997.
- [100] C. Studholme, V. Cardenas, N. Schuff, H. Rosen, B. Miller, and M. Weiner, "Detecting spatially consistent structural differences in Alzheimer's and fronto temporal dementia using deformation morphometry," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI '01)*, vol. 2208, Lecture Notes in Computer Science, W. J. Niessen and M. A. Viergever, Eds., Utrecht, The Netherlands, Oct. 2001, pp. 41–48.
- [101] C. Studholme, R. T. Constable, and J. S. Duncan, "Accurate alignment of functional EPI data to anatomical MRI using a physics-based distortion mode," *IEEE Trans. Med. Imag.*, vol. 19, pp. 1115–1127, Nov. 2000.
- [102] C. Studholme, D. L. G. Hill, and D. J. Hawkes, "Automated 3-D registration of truncated MR and CT images of the head," in *Proc. British Machine Vision Conf.*, vol. 1, D. Pycock, Ed., Sept. 1995, pp. 27–36.
- [103] —, "Multiresolution voxel similarity measures for MR-PET registration," in *Proc. 14th Int. Conf. Information Processing in Medical Imaging (IPMI'95)*, vol. 3, Computational Imaging and Vision, Y. Bizais, C. Barillot, and R. Di Paola, Eds., Ile de Berder, France, June 1995, pp. 287–298.
- [104] —, "Automated 3-D registration of MR and CT images of the head," *Med. Image Anal.*, vol. 1, no. 2, pp. 163–175, 1996.
- [105] —, "Incorporating connected region labeling into automated image registration using mutual information," *Proc. 2nd IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pp. 23–31, June 1996.
- [106] —, "An overlap invariant entropy measure of 3-D medical image alignment," *Pattern Recognit.*, vol. 32, no. 1, pp. 71–86, 1999.
- [107] C. Studholme, D. L. G. Hill, J. Wong, M. N. Naisey, and D. J. Hawkes, "Registration measures for automated 3-D alignment of PET and intensity distorted MR images," in *Proc. Image Fusion and Shape Variability Techniques*, K. V. Mardia, C. A. Gill, and I. L. Dryden, Eds., Leeds, U.K., July 1996, pp. 186–193.
- [108] T. S. Sumanaweera, J. R. Adler Jr., S. Napel, and G. H. Glover, "Characterization of spatial distortion in magnetic resonance imaging and its implications for stereotactic surgery," *Neurosurg.*, vol. 35, no. 4, pp. 696–703, 1994.
- [109] T. S. Sumanaweera, G. H. Glover, T. O. Binford, and J. R. Adler, "MR susceptibility misregistration correction," *IEEE Trans. Med. Imag.*, vol. 12, pp. 251–259, Feb. 1993.
- [110] P. Thévenaz and M. Unser, "Optimization of mutual information for multiresolution image registration," *IEEE Trans. Image Process.*, vol. 9, pp. 2083–2099, Dec. 2000.
- [111] J.-P. Thirion, "Non-rigid matching using demons," in *Proc. Int. Conf. Computer Vision and Pattern Recognition (CVPR'96)*, San Francisco, CA, June 1996, pp. 245–251.
- [112] —, "Image matching as a diffusion process: An analogy with Maxwell's demons," *Med. Image Anal.*, vol. 2, no. 3, pp. 243–260, 1998.
- [113] L. Thurfjell, Y. H. Lau, J. L. R. Andersson, and B. F. Hutton, "Improved efficiency for MRI-SPET registration based on mutual information," *Eur. J. Nucl. Med.*, vol. 27, no. 7, pp. 847–856, 2000.
- [114] I. Vajda, *Theory of Statistical Inference and Information*. Dordrecht, The Netherlands: Kluwer, 1989.
- [115] P. A. van den Elsen, J. B. A. Maintz, E.-J. D. Pol, and M. A. Viergever, "Grey value correlation techniques used for automatic matching of CT and MR brain and spine images," *Proc. SPIE, Visualization in Biomedical Computing (VBC'94)*, vol. 2359, pp. 227–237, Oct. 1994.
- [116] —, "Automatic registration of CT and MR brain images using correlation of geometrical features," *IEEE Trans. Med. Imag.*, vol. 14, no. 2, June 1995.
- [117] P. A. van den Elsen, E.-J. D. Pol, and M. A. Viergever, "Medical image matching—A review with classification," *IEEE Eng. Med. Biol. Mag.*, pp. 26–38, Mar. 1993.
- [118] K. Van Leemput, F. Maes, D. Vandermeulen, and P. Suetens, "Automated model-based bias field correction of MR images of the brain," *IEEE Trans. Med. Imag.*, vol. 18, no. 10, pp. 885–896, 1999.
- [119] —, "Automated model-based tissue classification of MR images of the brain," *IEEE Trans. Med. Imag.*, vol. 18, pp. 897–908, Oct. 1999.
- [120] J. Vansteenkiste, S. Stroobants, P. Dupont, P. De Leyn, W. De Wever, E. Verbeken, J. Nuyts, F. Maes, and J. Bogaert, "FDG-PET scan in potentially operable nonsmall cell lung cancer: Do anatomical PET-CT fusion images improve the localization of regional lymph node metastases?," *Eur. J. Nucl. Med.*, vol. 25, no. 11, pp. 1495–1501, 1998.
- [121] A. Venot, J. F. Lebruchec, and J. C. Roucauyrol, "A new class of similarity measures for robust image registration," *Comput. Vis. Graph. Image Process.*, vol. 28, no. 2, pp. 176–184, Nov. 1984.
- [122] P. Viola and W. M. Wells III, "Alignment by maximization of mutual information," in *Proc. 5th Int. Conf. Computer Vision*, Cambridge, MA, June 1995, pp. 16–23.
- [123] M. Y. Wang, C. R. Maurer Jr., J. M. Fitzpatrick, and R. J. Maciunas, "An automatic technique for finding and localizing externally attached markers in CT and MR volume images of the head," *IEEE Trans. Biomed. Eng.*, vol. 43, pp. 627–637, June 1996.
- [124] W. M. Wells III, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis, "Multimodal volume registration by maximization of mutual information," *Med. Image Anal.*, vol. 1, no. 1, pp. 35–51, Mar. 1996.
- [125] J. West, J. M. Fitzpatrick, M. Y. Wang, B. M. Dawant, C. R. Maurer Jr., R. M. Kessler, and R. J. Maciunas, "Retrospective intermodality registration techniques for images of the head: Surface-based versus volume-based," *IEEE Trans. Med. Imag.*, vol. 18, pp. 144–150, Feb. 1999.
- [126] J. West, J. M. Fitzpatrick, M. Y. Wang, B. M. Dawant, C. R. Maurer Jr., R. M. Kessler, R. J. Maciunas, C. Barillot, D. Lemoine, A. Collignon, F. Maes, P. Suetens, D. Vandermeulen, P. A. van den Elsen, S. Napel, T. S. Sumanaweera, B. Harkness, P. F. Hemler, D. L. G. Hill, D. J. Hawkes, C. Studholme, J. B. A. Maintz, M. A. Viergever, G. Malandain, X. Pennec, M. E. Noz, G. Q. Maguire Jr., M. Pollack, C. A. Pellizari, R. A. Robb, D. Hanson, and R. P. Woods, "Comparison and evaluation of retrospective intermodality brain image registration techniques," *J. Comput. Assist. Tomogr.*, vol. 21, pp. 544–566, 1997.
- [127] J. West, J. M. Fitzpatrick, M. Y. Wang, B. M. Dawant, C. R. Maurer Jr., R. M. Kessler, R. J. Maciunas, C. Barillot, D. Lemoine, A. Collignon, F. Maes, P. Suetens, D. Vandermeulen, P. A. van den Elsen, P. F. Hemler, S. Napel, T. S. Sumanaweera, B. Harkness, D. L. G. Hill, C. Studholme, G. Malandain, X. Pennec, M. E. Noz, G. Q. Maguire Jr., M. Pollack, C. A. Pellizari, R. A. Robb, D. Hanson, and R. P. Woods, "Comparison and evaluation of retrospective intermodality brain image registration techniques," in *Proc. SPIE, Medical Imaging: Image Processing*, 1996, vol. 2710, pp. 332–347.
- [128] G. Wolynski and F. Kruggel, "Computational cost of nonrigid registration algorithms based on fluid dynamics," *IEEE Trans. Med. Imag.*, vol. 21, pp. 946–952, Aug. 2002.
- [129] R. P. Woods, J. C. Mazziotta, and S. R. Cherry, "MRI-PET registration with automated algorithm," *J. Comput. Assist. Tomogr.*, vol. 17, no. 4, pp. 536–546, July/Aug. 1993.
- [130] Y.-M. Zhu and S. M. Cochoff, "Influence of implementation parameters on registration of MR and SPECT brain images by maximization of mutual information," *J. Nucl. Med.*, vol. 43, no. 2, pp. 160–166, 2002.



Frederik Maes received the M.S. degrees in electrical engineering from the Katholieke Universiteit (K.U.) Leuven, Leuven, Belgium, in 1991 and from Stanford University, Stanford, CA, in 1992, and the Ph.D. degree in electrical engineering from the K.U. Leuven in 1998.

Currently, he is a Postdoctoral Researcher at the Laboratory for Medical Image Computing [Radiology and Electrical Engineering (ESAT)] at K.U. Leuven. Since 2001, he has been a Guest Lecturer with the Faculty of Medicine and is involved in numerous imaging-related clinical research projects within the University Hospital Gasthuisberg, Leuven. His main research interests are multimodal image registration and model-based image segmentation.

Dr. Maes is a Postdoctoral Fellow of the Fund for Scientific Research—Flanders, Belgium (FWO—Vlaanderen).



Dirk Vandermeulen received the Ph.D. degree in electrical engineering from the Katholieke Universiteit (K.U.) Leuven, Leuven, Belgium, in 1991.

Since 1994, has been a Professor of Electrical Engineering at K.U. Leuven, where he teaches computer vision, medical image analysis, and biometrics. Currently, he is the Head of the Medical Image Analysis Group, Laboratory for Medical Image Computing [Radiology and Electrical Engineering (ESAT)] at K.U. Leuven.

His main research interest is fundamental research and basic algorithm development for medical image analysis and computer vision. He has authored more than 100 papers in international journals and conference proceedings about medical imaging and analysis and acts as member of the program committee of various international conferences in the field.



Paul Suetens received the Ph.D. degree in computer science from the Katholieke Universiteit (K.U.) Leuven, Leuven, Belgium, in 1983 with pioneering research in the domain of image-guided planning of neurosurgical interventions.

Since 1990, he has been a Professor of electrical engineering with K.U. Leuven, where he teaches medical imaging and medical image analysis. Currently, he is the Head of the Center for Processing Speech and Images (ESAT-PSI) and of the Laboratory for Medical Image Computing [Radiology and Electrical Engineering (ESAT)] at the University Hospital Gasthuisberg, Leuven. He has authored more than 400 papers in international journals and conference proceedings about medical imaging and analysis and acts as member of the program committee of various international conferences in the field.

Dr. Suetens is the Chairman of the Management Board of MEDICIM, a spinoff of the K.U. Leuven active in the field of image-guided surgery planning.