

# Unifying Maximum Likelihood Approaches in Medical Image Registration

Alexis Roche, Grégoire Malandain, Nicholas Ayache

INRIA, EPIDAURE Project, 2004, route des Lucioles, BP 93, 06902 Sophia Antipolis Cedex, France

**ABSTRACT:** Although intensity-based similarity measures are increasingly used for medical image registration, they often rely on implicit assumptions regarding the imaging physics. This paper clarifies the assumptions on which a number of popular similarity measures rely. After formalizing registration based on general image acquisition models, we show that the search for an optimal measure can be cast into a maximum likelihood estimation problem. We then derive similarity measures corresponding to different modeling assumptions and retrieve some well-known measures (correlation coefficient, correlation ratio, mutual information). Finally, we present results of rigid registration between several image modalities to illustrate the importance of choosing an appropriate similarity measure. © 2000 John Wiley & Sons, Inc. *Int J Imaging Syst Technol*, 11, 71–80, 2000

## I. INTRODUCTION

Registration is a problem common to many tasks in medical image analysis. It can be necessary to compare images acquired from the same patient at different times or with different sensors. One usually distinguishes registration problems according to whether the images are from the same modality (monomodal registration) or from different modalities (multimodal registration). In general, inpatient registration consists of estimating a rigid transformation between the images, but it can also involve a nonrigid transformation in order to compensate for tissue deformations or geometrical distortions inherent to the imaging processes.

Registration is also useful for comparing images acquired from different patients, e.g., to build statistical anatomical atlases. Much effort in this area has been devoted to the geometrical modeling of anatomical variations from one subject to another. In general, interpatient registration involves nonrigid transformations (Toga, 1999).

Reviews of medical image registration methods were written by van den Elsen et al. (1993), Lavalley (1995), and Maintz and Viergever (1998). Quite recently, a comparison of algorithms based on a retrospective evaluation was published by West et al. (1997) in the context of inpatient rigid registration. Registration methods are usually classified as being either feature or intensity based. Methods from the former class proceed in two sequential steps. The first is to segment homologous geometrical landmarks in the images; these can be points, lines, surfaces, or volumes. The problem then reduces to a purely geometrical task, i.e., to evaluate the transformation that best matches these landmarks. Because these methods are highly dependent on the algorithms that are used in the segmen-

tation step, they are often restricted to very specific registration problems. Likewise, when dealing with images from different modalities, finding homologous landmarks is a very challenging task due to the lack of redundancy in anatomical information.

Intensity-based techniques circumvent these difficulties because they do not deal with identifying geometrical landmarks. Their basic principle is to search, in a certain space of transformations, the one that maximizes a criterion measuring the intensity similarity of corresponding voxels. This paper focuses on this class of methods. Over the last few years, they have been applied to a number of registration problems, including monomodal, multimodal, rigid, and nonrigid registration (Maintz and Viergever, 1998). Common to the many proposed similarity measures is the idea that, when matched, the image intensities should verify a certain relationship. The similarity measure is intended to quantify how well this relationship is verified given a transformation between the images.

Choosing one measure adapted to a specific registration problem is not always straightforward for at least two reasons. First, it is often difficult to model the physical relation between the image intensities. Second, most of the similarity measures rely on imaging assumptions that are not fully explicit. We can roughly classify existing similarity measures according to four main kinds of hypotheses:

**Identity Relationship.** In this category, the basic assumption is the conservation of intensity from one image to the other. This includes a number of popular measures, e.g., the sum of squared intensity differences (SSD), the sum of absolute intensity differences, cross-correlation (Brown, 1992), and entropy of the difference image (Buzug and Weese, 1998). Although these measures are not equivalent in terms of robustness and accuracy, none of them is able to cope with relative intensity changes from one image to the other.

**Affine Relationship.** The step beyond is to assume that the two images  $I$  and  $J$  to be registered are related by an affine mapping, i.e.,  $I \approx \alpha J + \beta$ . The measures adapted to this situation are more or less variants on the correlation coefficient (Brown, 1992), defined as the ratio between the covariance of the images and the product of individual standard deviations:

$$\rho(I, J) = \frac{\text{Cov}(I, J)}{\sqrt{\text{Var}(I)} \sqrt{\text{Var}(J)}}. \quad (1)$$

---

Correspondence to: Alexis Roche

The correlation coefficient is generally useful for matching images from the same modality. Nevertheless, the affine hypothesis is hardly valid for images from different modalities, and thus it has not provided convincing results in multimodal registration.

**Functional Relationship.** For multimodal images, more complex relationships are involved. The approach we proposed in Roche et al. (1998b) was to assume that, at the registration position, one image could be approximated in terms of the other by applying some intensity function,  $I \approx \phi(J)$ . Making no assumption regarding the nature of the function, we derived a natural statistical measure, the correlation ratio:

$$\eta^2(I|J) = 1 - \frac{\text{Var}(I - \hat{\phi}(J))}{\text{Var}(I)}, \quad (2)$$

where  $\hat{\phi}(J)$  is the least squares optimal nonlinear approximation of  $I$  in terms of  $J$  (Papoulis, 1991). The correlation ratio is closely related to a very popular measure previously proposed by Woods et al. (1993) and generalized using robust metrics in Nikou et al. (1998).

**Statistical Relationship.** Finally, assuming a functional relationship is sometimes too restrictive. In these cases, it is more appropriate to use information theoretical measures; from this group, mutual information (Maes et al., 1997; Viola and Wells, 1997) is today probably the most popular:

$$\mathcal{I}(I, J) = \sum_i \sum_j p(i, j) \log \frac{p(i, j)}{p(i)p(j)}, \quad (3)$$

where  $p(i, j)$  is the intensity joint probability distribution of the images and  $p(i)$  and  $p(j)$  the corresponding marginal distributions. This category is not fundamentally different from the previous one, as the ideal case is still perfect functional dependence; mutual information is, however, theoretically more robust to variations with respect to this ideal situation.

A number of comparison studies have shown that similarity measures yield different performances depending on the considered modality combinations (West et al., 1997; Bro-Nielsen, 1997; Penney et al., 1998; Nikou et al., 1998; Roche et al., 1998b). There is probably no universal measure and, for a specific problem, the point is rather to choose the one that is best adapted to the nature of the images.

The link between explicit modeling assumptions and similarity measures has not yet been made clear. Some authors (Mort and Srinath, 1988; Costa et al., 1993) proposed that image registration could be seen as a maximum likelihood estimation problem. Others (Viola and Wells, 1997; Wells et al., 1996) suggested the analogy of this approach with registration based on information theory. Notably, other teams had motivated information-theoretical measures using different arguments (Maes et al., 1997; Studholme et al., 1996).

In Section 2, we propose to formulate image registration as a general maximum likelihood estimation problem, examining carefully the assumptions that are required. In Section 3, deriving optimal similarity measures from specific modeling assumptions, we retrieve the correlation ratio and mutual information. Section 4 proposes to illustrate the practical differences between these two

measures with results of rigid multimodal registration of three-dimensional (3D) brain images.

## II. FORMULATION

**A. Maximum Likelihood Registration.** Two images  $I$  and  $J$  to be registered, are related through the common anatomical reality that they measure. However, the way anatomical structures are represented depends on the physics of the imaging involved in each acquisition. Thus, modeling the relationship between the images requires the knowledge of both the underlying anatomy and the image formation processes. A convenient model of the anatomy will be an image called segmentation or scene: by definition, a scene is any image for which the intensity of a given voxel represents the tissue class to which it belongs.

Assuming that we know a scene, we can model indirectly the relationship between  $I$  and  $J$  based on image acquisition models. A standard approach in computer vision is to interpret an image as being a realization of a random process that corrupts the scene. This means that the relationship between  $I$  and  $S$  (respectively,  $J$  and  $S$ ) is defined in terms of a conditional probability density function  $P(I|S)$ . The two following assumptions are usually stated:

- (A1) The voxels of the image are conditionally independent knowing the scene, i.e.,

$$P(I|S) = \prod_{x_k \in \Omega_I} P(i_k|S),$$

where  $\Omega_I$  denotes the voxel grid of  $I$  and  $i_k \equiv I(x_k)$  is the intensity of the voxel with coordinates  $x_k$  in a given frame attached to the grid  $\Omega_I$ .

- (A2) The noise is context free. In other words, the intensity of a voxel depends only on its homologous in the scene:

$$P(i_k|S) = P(i_k|s_k^\downarrow), \quad \text{with } s_k^\downarrow \equiv S(T(x_k)) = (S \circ T)(x_k),$$

where  $T$  is the spatial transformation that relates the coordinate frames of  $\Omega_I$  and  $\Omega_S$ , the grid of  $S$ . In the case where  $I$  and  $S$  are not supposed to be aligned,  $T$  has no reason to be the identity. Of course, to be meaningful, the transformation  $T$  needs to be defined as a mapping from  $\Omega_I$  to  $\Omega_S$ , i.e., a grid point of  $I$  is supposed to match a grid point of  $S$ . In Section IID, we discuss how this may be achieved in practice.

Under these assumptions, the conditional probability of  $I$  knowing the scene  $S$  and the transformation  $T$  is easily seen to be:

$$P(I|S, T) = \prod_{x_k \in \Omega_I} P(i_k|s_k^\downarrow). \quad (4)$$

We can model the relationship between  $J$  and  $S$  in the same manner. However, as we are only interested in the relative displacement between  $I$  and  $J$ , we will consider  $J$  as a reference image already aligned with the scene, meaning that no transformation is involved in the conditional probability,

$$P(J|S) = \prod_{y_l \in \Omega_J} P(j_l|s_l), \quad \text{with } j_l = J(y_l), \quad s_l = S(y_l), \quad (5)$$

$\Omega_J \equiv \Omega_S$  being the voxel grid of  $J$ , which coincides with that of  $S$ . Without knowledge of the scene, the probability of the image pair  $(I, J)$  is obtained by integrating over all possible realizations of

S. Assuming that the two acquisitions are independent, we have  $P(I, J|S, T) = P(I|S, T) P(J|S)$ , and thus:

$$P(I, J|T) = \int P(I|S, T)P(J|S)P(S) dS. \quad (6)$$

The transformation  $T$  appears as a parameter of this joint probability function. It is natural to invoke the maximum likelihood principle to formulate registration, as already proposed (Viola and Wells, 1997; Leventon and Grimson, 1998; Bansal et al., 1998; Mort and Srinath, 1988; Costa et al., 1993). This simply states that the most likely transformation between  $I$  and  $J$  is the one that maximizes the joint probability of  $(I, J)$ ,

$$\hat{T} = \arg \max_T P(I, J|T).$$

Unfortunately, the integral in Eq. (6) may be intractable unless we assume that the voxels of the scene are independently distributed, so that  $P(S) = \prod_{y_i} P(s_i)$ . This appears as a minimal way to introduce prior anatomical information. Notice, however, this does not mean that the voxels are identically distributed, so that spatial dependencies may still be incorporated into the model. Once  $S$  is a coupled field, there might not exist an analytical expression of  $P(I, J|T)$ . Instead of the present maximum likelihood approach, a maximum a posteriori (MAP) strategy could be employed. This alternative, using for example a Gibbs-Markov random field, would require an explicit estimation of the scene that would be computationally very expensive in 3D images.

In order to simplify Eq. (6), we will also need the transformation  $T$  to be an injection mapping from  $\Omega_I$  to  $\Omega_J$ , i.e.,  $T$  maps distinct points from  $\Omega_I$  to distinct points from  $\Omega_J$ . Let us denote the subset of matched points,  $\mathcal{A} \equiv T(\Omega_I) = \{y_i \in \Omega_J, \exists x_k \in \Omega_I, T(x_k) = y_i\}$ . Recall that, because  $T$  is assumed to be a mapping from  $\Omega_I$  to  $\Omega_J$  (assumption A2), the matched points  $\mathcal{A}$  are assumed to lie entirely within  $\Omega_J$ . The conditional probability  $P(I|S, T)$  can then be evaluated in the same coordinate frame as  $P(J|S)$  and  $P(S)$ :

$$P(I|S, T) = \prod_{x_k \in \Omega_I} P(i_k | s_k^\downarrow) = \prod_{y_i \in \mathcal{A}} P(i_i^\uparrow | s_i),$$

where we have to be cautious that  $i_i^\uparrow \equiv I((T^{-1}(y_i)))$  represents the intensity of the unique voxel  $x_k$  such that  $T(x_k) = y_i$ : it is defined iff  $y_i \in \mathcal{A}$ . We are now in a position to rewrite the joint probability of  $(I, J)$ . Starting from Eq. (6), we have:

$$\begin{aligned} P(I, J|T) &= \int \prod_{y_i \in \mathcal{A}} P(i_i^\uparrow | s_i) \prod_{y_i \in \Omega_J} P(j_i | s_i) P(s_i) \prod_{y_i \in \Omega_J} ds_i, \\ &= \int \prod_{y_i \in \mathcal{A}} P(i_i^\uparrow | s_i) P(j_i | s_i) P(s_i) ds_i \\ &\quad \times \int \prod_{y_i \notin \mathcal{A}} P(j_i | s_i) P(s_i) ds_i, \end{aligned}$$

in which we have regrouped the points  $y_i$  according to whether they match a point in  $\Omega_I$  or not. Then, applying the Fubini theorem, we can invert the integral operands and the products, so that:

$$\begin{aligned} P(I, J|T) &= \prod_{y_i \in \mathcal{A}} \int P(i_i^\uparrow | s_i) P(j_i | s_i) P(s_i) ds_i \\ &\quad \times \prod_{y_i \notin \mathcal{A}} \int P(j_i | s_i) P(s_i) ds_i, \\ &= \prod_{y_i \in \mathcal{A}} P(i_i^\uparrow, j_i) \times \prod_{y_i \notin \mathcal{A}} P(j_i), \end{aligned}$$

Noting that  $P(i_i^\uparrow, j_i) = P(i_i^\uparrow | j_i) P(j_i)$ , we finally get:

$$\begin{aligned} P(I, J|T) &= \prod_{y_i \in \Omega_I} P(j_i) \times \prod_{y_i \in \mathcal{A}} P(i_i^\uparrow | j_i), \\ &= \underbrace{\prod_{y_i \in \Omega_I} P(j_i)}_{P(J)} \times \underbrace{\prod_{x_k \in \Omega_I} P(i_k | j_k^\downarrow)}_{P(I|J, T)}, \end{aligned} \quad (7)$$

where the last step is only a rewriting of  $P(I|J, T)$  in the coordinate frame of  $\Omega_J$ . In Eq. (7), the left term of the product is the marginal probability of  $J$  and is independent of the transformation  $T$ . Only the right term, the conditional probability of  $I$  knowing  $J$  and  $T$ , will play a role in the maximization with respect to  $T$ . We should note that  $P(I|J, T)$  is of the same factored form as  $P(I|S, T)$ :

$$\begin{aligned} P(I|J, T) &= \prod_{x_k \in \Omega_I} P(i_k | j_k^\downarrow), \\ \text{with } P(i_k | j_k^\downarrow) &= \frac{\int P(i_k | s_k^\downarrow) P(j_k^\downarrow | s_k^\downarrow) P(s_k^\downarrow) ds_k^\downarrow}{\int P(j_k^\downarrow | s_k^\downarrow) P(s_k^\downarrow) ds_k^\downarrow}. \end{aligned} \quad (8)$$

It turns out that the statistical relation between  $I$  and  $J$  is of the same form as that between  $I$  and  $S$ . This result is obtained under the assumption that the scene voxels are mutually independent. Therefore, the image  $J$  can be considered as a scene for  $I$  in the sense that  $P(I|J, T)$  verifies the assumptions (A1) and (A2) stated above. However, it is important to realize that the conditional densities  $P(i_k | s_k^\downarrow)$  and  $P(i_k | j_k^\downarrow)$  may have very different expressions. Because the  $P(i_k | s_k^\downarrow)$  are intended to model acquisition noise, they may generally be chosen as single-mode densities (e.g., Gaussian densities). On the other hand, the  $P(i_k | j_k^\downarrow)$  may have much more complicated forms because they incorporate the noise models corresponding to each image as well as the prior probability on  $S$ .

**B. Estimating the Probability Densities.** Until now, we have worked under the assumption that all the probability densities involved in our model were perfectly known. We address here the question of how to estimate them.

Because these densities stand for anatomical and image acquisition models, they should vary significantly from one data set to

another, due not only to interpatient anatomical variability but also to changes in acquisition settings. For example, a tissue such as white matter may have very different ranges of response in two differently acquired brain magnetic resonance (MR) scans. Very often in practice, we cannot use information from previous data sets to model the relationship between the images we want to register.

Therefore, the conditional densities have to be estimated online, in the same manner as we have to estimate the transformation  $T$ . We notice that estimating the densities would be easier if the images were aligned; on the other hand, the registration process needs density estimates to work. Thus, the trade-off is to alternate from registration to density estimation: given a current estimate of  $T$ , estimate the densities, and given current estimates of the densities, update the transformation by maximizing  $P(I, J|T)$ , hoping for convergence. Notice that because only the conditional densities  $P(i_k|j_k^\downarrow)$  play a role in the maximization with respect to  $T$  (see Eq. 7), we could forget about estimating the marginal densities  $P(j_i)$ .


Within analogous formulations of image registration, several methods have been proposed for the density estimation step. Viola and Wells (1997) use the method of Parzen windows to estimate the joint densities  $P(i_k, j_k^\downarrow)$ , from which it would be straightforward to obtain the conditional densities  $P(i_k|j_k^\downarrow)$ . Other authors (Maes et al., 1997; Studholme et al., 1996) construct the 2D histogram of the images, which may be seen as a particular case of Parzen windowing. In these techniques, the densities are not constrained by any model of anatomy or image acquisition; this is perhaps both a strength and a weakness. Moreover, we notice that they provide density estimates that are independent of the spatial positions, in the sense that two voxels  $x_k$  and  $x_{k'}$  are assumed to be identically distributed provided that they have the same intensities in both images, i.e.,  $i_k = i_{k'}$  and  $j_k^\downarrow = j_{k'}^\downarrow$ . A way to incorporate explicit spatial dependence was recently suggested by Bansal et al. (1998), who applied the maximum entropy principle to get nonstationary prior probabilities  $P(s_i)$  for the tissue classes, implying nonstationary densities  $P(i_k|j_k^\downarrow)$ .

In the framework where the transformation is found by maximum likelihood, the most natural way to estimate densities is also to use a maximum likelihood strategy. This means that we can search for the conditional densities  $P(i_k|j_k^\downarrow)$  that maximize exactly the same criterion as in Eq. (8). Basically, this is a parametric approach: we assume that the  $P(i_k|j_k^\downarrow)$  belong to a given class of distributions parameterized by a vector  $\theta$  (regardless, for the moment, of what  $\theta$  represents); then their maximum likelihood estimates, for a given estimate of the transformation  $T$ , are found by:

$$\hat{\theta}(T) = \arg \max_{\theta} P(I|J, T, \theta) = \arg \max_{\theta} \prod_{x_k \in \Omega_I} P_{\theta}(i_k|j_k^\downarrow).$$

The parametric form of  $P(I|J, T, \theta)$  may be derived from the modeling assumptions presented in Section IIA whenever all the components of the model,  $P(i_k|s_k^\downarrow)$ ,  $P(j_i|s_i)$ , and  $P(s_i)$ , are themselves chosen as parametric densities. Then, from Eq. (8), the form of  $P(I|J, T, \theta)$  can be known. We show in Section III that under some specific modeling assumptions, maximum likelihood density estimates can be computed explicitly.

**C. Registration Energy.** By substituting the estimated densities in Eq. (8), our registration criterion becomes the maximum of  $P(I|J, T, \theta)$  under  $\theta$  at fixed  $T$ . Actually, this is only a special way to maximize  $P(I|J, T, \theta)$  with respect to  $(T, \theta)$ . There is no formal



**Figure 1.** Effects of applying a continuous spatial transformation.

difference between the parameters  $T$  and  $\theta$ , except that the latter models the relation between the image intensities. In the context of registration,  $T$  is the parameter in which we are really interested. For practical optimization, it is often more convenient to consider the negative log-likelihood (to be minimized); thus, we will define the energy of a transformation  $T$  as:

$$U(T) = -\log \max_{\theta} P(I|J, T, \theta) = \min_{\theta} \left[ -\sum_{x_k \in \Omega_I} \log P_{\theta}(i_k|j_k^\downarrow) \right]. \quad (9)$$

**D. Practical Issues.** In Section IIA, we derived the likelihood registration criterion under the assumption that the transformation  $T$  is searched for among mappings from the floating image grid,  $\Omega_f$ , to the reference image grid,  $\Omega_r$ . In other words, a grid point of  $I$  was supposed to always match a grid point of  $J$ . The spatial resolution of the transformation is thus intrinsically limited by the resolution of the reference grid. Clearly, this assumption cannot deal with sub-voxel accurate registration.

In practice, we generally want to take into account continuous spatial transformations, not only for a question of accuracy but also because the motion to be estimated is continuous in nature. Thus, we would like the resolution of the reference grid to be as small as the computer working precision. This is achieved in practice by oversampling the image  $J$  using fast interpolation techniques such as trilinear or partial volume interpolation (Maes et al., 1997; Sarrut and Miguette, 1999; Sarut and Feschet, 1999; Plum et al., 1999). Notice that for evaluating the registration criterion Eq. (8), we do not actually have to interpolate every point in space, but only, for a given transformation, the points that are put into correspondence with voxels of  $I$ , i.e., the subset  $\mathcal{A}$  defined in Section IIA.

However, interpolation is possible only if the transformed position of a voxel falls inside the reference domain. Because this domain has a finite extension in space, other voxels may fall outside, so that there is not enough information to interpolate the intensity of their correspondent (Fig. 1). The problem of how to treat these outside voxels plays an important role in voxel-based image registration. They are generally ignored by the registration criterion, which necessitates some heuristic normalization to avoid nasty effects such as disconnecting the images (Studholme et al., 1998; Viola and Wells, 1997; Roche et al., 1998b).

Here, to keep consistent with the maximum likelihood framework, we definitely cannot ignore them. Doing so, we would no longer maximize the image likelihood,  $P(I|J, T)$ , but the likelihood

of a part of  $I$ , which is variable according to the considered transformation. There is always a risk to isolate small image regions that seem very likely to be aligned (typically, in the background). The algorithm might then converge to an aberrant solution.

To tackle this problem, a natural approach is simply to extend the reference domain by assigning the external points to an arbitrary intensity class  $J = j^*$  and defining a specific conditional density corresponding to this class. Although this sounds like a computational artifice, this enables us to take into account every voxel of  $I$  at each iteration of the registration process. Each voxel with intensity  $i_k$  falling outside the reference domain will have a nonzero contribution  $-\log P(i_k|j^*)$  to the registration energy. Thus, we can expect the registration energy to vary little with image overlap as is the effect achieved with classical normalization.

### III. FROM MODELING ASSUMPTIONS TO SIMILARITY MEASURES

The purpose of this section is to demonstrate the link between the general maximum likelihood approach that has been presented and two popular registration criteria: the correlation ratio and mutual information. We will show these measures can be derived from the above formalism using specific modeling assumptions.

**A. Gaussian Channel.** Perhaps the simplest model we can imagine is that the image  $J$  be a valid scene ( $J = S$ ) and the image  $I$  be a measure of  $J$  corrupted with additive and stationary Gaussian white noise:

$$I(x_k) = f(S(T(x_k))) + \epsilon_k,$$

where  $f$  is some unknown intensity function: each tissue class  $j$  is imaged in  $I$  with an average response value  $f(j) = f_j$ . Then, the conditional densities  $P(i_k|j_k^\downarrow)$  have the Gaussian form:

$$P(i_k = i|j_k^\downarrow = j) = \frac{1}{\sqrt{2\pi\sigma}} e^{-(i-f_j)^2/2\sigma^2},$$

and the parameter vector  $\theta = (f_0, f_1, \dots, \sigma)$  needs to be estimated. In order to minimize the negative log-likelihood (Eq. 9) with respect to  $\theta$ , we group the voxels  $x_k$  that match the same class. Letting  $N = \text{Card}\Omega_I$ ,  $\Omega_I^j = \{x_k \in \Omega_I, j_k^\downarrow = j\}$ , and  $N_j = \text{Card}\Omega_I^j$ , we have:

$$\begin{aligned} -\log P(I|J, T, \theta) &= N \log \sqrt{2\pi\sigma} + \frac{1}{2} \sum_{x_k \in \Omega_I} \frac{(i_k - f(j_k^\downarrow))^2}{\sigma^2} \\ &= N \log \sqrt{2\pi\sigma} + \frac{1}{2} \sum_j \sum_{x_k \in \Omega_I^j} \frac{(i_k - f_j)^2}{\sigma^2}. \end{aligned} \quad (10)$$

The optimal parameters are then easily found by differentiating the log-likelihood:

$$\begin{aligned} -\frac{\partial \log P}{\partial f_j} &= -\frac{1}{\sigma^2} \sum_{\Omega_I^j} (i_k - f_j) \Rightarrow \hat{f}_j = \frac{1}{N_j} \sum_{\Omega_I^j} i_k, \\ -\frac{\partial \log P}{\partial \sigma} &= \frac{N}{\sigma} - \frac{1}{\sigma^3} \sum_j \sum_{\Omega_I^j} (i_k - f_j)^2 \Rightarrow \hat{\sigma}^2 = \sum_j \frac{N_j}{N} \hat{\sigma}_j^2, \end{aligned}$$

where  $\hat{\sigma}_j^2 \equiv 1/N_j \sum_{x_k \in \Omega_I^j} (i_k - \hat{f}_j)^2$  is the image variance corresponding to the isoset  $\Omega_I^j$ . The registration energy  $U(T)$  is then obtained by substituting the optimal  $\theta$  parameter:

$$U(T) = \frac{N}{2} \log \left[ 2\pi e \sum_j \frac{N_j}{N} \hat{\sigma}_j^2 \right] = \frac{N}{2} \log [2\pi e \text{Var}(I - \hat{f}(J^\downarrow))].$$

This result has a satisfying interpretation:  $U(T)$  decreases with the variance of the difference image between  $I$  and the intensity corrected  $\hat{f}(J)$ . The intensity function  $\hat{f}$  is nothing but a least squares fit of the image  $I$  in terms of the reference  $J$ : it is in fact the same fitting function as in the definition of the correlation ratio (Eq. 2) (Roche et al., 1998a,b), and we see that the registration energy  $U(T)$  is related to the correlation ratio  $\eta^2(I|J^\downarrow)$  by:

$$\eta^2(I|J^\downarrow) = 1 - \frac{1}{k} e^{2U(T)/N}, \quad \text{with } k = 2\pi e \text{Var}(I).$$


In the original version of the correlation ratio (Roche et al., 1998b), the quantities  $N$  and  $\text{Var}(I)$  were computed only in the overlap between the images, and thus, they could vary according to the considered transformation. Their role was precisely to prevent the image overlap from being minimized. In the implementation proposed in Section IID,  $N$  and  $\text{Var}(I)$  are independent of the considered transformation. Minimizing  $U(T)$  is then strictly equivalent to maximizing the correlation ratio, although it is not strictly equivalent to maximizing the original version of the correlation ratio. In our experiments, this distinction seemed to have very little impact on the results. Still there are reasons to believe that differences may be observed in cases where the image overlap is susceptible to rapidly changing with pose. In our experiments, this was not the case. However, this question needs to be addressed with further testing.

Another remark is that, in practice, we may compute the correlation ratio using a reference image that is not a valid segmentation: then there are as many tissue classes as image isointensity sets, typically 256 for a 1-byte encoded image. For 2- or 4-byte images, this approach may be meaningless and we should impose constraints to the intensity function  $f$ . In the appendix, we generalize the notion of correlation ratio to polynomial imaging functions. Notably, if we constraint  $f$  to follow an affine variation with respect to  $j$ , i.e.,  $f(j) = \alpha j + \beta$ , we get a similar equivalence with the correlation coefficient defined in Eq. (1):

$$\rho^2(I, J^\downarrow) = 1 - \frac{1}{k} e^{2U(T)/N}, \quad \text{with } k = 2\pi e \text{Var}(I).$$

**B. Unspecified Channel.** A straightforward extension of the previous model would be to assume the reference image  $J$  to be also corrupted with Gaussian noise. Then, having defined the prior probabilities for the tissue classes, we could derive the analytical form of the conditional densities  $P(i_k|j_k^\downarrow)$  from Eq. (8). This case has been investigated by Leventon and Grimson (1998). It turns out that there is probably nothing much faster than an Expectation-Maximization (EM) algorithm to provide maximum likelihood estimates of the density parameters.

In order to get explicit density estimates, we can relax every formal constraint on the model. Then, the densities  $P(i_k|j_k^\downarrow)$  are totally unspecified, and we will only assume that they are stationary,



**Figure 2.** Multimodal registration by maximization of CR. Images from left to right: MR-T1, MR-T2, CT, and PET. The images are resampled in the same reference frame after registration. Contours extracted from the MR-T1 are superimposed on every other modality in order to better visualize the quality of registration.

i.e.,  $P(i_k = i | j_k^\downarrow = j)$  is independent of the position  $x_k$ . For the sake of simplicity, we consider the case of discrete densities, but the study is similar for continuous densities. The problem is now to minimize

$$-\log P(I|J, T, \theta) = \sum_{x_i \in \Omega_i} -\log f(i_k | j_k^\downarrow),$$

with respect to  $\theta = (f(0|0), f(1|0), \dots, f(1|1), \dots, f(2|0), \dots)$  and under the constraints:  $\forall_j, C_j = \sum_i f(i|j) - 1 = 0$ . We regroup the intensity pairs  $(i_k, j_k^\downarrow)$  that have the same values:

$$\Omega_{i,j} = \{x_k \in \Omega_i, I(x_k) = i, J(T(x_k)) = j\}, \quad N_{i,j} = \text{Card}\Omega_{i,j},$$

Then, the negative log-likelihood becomes:

$$-\log P(I|J, T, \theta) = -\sum_{i,j} N_{i,j} \log f(i|j).$$

Introducing Lagrange multipliers, there exists constants  $\lambda_0, \lambda_1, \dots$ , such that for any  $j$ :

$$0 = \frac{\partial \log P}{\partial f(i|j)} - \sum_j \lambda_j \frac{\partial C_j}{\partial f(i|j)} = \frac{N_{i,j}}{f(i|j)} - \lambda_j.$$

Thus, taking into account the constraints  $\sum_i f(i|j) = 1$ , the optimal parameters verify:

$$\hat{f}(i|j) = \frac{N_{i,j}}{N_j} = \frac{p(i,j)}{p(j)},$$

where  $p(i,j) \equiv N_{i,j}/N$  is the image-normalized 2D histogram and  $p(j) \equiv \sum_i p(i,j)$  the corresponding marginal distribution for  $J^\downarrow$ . Equation (3) shows that  $U(T)$  is nothing but a decreasing function of mutual information:

$$U(T) = -N \sum_{i,j} p(i,j) \log \frac{p(i,j)}{p(j)} = N[H(I) - \mathcal{F}(I, J^\downarrow)],$$

where  $H(I)$  is the entropy of image  $I$  and is constant in the implementation proposed in Section IID. The same remark as made in Section IIIA holds for the distinction between the usual implementation of mutual information and the one considered here.

**C. Comparison of Measures.** In the derivation of the correlation ratio (CR), it was assumed that the image to be registered is a measure of the reference corrupted with additive and stationary Gaussian white noise. In contrast, for deriving mutual information (MI), no assumption was made apart from stationarity and, of course, the assumptions (A1) and (A2) stated in Section II. Does it make MI necessarily a better registration measure than CR?

In principle, the answer is no whenever the assumptions of CR are verified by the images. Basically, these are reasonable if the reference image can be considered as a good anatomical model: in practice, this is often a rough approximation. The problem then is to determine what is better between an overconstrained and an underconstrained measure, a question to which experiments can yield some insight, as will be illustrated in the next section.

#### IV. EXPERIMENTS OF RIGID REGISTRATION

This section illustrates the practical differences between the CR and MI measures in the context of 3D rigid registration of brain images acquired from different modalities. Following the ideas proposed by Maes et al. (1997), we implemented Powell's method (Press et al., 1992) to optimize the measures with respect to the transformation parameters. Partial volume interpolation (PV) was used in all the experiments.

**A. Vanderbilt Database.** The registration algorithm was tested using image data sets from 10 patients. For each patient, the following images were available:

- MR, T1 weighted (256 × 256 × 20/26 voxels of 1.25 × 1.25 × 4 mm<sup>3</sup>)
- MR, T2 weighted (256 × 256 × 20/26 voxels of 1.25 × 1.25 × 4 mm<sup>3</sup>)
- Computed tomography (CT; 512 × 512 × 28/34 voxels of 0.65 × 0.65 × 4 mm<sup>3</sup>)
- Positron emission tomography (PET; 128 × 128 × 15 voxels of 2.59 × 2.59 × 8 mm<sup>3</sup>)

The gold standard transformations between the modalities were known thanks to a prospective, marker-based registration method (West et al., 1997). We performed three kinds of registrations: T1 to T2, CT to T1, and PET to T1 (Fig. 2). In all the experiments, the transformation was initially set either to the identity or to the gold standard: this was done to test if the algorithm was sensitive to initialization. However, because the results were almost the same for

**Table I.** Rigid registration errors obtained over 10 inpatient experiments.

Experiment	Measure	RMS		
		$\Delta\theta$ (deg)	$\Delta t$ (mm)	$\epsilon$ (mm)
T1/T2	CR	0.31	2.28	1.91
	MI	0.58	2.19	2.16
CT/T1	CR	2.91	11.27	6.75
	MI	0.77	3.98	3.31
PET/T1	CR	1.53	5.49	5.16
	MI	1.42	7.55	7.84

both types of initialization, we present only those obtained when starting from the identity.


After each registration, a typical error  $\epsilon$  was computed by taking the average registration error of the eight vertices of a bounding box corresponding approximately to the head volume. Thus,  $\epsilon$  represents the error to be expected in the region of interest. We also computed the intrinsic rotation and translation errors (Pennec and Thirion, 1997). Let  $R_*$  and  $t_*$  be the ground truth rotation matrix and translation vector. For a rotation matrix  $R$  and a translation vector  $t$  found by the registration algorithm, the intrinsic rotation error  $\Delta\theta$  is the norm of the rotation vector corresponding to the residual rotation

matrix  $R_*'R$  and the translation error is the Euclidean distance between  $t_*$  and  $t$ , i.e.,  $\Delta t = \|t - t_*\|$ . Intrinsic errors were chosen because these are objective measures that are independent of any region of interest.

Table I shows Root Mean Squares (RMS) of  $\epsilon$ ,  $\Delta\theta$ , and  $\Delta t$  for the 10 patients, for each modality combination. These have to be compared to the image resolution, which is quite poor (4 mm in the  $z$ -axis for the MR and CT data sets and 8 mm for the PET). The reader may notice that errors reported here are higher than target registration errors reported by other groups in the retrospective registration evaluation project (Woods et al., 1993). Our guess is that this is simply because our respective methods for computing errors are different.

In T1/T2 registration, CR and MI give good and similar results. For the other combinations, more significant differences are observed. MI does a much better job for matching CT to T1. This might be due to the fact that functional dependence is a crude hypothesis in the CT/MR case (Wells et al., 1997). On the other hand, CR tends to give slightly better results for PET/T1 registration.

**B. US/MR.** A very challenging registration problem consists of aligning an intraoperative ultrasound (US) image with a preoperative image such as an MR scan. We tested the registration algorithm



**Figure 3.** (Top left) Three orthogonal views of the MR image. (Top right) Corresponding views of the US image in a random initial position. (Bottom left) Display of the initial US with contours from the MR superimposed. (Bottom right) Same display with the registered US.

**Table II.** RMS errors and percentages of failures in 3D US-MR rigid registration.

Reference Image	Similarity Measure	RMS		RMS (successes)		Failures (%)
		$\Delta\theta$ (deg)	$\Delta t$ (mm)	$\Delta\theta$ (deg)	$\Delta t$ (mm)	
Original MR	CR	11.49	23.33	1.11	0.42	14.0
	MI	19.07	47.14	1.27	0.64	51.0
Filtered MR (anisotropic diffusion)	CR	12.64	26.29	0.92	0.52	12.5
	MI	17.35	27.41	1.35	0.82	28.0
Distorted MR ( $\sigma = 10\%$ )	CR	28.51	18.08	3.21	2.04	36.0
	MI	44.23	45.06	1.84	1.36	90.0

with an MR, T1 weighted scan ( $256 \times 256 \times 124$  voxels of  $0.9 \times 0.9 \times 1.1$  mm) and an intraoperative 3D US image ( $180 \times 136 \times 188$  voxels of  $0.95^3$  mm<sup>3</sup>).

Because the US image was acquired before opening the dura-mater, we neglect the brain shift phenomenon. Thus, there is essentially a rigid displacement to find. The correct registration position was found manually using an interactive matching tool, and validated by a clinician. The estimated accuracy was  $2^\circ$  in rotation and 2 mm in translation. We took this first result as a ground truth for subsequent experiments.

We then performed 200 automatic registrations by initializing the algorithm with random displacements from the ground truth position (Fig. 3): a rotation vector  $\Delta r$  with random direction and constant magnitude  $\|\Delta r\| = 15^\circ$  and a translation  $\Delta t$  with random direction and constant magnitude  $\|\Delta t\| = 20$  mm. These values correspond to the variation between the ground truth and the original position. For each random transformation, two registrations were performed using, alternatively, CR and MI. To avoid interpolation artifacts due to resampling, the algorithm did not take as an input the US resampled by the ground truth transformation, but always the original US itself. Another advantage of doing so is that the ground truth corresponds to a partial overlap between the MR volume and the original US volume. Therefore, there is no reason to expect the registration results to be biased toward the ground truth due to the problems associated with changing the image overlap (see Section IID).

We observe two kinds of results: either the algorithm retrieves the ground truth transformation (yielding errors systematically lower than  $\|\delta r\| = 2^\circ$  and  $\|\delta t\| = 2$  mm) or it converges to a local maximum (yielding errors systematically larger than  $\|\delta r\| = 10^\circ$  and  $\|\delta t\| = 10$  mm). The main result is that CR fails in 14% of cases whereas MI fails in 51% of cases (Table II). The RMS errors computed on successful registrations are lower than the expected accuracy of the ground truth; thus, they prove nothing but the fact that both CR and MI have a maximum in the neighborhood of the ideal registration transformation (this is probably also a global maximum). However, the percentages of success indicate that CR may have a wider attraction basin, an observation consistent with previous experiments with other modality combinations (Roche et al., 1998b).

To study the effect of noise in the data, we repeated the same experimental protocol twice, using as a reference image the MR presegmented by anisotropic diffusion (Perona and Malik, 1990) and the MR corrupted with Gaussian noise. The number of failures for both measures is clearly affected by the amount of noise (Table II). This comes as no surprise in the case of CR, because this measure has been derived under the assumption that there is no noise in the reference image (see Section IIIA). This is more surprising for MI, as no such assumption was made.

We conclude that the attraction basin of the measures could be extended by denoising the MR image in a preprocessing step.

Studying the effects on accuracy would have been of great interest too, but this was not possible here because the ground truth could not be considered accurate enough.

We believe that the registration algorithm would greatly benefit from reducing noise in the US. Unfortunately, applying classical anisotropic diffusion to the US did not provide convincing results. It tended to blur the image and did not remove speckle artifacts. Specific filtering tools need to be developed for US images. This is still an open research track.

## V. CONCLUSION

We have formalized image registration as a general maximum likelihood estimation problem and shown that several existing similarity measures may be reinterpreted in this framework. This enables us to better understand the implicit assumptions we make when using a particular measure, and hopefully, helps the selection of an appropriate strategy given a certain problem.

Experimental results of rigid registration confirm (if needed) that similarity measures relating to different assumptions yield different performances. The CR measure was shown to be more efficient than MI in the case of PET/MR and US/MR registration. As CR relies on more restrictive hypotheses than MI, this suggests the importance of constraining the relationship between the images. On the other hand, the assumptions should also be founded. We are aware that CR relies on a model that, although simpler, may not be realistic.

Because the presented work allows us to derive systematically the similarity measures from explicit modeling assumptions, this is a step toward taking into account more realistic models of image acquisition and anatomy. In the future, we plan to develop this approach for the challenging problem of US/MR registration.

## ACKNOWLEDGMENTS

The authors thank Sébastien Ourselin, Alexandre Guimond, and Sylvain Prima for constant interaction and countless suggestions. Also many thanks to Janet Bertot for the proofreading of this article. The images and the standard transformations used in Section IVA were provided as part of the project, "Evaluation of Retrospective Image Registration," National Institutes of Health, Project Number 1 R01 NS33926-01, Principal Investigator, J. Michael Fitzpatrick, Vanderbilt University, Nashville, TN. The images used in Section IVB were provided by ISM-Austria, Salzburg, Austria, for the US data sets, and the Max Planck Institute for Psychiatry, AG-NMR, Munich, Germany, for the MR datasets, as part of the EC-funded ROBOSCOPE project HC 4018, a collaboration among The Fraunhofer Institute (Germany), Fokker Control System (Netherlands), Imperial College (United Kingdom), INRIA Sophia Antipolis (France), and ISM-Salzburg and Kretz Technik (Austria). Part of this work was supported by la Région Provence-Alpes-Côtes d'Azur (France).



## APPENDIX

**A. Generalization of the CR.** For the problem considered in Section IIIA, we could define the unknown imaging function  $f$  to be polynomial with degree  $d$ :

$$f(j) = \sum_{p=0}^d \alpha_p j^p.$$

Then, we aim at minimizing the log-likelihood,

$$\begin{aligned} -\log P(I|J, T, \theta) &= N \log \sqrt{2\pi\sigma} + \frac{1}{2} \sum_{x_k \in \Omega_I} \frac{(i_k - f(j_k^\downarrow))^2}{\sigma^2} \\ &= N \log \sqrt{2\pi} \sigma + \frac{1}{2\sigma^2} \sum_{x_k \in \Omega_I} [i_k - \sum_{p=0}^d \alpha_p j_k^{\downarrow p}]^2, \end{aligned} \quad (11)$$

with respect to  $\theta = (\alpha_0, \alpha_1, \dots, \alpha_p, \sigma)$ . By differentiating Eq. (11), we get:

$$\begin{aligned} -\frac{\partial \log P}{\partial \alpha_q} &= -\frac{1}{\sigma^2} \sum_{x_k \in \Omega_I'} j_k^{\downarrow q} [i_k - \sum_{p=0}^d \alpha_p j_k^{\downarrow p}], \\ &= -\frac{1}{\sigma^2} \left[ \sum_{x_k \in \Omega_I'} j_k^{\downarrow q} i_k - \sum_{x_k \in \Omega_I'} \sum_{p=0}^d \alpha_p j_k^{\downarrow p+q} \right], \\ &= -\frac{1}{\sigma^2} \left[ \sum_{x_k \in \Omega_I'} j_k^{\downarrow q} i_k - \sum_{p=0}^d \alpha_p \left( \sum_{x_k \in \Omega_I'} j_k^{\downarrow p+q} \right) \right]. \end{aligned}$$

The polynomial coefficients are then seen to be solutions of the linear system  $AX = B$ , with

$$\begin{aligned} X &= (\alpha_0 \quad \alpha_1 \quad \alpha_2 \quad \dots \quad \alpha_d)^T \\ A &= \begin{pmatrix} N & \sum j_k^\downarrow & \sum j_k^{\downarrow 2} & \dots & \sum j_k^{\downarrow d} \\ \sum j_k^\downarrow & \sum j_k^{\downarrow 2} & \sum j_k^{\downarrow 3} & \dots & \sum j_k^{\downarrow d+1} \\ \sum j_k^{\downarrow 2} & \sum j_k^{\downarrow 3} & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \sum j_k^{\downarrow 2d-1} \\ \sum j_k^{\downarrow d} & \sum j_k^{\downarrow d+1} & \dots & \dots & \sum j_k^{\downarrow 2d} \end{pmatrix} \\ B &= \left( \sum i_k \quad \sum j_k^\downarrow i_k \quad \sum j_k^{\downarrow 2} i_k \quad \dots \quad \sum j_k^{\downarrow d} i_k \right)^T \end{aligned}$$

In practice, we invert the  $(d+1) \times (d+1)$  matrix  $A$  by the method of singular value decomposition (SVD). This avoids numerical explosion when  $A$  comes close to singularity. To solve for the standard deviation  $\sigma$ , we differentiate the log-likelihood:

$$-\frac{\partial \log P}{\partial \sigma} = \frac{N}{\sigma} - \frac{1}{\sigma^3} \sum_{x_k \in \Omega_I} (i_k - f(j_k^\downarrow))^2 \Rightarrow \hat{\sigma}^2$$

$$= \frac{1}{N} \sum_{x_k \in \Omega_I} (i_k - \hat{f}(j_k^\downarrow))^2.$$

Thus, the optimal  $\sigma$  is equal to the standard deviation of the difference image between  $I$  and the optimally corrected image  $J^\downarrow$ . This yields the registration energy,

$$U(T) = \frac{N}{2} \log[2\pi e \text{Var}(I - \hat{f}(J^\downarrow))],$$

a measure that directly generalizes the CR.

## REFERENCES

- R. Bansal, L.H. Staib, Z. Chen, A. Rangarajan, J. Knisely, R. Nath, and J.S. Duncan, A novel approach for the registration of 2D portal and 3D CT images for treatment setup verification in radiotherapy, Proc MICCAI '98, 1998, pp. 1075–1086.
- M. Bro-Nielsen, Rigid registration of CT, MR and cryosection images using a GLCM framework, Proc CVRMed-MRCAS '97, 1997, pp. 171–180.
- L.G. Brown, A survey of image registration techniques, ACM Comput Surv 24 (1992), 325–376.
- T.M. Buzug and J. Weese, Voxel-based similarity measures for medical image registration in radiological diagnosis and image guided surgery, J Comput Inform Technol 1998, 165–179.
- W.L.S. Costa, D.R. Haynor, R.M. Haralick, T.K. Lwells, and M.M. Graham, A maximum-likelihood approach to PET emission/attenuation image registration, IEEE Nuclear Science Symp and Medical Imaging Conf, 1993.
- S. Lavallee, “Registration for computer integrated surgery: Methodology, state of the art,” Computer integrated surgery, R. Taylor, S. Lavallee, G. Burdea, and R. Moesges (Editors), MIT Press, Cambridge, MA, 1995, pp. 77–97.
- M.E. Leventon and W.E.L. Grimson, Multi-modal volume registration using joint intensity distributions, Proc MICCAI '98, 1998, pp. 1057–1066.
- F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, Multimodality image registration by maximization of mutual information, IEEE Trans Med Imaging 16 (1997), 187–198.
- J.B.A. Maintz and M.A. Viergever, A survey of medical image registration, Medical Image Analysis 2 (1998), 1–36.
- M.S. Mort and M.D. Srinath, Maximum likelihood image registration with subpixel accuracy, Proc SPIE, Vol. 974, 1988, pp. 38–45.
- C. Nikou, F. Heitz, J.-P. Armspach, and I.-J. Namer, Single and multimodal subvoxel registration of dissimilar medical images using robust similarity measures, SPIE Conf on Medical Imaging, Vol. 3338, 1998, pp. 167–178.
- A. Papoulis, Probability, random variables, and stochastic processes (3rd ed.), McGraw-Hill, New York, 1991.
- X. Pennec and J.P. Thirion, A framework for uncertainty and validation of 3D registration methods based on points and frames, Int J Computer Vision 25 (1997), 203–229.
- 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, Proc MICCAI '98, 1998, pp. 1153–1161.
- P. Perona and J. Malik, Scale-space and edge detection using anisotropic diffusion, IEEE Trans Pattern Analysis Machine Intell 12 (1990), 629–639.
- J.P.W. Pluim, J.B.A. Maintz, and M.A. Viergever, Mutual information matching and interpolation artefacts, Proc SPIE 3661 (1999), 56–65.
- W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling, Numerical recipes in C (2nd ed.), University Press, Cambridge, England, 1992.

- A. Roche, G. Malandain, N. Ayache, and X. Pennec, Multimodal image registration by maximization of the correlation ratio, Technical Report 3378, INRIA, August 1998a.
- A. Roche, G. Malandain, X. Pennec, and N. Ayache, The correlation ratio as a new similarity measure for multimodal image registration, Proc MICCAI '98, Vol. 1496 of Lecture Notes in Computer Science, 1998b, pp. 1115–1124.
- D. Sarrut and F. Feschet, "The partial intensity difference interpolation," International conference on imaging science, systems and technology, H.R. Arabnia (Editor), CSREA Press, Las Vegas, USA, pp. 46–51, 1999.
- D. Sarrut and S. Miguet, Fast 3D images transformations for registration procedures, 10th Int Conf on Image Analysis and Processing, 1999, pp. 446–452.
- C. Studholme, D.L.G. Hill, and D.J. Hawkes, Automated 3-D registration of MR and CT images of the head, *Med Image Analysis* 1 (1996), 163–175.
- C. Studholme, D.L.G. Hill, and D.J. Hawkes, An overlap invariant entropy measure of 3D medical image alignment, *Pattern Recogn* 1 (1998), 71–86.
- A.W. Toga, *Brain warping*, Academic Press, San Diego, 1999.
- P.A. van den Elsen, E.J.D. Pol, and M.A. Viergever, Medical image matching: A review with classification, *IEEE Eng Med Biol* 12 (1993), 26–39.
- P. Viola and W.M. Wells, Alignment by maximization of mutual information, *Int J Computer Vision* 24 (1997), 137–154.
- W.M. Wells, P. Viola, H. Atsumi, and S. Nakajima, Multi-modal volume registration by maximization of mutual information, *Med Image Analysis* 1 (1996), 35–51.
- J. West et al., Comparison and evaluation of retrospective intermodality brain image registration techniques, *J Comp Assist Tomogr* 21 (1997), 554–566.
- R.P. Woods, J.C. Mazziotta, and S.R. Cherry, MRI-PET registration with automated algorithm, *J Comp Assist Tomogr* 17 (1993), 536–546.