Generalized Image Models and Their Application as Statistical Models of Images

Miguel Ángel González Ballester *, Xavier Pennec, Marius George Linguraru and Nicholas Ayache

INRIA, Epidaure Project, 2004 route des lucioles - B.P. 93, FR-06902 Sophia Antipolis, France

Abstract

A Generalized Image Model (GIM) is presented. Images are represented as sets of 4-dimensional sites combining position and intensity information, as well as their associated uncertainty and joint variation. This model seamlessly allows for the representation of both images and statistical models (such as those used for classification of normal/abnormal anatomy and for inter-patient registration), as well as other representations such as landmarks or meshes. A GIM-based registration method aimed at the construction and application of statistical models of images is proposed. A procedure based on the Iterative Closest Point (ICP) algorithm is modified to deal with features other than position and to integrate statistical information. Furthermore, we modify the ICP framework by using a Kalman filter to efficiently compute the transformation. The initialization and update of the statistical model are also described. Preliminary results show the feasibility of the approach and its potentialities.

Key words: Generalized image models, statistical shape models, registration, ICP, Kalman filter.

* Corresponding author. Phone: +33-492387155, Fax: +33-492387669.

Email address: Miguel.Gonzalez@sophia.inria.fr (Miguel Ángel González Ballester).

URL: http://www-sop.inria.fr/epidaure/personnel/Miguel.Gonzalez/ (Miguel Ángel González Ballester).

1 Introduction

A common taxonomy of image analysis techniques consists in establishing two main groups of algorithms: those dealing with images (voxel-based, iconic, intensity-based) and those incorporating some form of geometric information (landmark-based, active surface methods, etc.). In order to bridge the gap between these two classes, this paper introduces a Generalized Image Model (GIM) aimed at unifying the representation of images, landmarks, point sets, and other shape representations.

Central to the idea of the GIM is the explicit modeling of the uncertainty inherent to the available data. Such uncertainty derives not only from noise and limitations on device resolution; we also consider the explicit representation of prior knowledge about the anatomical variability across patients and/or acceptable ranges for data values. Thus, both the issue of error propagation (e.g. [1,2]) and that of statistical shape modeling [3,4] are considered.

In this paper, we introduce the GIM and discuss its benefits. We describe how it can be used to seamlessly model images and statistical models of image data in the same representation. In order to make them interact, i.e. exploit the statistical model to aid in the analysis of new images and, conversely, update the model with the contents of such new images, an image/model registration method based on the iterative closest point (ICP) algorithm [5] is presented. We follow the principle of the 4D-ICP algorithm [6] in regards to taking into account intensity as well as position. Most importantly, the method has been extended to deal with the statistical nature of the data contained in the GIM. Thus, it applies without modification to the registration of either two images, an image and a statistical model, or two models. Optimization of the registration criterion is improved by modifying the classical ICP framework by incorporating an efficient Kalman filter.

Finally, the initialization and update of the statistical model are dealt with by computing simple statistical measures on the Voronoi tessellation defined by the elements of the GIM. This provides a way to dynamically refine the statistical model to incorporate new information.

A description of the GIM and its use to represent in the same framework images and statistical models of image data are given in section 2. The ICP-based image/model registration algorithm follows in section 3, and the use of Kalman filters to efficiently find the solution is described in section 4. Finally, the initialization and update of the statistical model are given in section 5. Section 6 shows some illustrative results, and discussion and conclusions are provided in section 7.

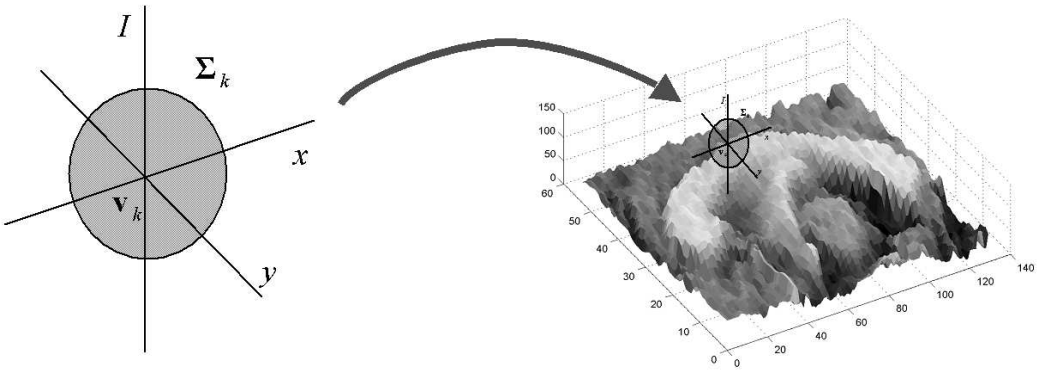

Fig. 1. Illustration of the Generalized Image Model. Sites contain position and intensity information; a covariance matrix is associated to each site representing variability and correlation of the position and intensity components. The surface on the right corresponds to a 2D image of the corpus callosum in MRI, where the elevation of the surface represents image intensity.

2 Generalized Image Model (GIM)

The essence of the GIM is the explicit representation of position along with intensity information, and the inclusion of a covariance term to model uncertainty and correlation. In particular, an image A is modeled as a set of N_A sites v_k, each one with an associated covariance matrix Σ_k:

$$A = \{(v_k, \Sigma_k)\}_{k=1..N_A} \tag{1}$$

Each site contains position and intensity information. In the general case, the position component of a site v_k is a D_p-dimensional (typically $D_p = 3$) vector p_k. Similarly, the intensity component can be generalized as a vector $i \in \Re^{D_i}$, with $D_i = 1$ in the case of scalar-valued images. Thus:

$$v_k = (p_k, i_k) \tag{2}$$

The covariance matrix Σ_k, of dimensions $(D_p + D_i) \times (D_p + D_i)$, models the variance in position and intensity, and their possible correlation, at each site.

The choice of this image model is part of an attempt to seamlessly integrate information coming from different shape representations. Thus, an image is a particular case of GIM in which the positions are situated in a regular grid. Additionally, images with different resolutions can be combined in a natural fashion, as the uncertainty components can be used to model voxel sizes. More importantly, this representation can be used to include other models, such as point sets, landmarks or meshes (cf. discussion in section 7).

The use of the GIM to build statistical models of images is described in this paper. The same GIM framework is used to represent both an image $A = \{(v_k, \Sigma_k)\}_{k=1..N_A}$ and a statistical model $M = \{(\bar{v}_{k'}, \bar{\Sigma}_{k'})\}_{k'=1..N_M}$. The key difference between them lies within the nature of the covariance matrices. The covariance matrix for an image models the uncertainty in the localization of the information (which can be related for instance to the discretization steps, i.e. the voxel size) and in the intensity information (which can be related to the noise in the imaging process). On the other hand, the covariance matrix for a model contains information about the variability in position and intensity of the set of samples used for training, i.e. it contains the shape variation model.

The following two sections describe an algorithm for the registration of images represented using the GIM, with a particular emphasis to the registration of an image and a statistical model. Section 5 will deal on the construction of such statistical models.

3 ICP-Based Registration of Images and Statistical Models

3.1 Introduction

In this section we will describe how to compute the affine transformation that puts into correspondence two images represented following the GIM described in section 2. Due to the generality of the GIM, this method applies without modification for the registration of two images, an image and a statistical model, or two statistical models. The registration problem, illustrated for the case of the registration of an image and a statistical model, is as follows. Let A be an image, with sites v_k, $k = 1..N_A$, and M a model, whose sites are $\bar{v}_{k'}$, $k' = 1..N_M$. We aim to find the transformation T and the correspondence function π mapping points in the image to points in the model: $\pi(k) = k'$, so that a certain similarity criterion is met. In the classical Euclidean distance case, the least-squares criterion is:

$$\hat{T} = \arg\min_{T,\pi} \sum_k \left[(Tv_k - \bar{v}_{\pi(k)})^t (Tv_k - \bar{v}_{\pi(k)}) \right] \tag{3}$$

Note that T is also applied to the covariance matrix, i.e. the covariance associated to Tv_k is $T\Sigma_k T^t$ (if T is linear, which is the case when using homogeneous coordinates). Let us keep in mind that the information at each site contains not only location but also intensity (and potentially other features).

Our method is based on the Iterative Closest Point (ICP) algorithm [5]. This method was extended to 4-D as described in [6] in order to jointly match

location and intensity (a technique known as 4D-ICP). We modify this technique in order to handle the statistical information contained in the GIM, thus allowing it to deal with statistical models.

3.2 Criterion to be minimized

The classical Euclidean distance measure employed in eq. 3 must be replaced by a distance measure that takes into account second order statistics (i.e. the covariance matrices in the GIM). One such distance measure is the Mahalanobis distance. We thus minimize the Mahalanobis distance between the current estimate of Tv_k and each point in the model (which defines a Gaussian distribution). Thus:

$$\hat{T} = \arg\min_{T,\pi} \sum_k \left[(Tv_k - \bar{v}_{\pi(k)})^t \bar{\Sigma}^{-1}_{\pi(k)} (Tv_k - \bar{v}_{\pi(k)}) \right] \tag{4}$$

Alternatively, the uncertainty in the location of v_k can also be included in the criterion by measuring the distance between the distributions defined at each location of the model and the image. In the case of the Mahalanobis distance, the criterion is then:

$$\hat{T} = \arg\min_{T,\pi} \sum_k \left[(Tv_k - \bar{v}_{\pi(k)})^t (T\Sigma_k T^t + \bar{\Sigma}_{\pi(k)})^{-1} (Tv_k - \bar{v}_{\pi(k)}) \right] \tag{5}$$

Other distance measures, such as the Bhattacharyya distance and the Kullback-Leibler divergence, are available. Details about these and a discussion of their performance are available in [7,8].

Eventually, additional information can be included at each site. For example, in the 2D case curvature can be used to further constrain the ICP evolution. A particularly interesting case is that of an independent curvature term, which for the case of the Mahalanobis distance gives the following criterion:

$$\hat{T} = \arg\min_{T,\pi} \sum_k \left[(Tv_k - \bar{v}_{\pi(k)})^t \bar{\Sigma}^{-1}_{\pi(k)} (Tv_k - \bar{v}_{\pi(k)}) + \frac{\kappa_{v_k} - \kappa_{\bar{v}_{\pi(k)}}}{\sigma_{\bar{v}_{\pi(k)}}} \right] \tag{6}$$

where κ_v is the curvature at site v and σ_v its standard deviation in the model. Similar criteria can be employed by incorporating information about the normal at each point (cf. results section 6.2).

Given an initial estimate of the transformation T, the classical ICP framework [5] works in an iterative fashion repeating the following three steps:

(1) Apply T to the image points;
(2) Establish correspondences between the transformed image points and the model points by finding the closest model point according to the chosen distance;
(3) Compute a new least squares estimate of T based on these correspondences.

This algorithm can be justified as an alternated optimization of the above criteria over the correspondence function π and the transformation T. Although simple, this method has proven to be very effective in practice. Typically, ICP is used to compute a rigid transformation between point sets. Here, we use it to compute an affine transformation that will act jointly on position and intensity.

A modification of the classical approach for its use with statistical models could be envisaged. Searching for the closest point would no longer rely on Euclidean distance, but it would utilize one of the distance measures listed above. However, the minimization of a term encompassing one of such distances provides no closed form solution for non-stationary covariance matrices and thus one has to rely on iterative gradient methods [2,9]. An alternative, efficient method to minimize the criterion is given in the following section.

4 Solving for the Transformation via Kalman Filtering

4.1 Registration as a data assimilation task

The computation of the transformation in the classical ICP algorithm is based on the full set of point correspondences found for every point in the data set. This is clearly not optimal from a computational point of view, as the incremental updates of the estimate of the transformation at each iteration are released at a very slow rate. Here, we propose a method for updating the transformation every time a correspondence is found, thus speeding up the computation (see [10] for a similar approach applied to tracking).

The computation of the transformation is formalized as a data assimilation task. A Kalman filter [11] is set up to this effect. Kalman filters are classi-

cally employed for the prediction of the internal state $x^{(t)}$ of a discrete-time controlled process governed by the linear stochastic difference equation:

$$x^{(t)} = Ax^{(t-1)} + Bu^{(t)} + w^{(t)} \tag{7}$$

where A relates the state $x^{(t-1)}$ with the current state at time t, B relates the optional control input u to the state x, and w is the process noise, assumed Gaussian with zero mean and covariance matrix Q.

The external evidence of the internal state of the system is the measurement:

$$z^{(t)} = Hx^{(t)} + v^{(t)} \tag{8}$$

where H relates the internal state with the measurement, and v is the measurement noise, assumed Gaussian with zero mean and covariance matrix R.

In our case, the system models the transformation to be estimated, which we assume constant for every point. We therefore use the following system evolution equation:

$$x^{(t)} = x^{(t-1)} + w \tag{9}$$

with the state vector x containing the parameters of the affine transformation T. For the case of a 2D scalar image the transformation parameters are:

$$T = \begin{bmatrix} x_1 & x_2 & x_3 & x_4 \\ x_5 & x_6 & x_7 & x_8 \\ x_9 & x_{10} & x_{11} & x_{12} \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{10}$$

The image data point v_k is the measure. Each data point v_k is treated individually, and an estimate of the transformation $x^{(t)}$ is computed each time a new point is treated. The measurement function is as follows:

$$\hat{v}_k(x; \bar{v}_{k'}) = H(\bar{v}_{k'})x + v \tag{11}$$

where the matrix $H(\bar{v}_{k'})$ is set up such that $H(\bar{v}_{k'})x = T\bar{v}_{k'}$ and v is drawn from a zero mean Gaussian distribution with covariance matrix R, modeling the measurement noise. $\bar{v}_{k'}$ and $\bar{\Sigma}_{k'}$ are treated as external parameters.

The covariance matrix of the process noise Q is taken to be diagonal with very small values for variances. In fact, it could simply be removed, as we

assume that the transformation being estimated remains constant throughout the process; however, we found that initializing the process variance to a very small value helps reaching convergence faster. As for the matrix R, it models the noise in the measurement process. Therefore it seems natural that the covariance matrix Σ_k, which contains information about the uncertainty of the measure, should be used in its place.

4.2 Algorithm

The resulting Kalman filter is a linear one, which is solved efficiently. Furthermore, the size of the matrices to be inverted is very small, as each data point is treated separately. The update equations for the Kalman filter are as follows [11]:

(1) Provide an initial estimate of the transformation parameters x^0. We can either initialize the transformation using a subsample of the data points and computing a least-squares first estimate, or simply initialize it to the identity matrix.
(2) Give an initial value to the Kalman error covariance P^0. This matrix models the certainty of the estimate: it initially has a large variance, so the measurement has more weight than the current estimate, and it decreases as new data are incorporated and the estimate is refined.
(3) Prediction step:
   - Project the state ahead: $\hat{x}^{(t)} = x^{(t-1)}$, where the hat means *prediction.*
   - Project the error covariance ahead: $\hat{P}^{(t)} = P^{(t-1)} + Q$
(4) Measurement update (correction):
   - Compute the Kalman gain: $K^{(t)} = \hat{P}^{(t)} H^t (H \hat{P}^{(t)} H^t + R)^{-1}$
   - Update estimate with measurement z_k: $x^{(t)} = \hat{x}^{(t)} + K^{(t)}(z_k - H\hat{x}^{(t)})$
   - Update error covariance: $P^{(t)} = (I - K^{(t)} H)\hat{P}^{(t)}$
(5) Increment k and go to step 3 until all points have been treated or the error covariance reaches a minimum value.

This method is easily extendible to non-linear systems, such as those arising from the introduction of transformation classes more complex than the affine case. The use of Extended Kalman Filters (EKF) for such non-linear systems is explained, for example, in [11,12].

4.3 Robustness issues

Starting from an initial estimate of the transformation using a few point correspondences, the method refines this estimate for each new point. One could start from the first point with a large covariance matrix on the transformation

but it often leads to unstable results. Each update step is extremely fast, as the size of the matrices to be inverted is very small. Thus, we obtain a good estimate of the transformation in the order of a few iterations. However, the system evolution equation assumes that we are considering an affine transformation that puts into correspondence every point in the two images, up to a certain white Gaussian noise. The whole process can easily be robustified using a M-estimator of the distance, for instance by ignoring matches that have a Mahalanobis distance greater than a given χ^2 threshold, e.g. following the 3σ rule. Since this amounts to saturating the distance at that value in the original criterion, the convergence of the algorithm is still ensured.

On the contrary, systematic deviations from the assumption of a given affine transformation being able to match the two images do affect the results. For example, if the first half of the points being treated corresponds to a given transformation and the second half to a different one, the final result will depend on the values given to the Kalman covariance, which weighs the relative contribution of the current estimate versus the new measure.

To make the method more robust to drifts in certain parts of the image, a sampling strategy that considers points in a random order can be employed. Another interesting alternative to be considered is the possibility of updating the estimate based in correspondence measures on blocks of the image [13] instead of single voxels.

5 Initializing and Updating the Statistical Model

The previous section showed how to perform the registration of a statistical model and an image. Once the registration has been successfully estimated, the information contained in the image can be used to update the model, or in the case of two models being registered, the information contained in both of them can be fused. In this section, we comment on the process of construction of statistical models as represented by the GIM.

The initialization of a model is obtained by registering a set of training images to one reference image which defines the number of sites used in the GIM. To this aim, the method described in the previous section can be used (taking covariance matrices proportional to the voxel size of each image). Alternatively, any other registration method can be applied.

The mean position and intensity values $\bar{v}_{k'}$, and covariance matrices $\bar{\Sigma}_{k'}$ are then estimated from the set of registered images. The Voronoi tessellation defined by the elements of the GIM determines which sites of the images are used to compute each site in the model. That is, the set of points that are

closest to each GIM site according to the Mahalanobis distance are searched for, and statistics are computed on them. Note that, as seen in section 3 the transformation T is also applied to the covariance matrix, i.e. the covariance associated to Tv_k is $T\Sigma_k T^t$.

The reference image to which all images are registered can be a rough model representing image structure, or one of the images in the training set. In this last case, in order to eliminate biases due to the choice of an arbitrary reference image [3], the process is iterated using the new average image as reference.

This same procedure can be used to update the model upon arrival of a new registered image. The means and covariances can be computed in an incremental, efficient manner.

6 Results

6.1 The role of covariance matrices in registration

As a first example, we provide an illustration of the workings of the GIM to incorporate prior knowledge into the registration process. Let us consider the simple synthetic data in figure 2(a). It consists of a set of 2D points arranged in the shape of a 3×3 grid. The covariance matrices are shown as 3σ ellipses. A second grid, consisting of a simple translation of the first one, is shown in figure 2(b). Registration results using Euclidean and Mahalanobis distances are shown respectively in figures 2(c) and 2(d). The covariance matrices encode the fact that there is a systematic tendency to move in diagonal directions, and this is used by the registration algorithm to find the correct transformation. Clearly, this very simplistic example is intended only as an illustration of the concept.

6.2 Landmark data

We now consider a more realistic case, and illustrate the construction and application of a GIM. The data employed for this example was obtained from a manual delineation of the corpus callosum (CC) in a slice of an MRI. A set of landmarks were identified. In order to generate sufficient data for our experiment, we perturbed these landmarks randomly as follows. First, we assigned covariance matrices to a subsample of the landmarks, thus defining Gaussian distributions on their location. We then sampled these distributions and interpolated the displacements (i.e. the difference between the original data set

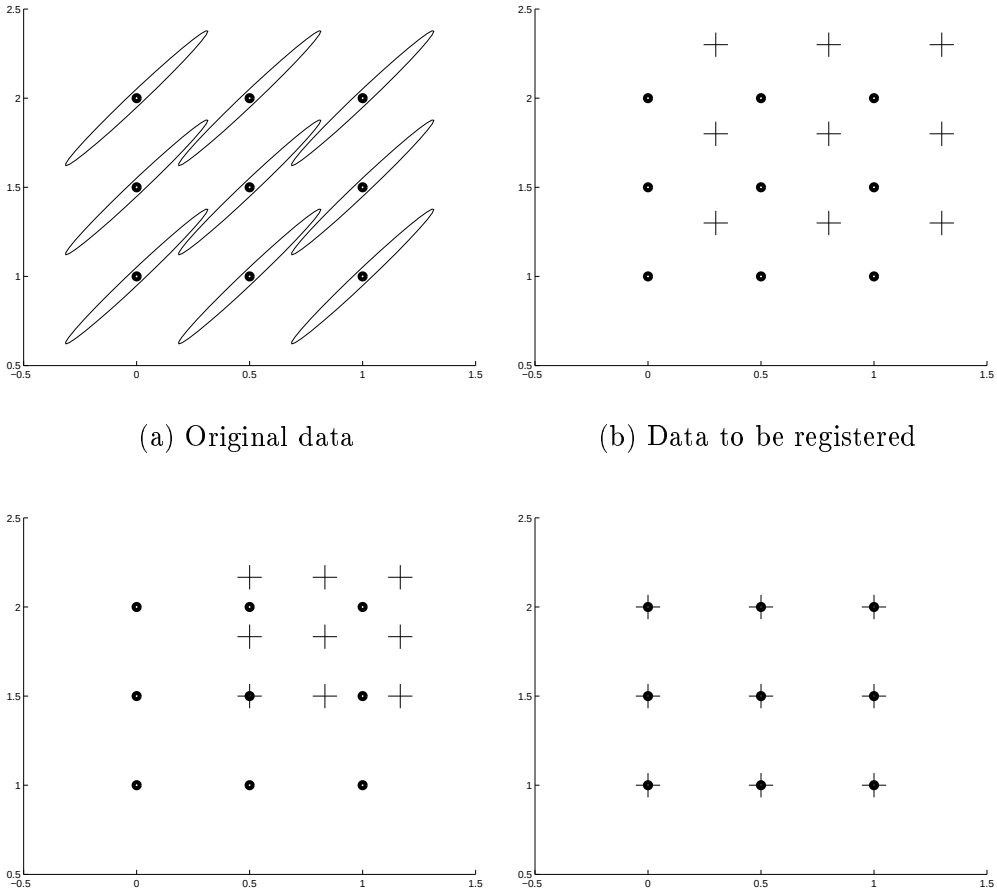

(a) Original data

(b) Data to be registered

(c) Euclidean matching

(d) Mahalanobis matching

Fig. 2. **(a)** Original data (circles), consisting of a simple grid with associated covariance matrices, **(b)** second grid, to be registered (plus signs), **(c)** results of registration (multiplication signs) using Euclidean distance, and **(d)** registration using Mahalanobis distance. The covariance matrices provide a priori knowledge about a systematic tendency to move in diagonal directions, which leads the algorithm to find the correct transformation.

and the new sample) using a thin-plate spline [14] to compute the remaining points. This effectively simulates anatomical variation in the shape of the CC. In order to simulate differences in geometry due to patient positioning and gross anatomy, each point set was perturbed by a randomly generated affine transformation. Figure 3(a) shows 50 data sets constructed in this way.

We now consider the construction of the GIM, which will encode the artificially-generated anatomical shape variation. First, we register the shapes by finding the best affine transformation. The results of the ICP algorithm applied to register all shapes to the first data set is shown in figure 3(b). Clearly, the thin shape of the CC results in ambiguous point matching, which leads the ICP to local minima. To solve this, information about the normal at each point

was added and used to discriminate good matches between points. Figure 3(c) shows the improved results.

Using the sites of the first data set as starting point, the landmarks of all data sets were clustered and their mean positions and covariance matrices computed. The resulting GIM is shown in figure 3(d). Note that the ellipses depict 3σ confidence bounds, i.e. there is a probability of 99.86% that the landmark be located within the ellipse.

The usefulness of the information encoded by the covariance matrices is shown next. A new point set is generated as described above. We first attempt to register it to the mean locations of the GIM using simple Euclidean distance measures in an affine ICP framework. No information about the normals at each point was employed. As shown in figure 3(e), this method fails to correctly register the frontal part (left in the figure) of the CC. On the contrary, the use of the covariance matrices, results in a good registration, as shown in figure 3(f). In this case, we used an ICP algorithm which utilizes Mahalanobis distance measures between point sites and GIM sites, again not using information about the normal. The new data set can now be used to update the GIM by computing local statistics.

6.3 Image data

In the previous examples, what we presented could be seen as a kind of robust and efficient ICP with a Mahalanobis distance. The main added value comes with the introduction of image intensities in the same framework. This is illustrated by the following example, where we tackle the construction of a GIM from image data. In this case, we consider the construction of a model of the appearance of the CC as seen on T1-weighted MRI. A rectangular image region containing the CC was isolated from a central sagittal slice of each of the 9 MRI data sets available. The sizes of such images were different for each data set. Figures 4(a) and 4(b) show two such images.

We aim at constructing a model of shape variability of the CC after factoring out affine transformations. The first step we perform is registering all images using the best affine transformation. To this end, we employed the algorithm described in [13], a block-matching multi-scale robust method. The result of registering the image in figure 4(b) to the image 4(a) is shown in figure 4(c).

The search for corresponding points between the affinely registered images can be seen as a non-rigid registration process. Given a voxel in an image, taken as the reference image, we aim at searching in the vicinity of such voxel in the rest of the images the most similar voxel. We employ the non-rigid registration method described in [15] to robustly compute these correspondences, which

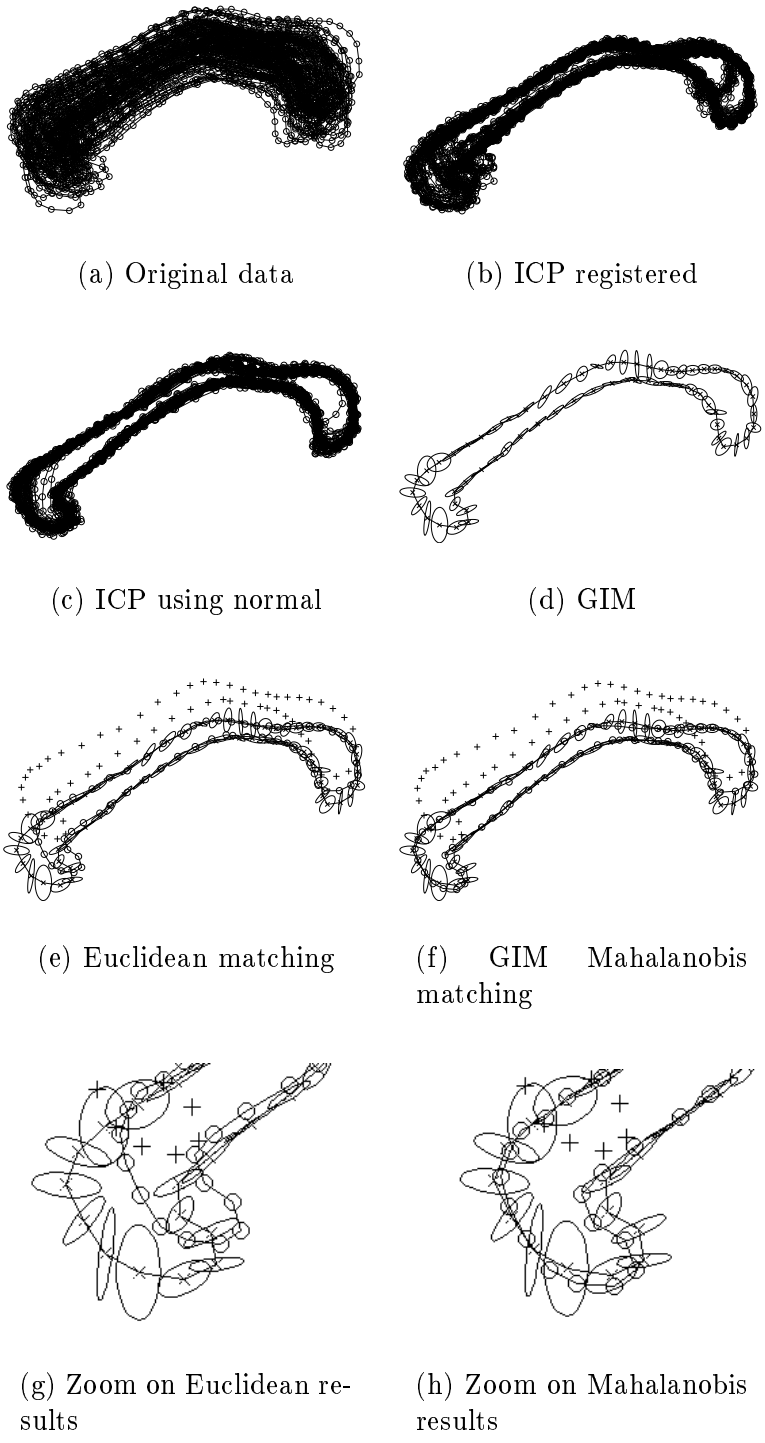

(a) Original data

(b) ICP registered

(c) ICP using normal

(d) GIM

(e) Euclidean matching

(f) GIM Mahalanobis matching

(g) Zoom on Euclidean results

(h) Zoom on Mahalanobis results

Fig. 3. **(a)** 50 landmark point sets representing corpora callosa contours; **(b)** their affine registration using ICP; **(c)** improved affine ICP registration using information about the normal at each point; **(d)** GIM computed from the registered data; **(e)** new shape registered to the positions of the GIM using Euclidean distance; **(f)** new shape registered to the GIM, but this time employing the information in the covariance matrices (Mahalanobis distance). **(g)** and **h** show enlarged views of the results in the anterior part of the CC, which is much improved in the Mahalanobis case.

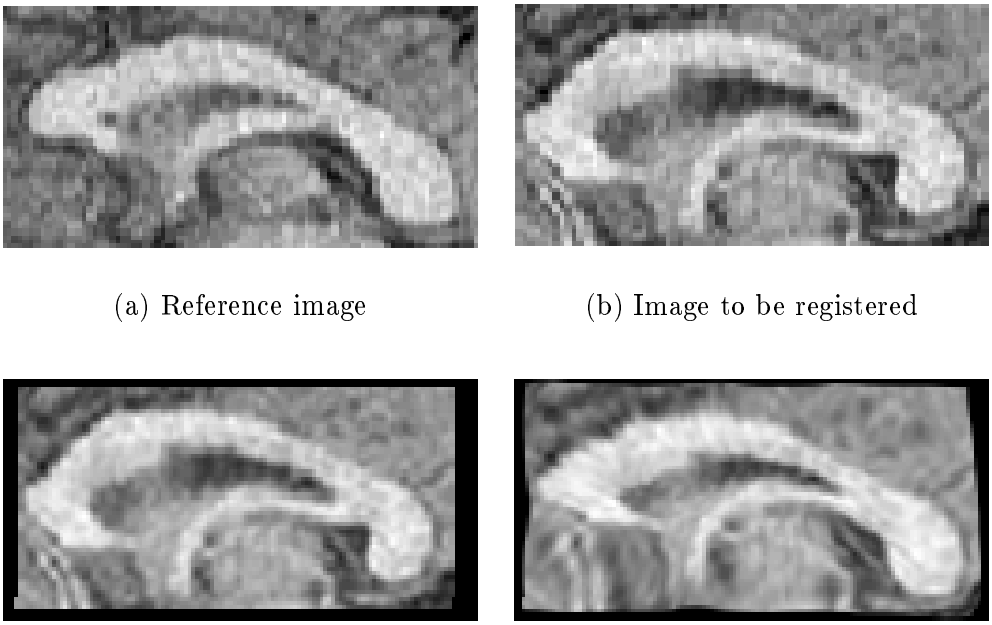

(a) Reference image (b) Image to be registered

(c) Affine registration (d) Non-rigid registration

Fig. 4. **(a)** and **(b)** Two images containing the corpus callosum as seen on T1-weighted MRI; **(c)** image B registered to A using an affine transformation; **(d)** image C registered to A using a non-rigid transformation.

will define a local deformation field at each voxel of the reference image, which will in turn be used to compute the statistics needed to build the GIM. The result of registering image 4(c) to 4(a) is shown in figure 4(d).

For the construction of the GIM we consider only a subset of the voxels in the reference image as model sites, in order to build sites that contain information about the shape variability of a set of voxels. In our case, a subsampling step of 100 was used. Local statistics were computed as described in section 5 on the set of closest points to each model site, after applying the corresponding non-rigid transformations. Two sites of the GIM are shown in figure 5. The left figure shows a site with a marked anisotropy in position (top), showing the local structure. The variation in intensity (bottom) is quite isotropic. Another site is shown in the right. In this case, the variability in position is isotropic (top), but the variation in intensity in not (bottom), showing a correlation of the intensity with respect to the position in the image.

7 Discussion and Conclusions

Notwithstanding the methodological contributions in the registration method, we consider that the most important element in our framework is the general-

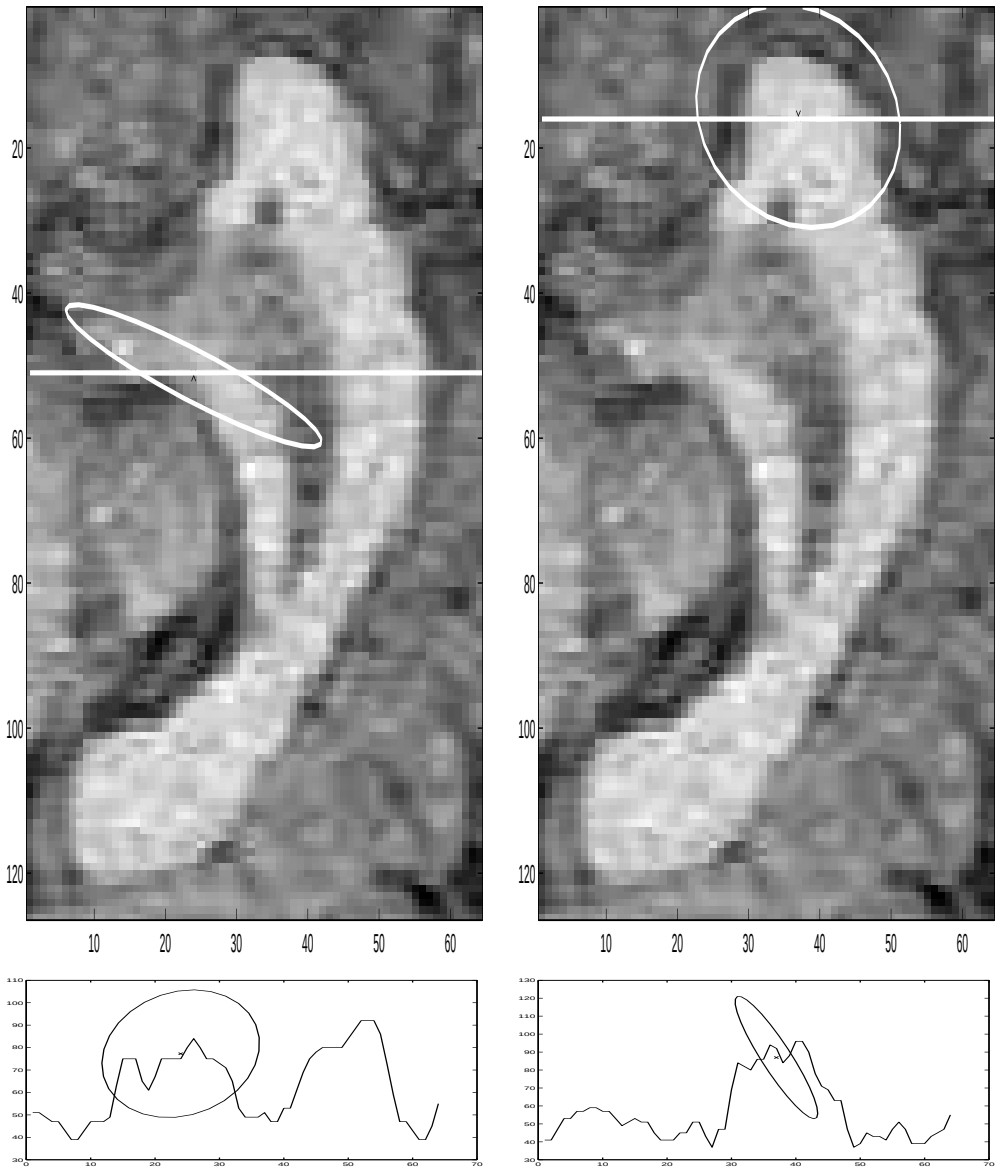

Fig. 5. Two sites (left and right), showing the spatial variation (top) and intensity variation along one axis (bottom). The site on the left shows a very anisotropic variation in position, reflecting the local structure, but a very isotropic variation in intensity. Conversely, the site on the right has a very isotropic variation in position, but the local variation in intensity is highly correlated to the position.

ized image model presented in section 2. Conceptually simple, it allows for the seamless integration of different types of information, and we are but starting to evaluate its real possibilities.

We believe that GIMs have the potential to pave the way for new approaches for image processing and analysis techniques that maximize the use of underlying image statistics. In the particular case of medical data sets, this technique is relevant to the construction of statistical shape models to drive segmenta-

tion methods or detect anatomical abnormalities. GIM-based registration, as illustrated by the ICP technique presented in this paper, allows the incorporation of prior knowledge into the registration process, thus making it more robust, potentially improving convergence rates and avoiding local minima.

The computation of a statistical model encompassing location and intensity information is the basis of the GIM. Several promising directions of future research are to be explored. The registration of point sets and meshes to images, as well as techniques for non-distorting resampling based on the GIM are being studied. The Kalman framework can be extended to model neighborhood interactions, and non-trivial system evolutions.

Regarding other works that share similarities with our method, we think that the closest one is [16]. In their work they considered the creation of shape models from unlabeled point sets. As in our case, no explicit correspondences between points is necessary. Their method, however, does not contemplate the construction of a representation of local structure, which is the base of our GIM. Furthermore, they did not consider the incorporation of intensity alongside with position at each site. Their registration method was based on thin-plate splines, which can be incorporated in our framework quite easily (results are forthcoming). Also, we provide a technique for incremental computation of the transformation via Kalman filtering.

Active appearance models [17] also share some similarities with our representation. Namely, the joint representation of intensity and position variability. However, their technique is based on principal components analysis and assumes explicit correspondences between points. On the other hand, it implicitly contains the notion of correlation between sites, which will be addressed in future work. Another work that employs a mixed position-signal representation, applied to fMRI signal analysis, can be found in [18]. This method is used to summarize in a compact representation the fMRI signal, but no correlation is considered between spatial and signal components.

Other future improvements of the method will focus on the extension of our Kalman framework to other types of transformation, such as thin-plate splines. More importantly, the issue of interpolation of the uncertainty contained in the covariance matrices will be studied. For this, we will employ concepts related to geostatistics, such as kriging and variogram analysis [19].

References

[1] M. A. González Ballester, A. Zisserman, M. Brady, Estimation of the partial volume effect in MRI, Medical Image Analysis 6 (4) (2002) 389–405.

[2] X. Pennec, J.-P. Thirion, A framework for uncertainty and validation of 3D registration methods based on points and frames, International Journal of Computer Vision 25 (3) (1997) 203–229.

[3] I. L. Dryden, K. V. Mardia, Statistical Shape Analysis, Wiley, 1998.

[4] T. F. Cootes, C. J. Taylor, D. H. Cooper, J. Graham, Active shape models - their training and applications, Computer Vision and Image Understanding 61 (2).

[5] P. Besl, N. McKay, A method for registration of 3D shapes, IEEE Transactions of Pattern Analysis and Machine Intelligence 18 (14) (1992) 239–256.

[6] J. Feldmar, J. Declerck, G. Malandain, N. Ayache, Extension of the ICP algorithm to nonrigid intensity-based registration of 3D volumes, Computer Vision and Image Understanding 66 (2) (1997) 193–206.

[7] J. P. Campbell, Speaker recognition: a tutorial, Proceedings of the IEEE 85 (9) (1997) 1437–1462.

[8] J. Sooful, E. Botha, An acoustic measure for automatic cross-language phoneme mapping, in: Twelfth Annual Symposium of the South African Pattern Recognition Association, 2001, pp. 99–102, http://www.prasa.uct.ac.za/Docs/Archive/PRASA2001/Speech%20Processing%20I/PDF/S6JJSooful.pdf.

[9] P. G. Batchelor, J. M. Fitzpatrick, A study of the anisotropically weighted Procrustes problem, in: Workshop on Mathematical Methods in Biomedical Image Analysis, 2000, pp. 212–218.

[10] G. Welch, G. Bishop, SCAAT: incremental tracking with incomplete information, in: SIGGRAPH, 1997.

[11] G. Welch, G. Bishop, An introduction to the Kalman filter, Tutorial at SIGGRAPH, http://www.cs.unc.edu/ welch/kalman/ (2001).

[12] N. Ayache, Atificial Vision for Mobile Robots - Stereo-Vision and Multisensor Perception, MIT Press, 1991.

[13] S. Ourselin, A. Roche, S. Prima, N. Ayache, Block matching: a general framework to improve robustness of rigid registration of medical images, in: MICCAI'2000, Vol. 1935 of Lecture Notes in Computer Science, Springer, 2000, pp. 557–566.

[14] F. L. Bookstein, Principal warps: thin-plate splines and the decomposition of deformations, IEEE Transactions on Pattern Analysis and Machine Intelligence 11 (6) (1989) 567–585.

[15] P. Cachier, E. Bardinet, D. Dormont, X. Pennec, N. Ayache, Iconic feature nonrigid registration: the PASHA algorithm, Computer Vision and Image Understanding 24 (4-5) (2003) 801–813.

[16] H. Chui, A. Rangarajan, Learning and atlas from unlabeled point-sets, in: IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA), IEEE Press, 2001, pp. 58–65.

[17] T. F. Cootes, G. J. Edwards, C. J. Taylor, Active appearance models, IEEE PAMI 23 (6) (2001) 681–685.

[18] G. Flandin, F. Kherif, X. Pennec, D. Rivière, N. Ayache, J.-B. Poline, Parcellation of brain images with anatomical and functional constraints for fMRI data analysis, in: IEEE International Symposium on Biomedical Imaging, Washington, USA, 2002, pp. 907–910.

[19] N. Cressie, Statistics for Spatial Data, Wiley, 1993.