

Gradient Intensity-Based Registration of Multi-Modal Images of the Brain

Ramtin Shams¹, Rodney A. Kennedy¹, Parastoo Sadeghi¹ and Richard Hartley^{1,2}
¹The Australian National University and ²NICTA*
Canberra, Australia

Abstract

We present a fast and accurate framework for registration of multi-modal volumetric images based on decoupled estimation of registration parameters utilizing spatial information in the form of ‘gradient intensity’. We introduce gradient intensity as a measure of spatial strength of an image in a given direction and show that it can be used to determine the rotational misalignment independent of translation between the images. The rotation parameters are obtained by maximizing the mutual information of 2D gradient intensity matrices obtained from 3D images, hence reducing the dimensionality of the problem and improving efficiency. The rotation parameters along with estimations of translation are then used to initialize an optimization step over a conventional pixel intensity-based method to achieve sub-voxel accuracy. Our optimization algorithm converges quickly and is less subject to the common problem of mis-registration due to local extrema. Experiments show that our method significantly improves the robustness, performance and efficiency of registration compared to conventional pixel intensity-based methods.

1. Introduction

Registration is a problem frequently encountered in medical image analysis. It is often required to align images acquired by different imaging devices and techniques (accounting for multi-modalities) and at different times (including long periods resulting in multi-temporal images). Registration allows for comparison of complimentary features, natural or pathological differences, data fusion and compiling atlases and is used in diagnosis, treatment planning and image guided surgery to name a few applications.

For certain applications it is impractical or impossible to use fiducial markers or landmarks or simply the images may have been acquired without such planning. The group of automatic registration methods align images by maximizing a

*NICTA is a research centre funded by the Australian Government through Backing Australia’s Ability and the ICT Research Centre of Excellence programs.

similarity measure using image information over a possible range of transformation parameters, or formally

$$T_{opt} = \underset{T}{\operatorname{argmin}} -S(I^F; I^{M(T)}), \quad (1)$$

where S is the similarity function to be maximized ¹, and I^F and $I^{M(T)}$ are the fixed and transformed moving images, respectively.

Since it is virtually impossible to try all combinations of transformation parameters, an optimization method such as Powell, simplex, gradient descent, etc has to be used. There are two major limitations with automatic registration methods: (a) the robustness of the optimization is dependent on the shape of the cost function and can often lead to mis-registration due to presence of local minima, (b) the calculation of the cost function involves image transformation and calculation of the similarity measure which can be time consuming.

In this paper, we present a method to overcome the above limitations and allow for accurate and efficient automatic registration of multi-modal images of the brain, including MR, PET, and CT scans.

2. Previous Work

Automatic retrospective registration has been extensively researched in the past decade with focus on MI-based methods. Plum *et al.* provide a comprehensive survey of MI-based methods in [7]. MI-based registration, pioneered by Collignon *et al.* [2] and Viola *et al.* [13], is widely regarded as a suitable similarity measure for multi-modal and multi-temporal images.

There are two broad approaches to solving the local minima problem for MI-based registration of images: (a) reducing local minima by improving the shape of the similarity function (b) skipping local minima by starting the optimization step closer to the actual alignment.

Studholme *et al.* [12] and Maes *et al.* [5] have proposed (virtually equivalent) normalized mutual information

¹In line with the optimization literature, our notation shows minimizing the equivalent cost function, $-S$.

(NMI) measures to reduce sensitivity to the amount of overlap between the images. Plum *et al.* [6] introduce spatial information to an MI-based cost function using a correction factor based on co-occurrence of gradient vectors. Liu *et al.* [4] propose an adaptive combination of intensity MI and gradient field MI with a multi-resolution approach to improve the shape of the cost function. The improved performance comes at the cost of increasing the computational complexity of the cost function calculation, which requires time consuming gradient image computation in addition to standard transformation and MI computation. Both methods have shown improved performance over NMI particularly for lower-resolution and down-sampled images. Despite favorable results, it cannot be analytically established whether the combined measures will always improve the smoothness of the cost function and can sufficiently cancel one-another's local minima and prevent formation of further ripples. Moreover, selection of coefficients, used in combined measures, is non-trivial and may well depend on the domain of the registration problem.

Mutual information of image gradients together with centroid-based estimation of scale and translations parameters have been used in [11] to improve the efficiency of the registration. The method is limited to 2D images and the centroid-based estimations are sensitive to partial overlap.

The performance and efficiency of MI-based registration can be improved by using multi-resolution methods, which iteratively register the images at lower resolutions and initialize the optimization for the later steps from the estimated parameters at the lower resolution. MI optimization at lower resolutions can still be trapped by local minima or even worse, an incorrect global minimum may be formed at lower resolutions, which can mislead the algorithm at higher levels. Multi-resolution methods can extend the capture range of the translation parameters, but do not usually improve the capture range of the rotation and scale parameters [10]. Registration at lower resolutions has also been reported to be more sensitive to the order in which the parameters are being optimized [5].

3. Approach and Contributions

We propose a method for registration of 3D images, comprising two main steps: first we acquire close estimates of the transformation parameters using image gradients, hence incorporating spatial information and second we use estimated parameters to initialize the optimization algorithm in the vicinity of the actual alignment and achieve sub-voxel accuracy while significantly improving the robustness and efficiency of the registration. Most other approaches improve the performance at the cost of increased computational complexity, whereas we improve the efficiency and performance at the same time.

The three rotation parameters are obtained by maximiz-

ing MI of 2D gradient intensity (GI) matrices of 3D volumes. The GI matrices are invariant to scale and translation between the images and allow us to determine rotation parameters on a 2D data-set. The reduction in dimensionality of the optimization problem and the fact that we only need to optimize 3 parameters significantly improves the efficiency of the algorithm. However, optimization of the GI-based cost function is challenging due to the shape of the cost function which is not very smooth and the presence of several local minima. This is a natural result of increased noise, quantization error, interpolation error, and limited size of the gradient kernels as well as reduction of the data-set from a 3D volume to a 2D matrix. We overcome this limitation by introducing a histogram calculation method which is not as sensitive to noise and gradient errors and also by proposing a robust optimization algorithm to retrieve the global minimum in the presence of local minima.

Given the rotation parameters, the images are brought into partial alignment. Assuming that the rotational estimation is close enough, the images will now differ by translation. The translation parameters are found by optimizing an MI-based cost function over pixel intensities of three 2D images, each obtained by reducing the 3D volume along a principal axis. This again allows us to obtain registration parameters efficiently by working on a much smaller 2D data-set.

We finally use a guided optimization algorithm on a conventional pixel intensity (PI)-based MI function to achieve sub-voxel accuracy. Our guided optimization algorithm is designed to converge quickly, leveraging from the knowledge that it has been initialized close to the actual alignment.

We claim two main novelties in this paper:

- We introduce *gradient intensity* (GI) as a measure of directional strength of an image, invariant to scale and translation between the images. The GI matrices are 2D data-sets calculated for 3D volumes and computationally efficient.
- We propose the *uniform volume histogram* method for generating the GI matrices, which is robust against noise and gradient calculation errors and results in better shaped cost functions and improves convergence and success rate of image registration.

Other contributions of this paper include:

- Proposing an optimization algorithm for finding the global minimum in the presence of local minima based on the Sobol quasi-random sequence and the simplex optimization.
- Implementation of an improved Powell-based optimization algorithm designed to converge quickly in the vicinity of the actual alignment for the final registration step.

4. Method

4.1. Gradient Map

Let \mathbf{I} be a 3D image and \mathbf{K}_x , \mathbf{K}_y and \mathbf{K}_z be 3D differentiating kernels in x , y and z directions, respectively. The gradient images are then calculated using

$$\mathbf{G}_x = \mathbf{K}_x \otimes \mathbf{I}, \quad \mathbf{G}_y = \mathbf{K}_y \otimes \mathbf{I}, \quad \mathbf{G}_z = \mathbf{K}_z \otimes \mathbf{I}, \quad (2)$$

where \mathbf{G}_x , \mathbf{G}_y and \mathbf{G}_z are 3D image gradient arrays in x , y and z directions and \otimes denotes 3D convolution. We used $9 \times 9 \times 9$ derivative of Gaussian kernels with standard deviation $\sigma = 2.0$ and used the separability feature of the kernels to speed up the convolutions. The resulting vector field is then expressed in spherical coordinates with $\mathbf{g}(\mathbf{x}) = [\rho_g(\mathbf{x}) \quad \phi_g(\mathbf{x}) \quad \theta_g(\mathbf{x})]^T$ components, where ρ_g , ϕ_g and θ_g represent magnitude, zenith and azimuth angles, respectively for the gradient vector at spatial location \mathbf{x} .

The GI matrices, as described in the next section, are based on gradient angles, so we would like to remove gradient vectors whose phases cannot be relied upon. To this end, we define a binary gradient map function f derived for 3D images using (2)

$$f(\mathbf{x}) = \begin{cases} 1, & \rho_g(\mathbf{x}) > t \\ 0, & \rho_g(\mathbf{x}) \leq t \end{cases}, \quad \mathbf{x} \triangleq [x \quad y \quad z]^T, \quad (3)$$

where t is the gradient magnitude threshold. Gradient angles are defined everywhere except where $\rho_g(\mathbf{x}) = 0$. However, smaller gradients are also more sensitive to noise, quantization and interpolation errors and may substantially change between transformed images.

Conventionally, an empirical or a statistical value such as RMS is used as the threshold. However lack of a principled method for selecting t may either result in removal of important gradient vectors or leave too many sensitive vectors behind. To resolve this dilemma, we use a noise-resistant phase histogram method, described in 4.3, which also removes sensitivity to the gradient threshold and allows us to conveniently set $t = 0$.

4.2. Gradient Intensity

Gradient intensity is a measure of directional strength of an image and is defined as the number of significant gradient vectors which pass through a given area on the surface of the unit sphere. Fig. 4.2 depicts gradient intensity mapped on the surface of the unit sphere with sample gradient vectors.

We define the gradient intensity at the center of the area specified by the angular resolution $2\delta\phi$ and $2\delta\theta$

$$\Gamma(\phi, \theta) = \int \int \int h_\phi(\mathbf{x}) h_\theta(\mathbf{x}) f(\mathbf{x}) dx dy dz, \quad (4)$$

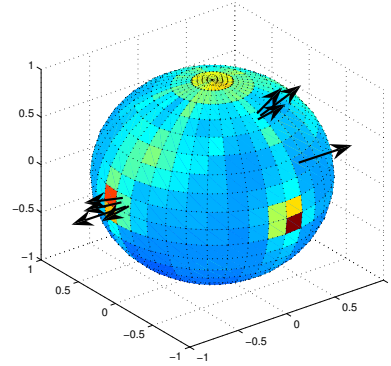


Figure 1. Gradient intensity mapped on the surface of the unit sphere. Intensity of each cell is a function of the number of gradients that pass through the cell. The gradient vectors are displayed for three cells.

where Γ is the gradient intensity at the center of the bin specified by (ϕ, θ) and $h_\phi(\mathbf{x})$ and $h_\theta(\mathbf{x})$ are defined as

$$h_\phi(\mathbf{x}) = \begin{cases} 1, & \phi - \delta\phi \leq \phi_g(\mathbf{x}) < \phi + \delta\phi \\ 0, & \text{elsewhere} \end{cases}, \quad (5)$$

$$h_\theta(\mathbf{x}) = \begin{cases} 1, & \theta - \delta\theta \leq \theta_g(\mathbf{x}) < \theta + \delta\theta \\ 0, & \text{elsewhere} \end{cases}, \quad (6)$$

To compare gradient intensity of different volumes we use a normalized version of (4)

$$\tilde{\Gamma}(\phi, \theta) = \frac{\Gamma(\phi, \theta)}{\max \Gamma(\phi, \theta)}. \quad (7)$$

Gradient intensity is different from a Gaussian sphere [3] representation of the orientation. A Gaussian sphere representation only captures the orientation of the surface normals, ignoring information inside the object (*e.g.* internal brain organs). It also assumes that there is an object to start with, hence implying that the object of interest is already segmented. The gradient intensity does not require segmentation and utilizes all the information which is in a 3D medical image and as such, is more suitable for the registration problem.

4.3. Uniform Volume Histogram

Equation (4) assumes that a continuous (and integrable) representation for gradient phases (ϕ_g and θ_g) of the image volume exists. In the absence of a continuous representation, equation (4) can be replaced by a standard 2D histogram. Standard histogram representations are sensitive to the number of bins, gradient threshold selection, noise, quantization, and interpolation errors.

In [9], the authors propose a probability density estimation for 2D image pixel-intensities based on a continuous representation of image intensities and report improved

robustness of the probability distribution function against noise and selection of the number of bins. The method in [9] assumes that the intensity in a triangle comprised of three adjacent pixels can be estimated as a linear function of the intensity of the vertices. This in effect fits a plane to the vertices. This method could be adapted to improve gradient histograms as well, however, the extension to 3D volumes requires calculation of hyper-planes and simplex structures and is unnecessarily time consuming.

We use a histogram algorithm called uniform volume, which assumes a continuous representation for the gradient field $\mathbf{g}(\mathbf{x}) = [\rho_g(\mathbf{x}) \ \phi_g(\mathbf{x}) \ \theta_g(\mathbf{x})]^T$ exists locally within each cubic volume \mathbf{C} comprising 8 adjacent vertices of the gradient image. Let $V = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_8\}$ represent the set of vertices. We assume that the gradient phase function within \mathbf{C} is bounded by the values of function at the vertices of \mathbf{C} , *i.e.*, for all $\mathbf{x} \in \mathbf{C}$

$$\begin{aligned} \min_{x \in V} \phi_g(\mathbf{x}) &\leq \phi(\mathbf{x}) \leq \max_{x \in V} \phi_g(\mathbf{x}), \\ \min_{x \in V} \theta_g(\mathbf{x}) &\leq \theta(\mathbf{x}) \leq \max_{x \in V} \theta_g(\mathbf{x}). \end{aligned} \quad (8)$$

Now assume a gradient intensity binning scheme with angular resolution of $2\delta\phi$ and $2\delta\theta$. Let (ϕ_i, θ_j) be the center of a bin which falls within the extents of the cubic volume \mathbf{C} , *i.e.*,

$$\begin{aligned} \min_{x \in V} \phi_g(\mathbf{x}) &\leq \phi_i \leq \max_{x \in V} \phi_g(\mathbf{x}), \\ \min_{x \in V} \theta_g(\mathbf{x}) &\leq \theta_j \leq \max_{x \in V} \theta_g(\mathbf{x}), \end{aligned} \quad (9)$$

if the gradient field has the property that

$$\int \int \int_{\mathbf{C}} h_\phi(\mathbf{x}) h_\theta(\mathbf{x}) f(\mathbf{x}) dx dy dz \quad (10)$$

is the same for all 2D bin pairs represented by their center (ϕ_i, θ_j) , then the histogram contributions have to be uniformly distributed between bins that satisfy (9). Simply put, the condition in (10) allows us to uniformly distribute histogram contributions to the bins which fall within the extents of the cubic volume vertices.

One benefit of the uniform volume distribution as described above is that it is computationally inexpensive, since it does not explicitly require calculation of $\mathbf{g}(\mathbf{x})$ except at the bounding vertices. The histogram distribution favors areas of the image where the second derivative of the intensity is smaller or in other words the gradient function is smoother. This is in line with the findings in [9] for a robust histogram with less sensitivity to noise. We also designed the histogramming method to suppress the contributions of the areas of the image where there is a large discrepancy between the function values at the adjacent vertices.

Compare the 1D histograms of gradient vectors' azimuth angle θ for an MR-T1 volume and its noisy version in

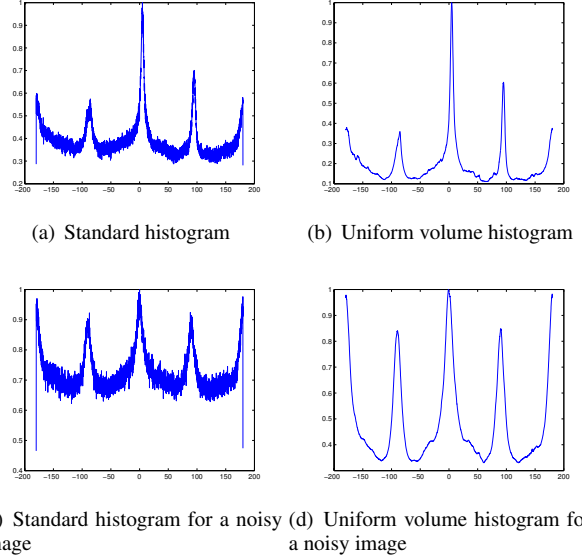


Figure 2. Uniform volume histograms exhibit smoother curves and better resistance to noise.

Fig. 2, where white Gaussian noise with $\sigma = 0.1$ is added. Notice that uniform volume histograms are smoother and less affected by noise compared to standard histograms. These properties translate to smoother cost functions with fewer local minima, in the later stages of the algorithm (see Fig. 3), and allow our optimization method to converge more easily.

4.4. Mutual Information of Gradient Intensity

MI of the gradient intensity of two images is a measure of directional similarity of the images, regardless of their relative size (scale) and position (translation) and is maximum, where images are rotationally aligned. Since the gradient intensity is a 2D function of azimuth and zenith angles, we can find the rotational misalignment between the images by calculating MI over a 2D data-set for 3D images, hence reducing the dimensionality of the problem.

4.5. Estimating Rotation Parameters

The 2D GI matrices are calculated for the fixed and moving images. We use an angular resolution of $2\delta\phi = 2\delta\theta = 1.0^\circ$, which results in GI matrices of 180×360 size regardless of the size of the original 3D volumes. To find the optimal rotation parameters, we apply the transformation directly to the GI matrix of the moving image. As illustrated by Fig. 4, the GI matrix is transformed from spherical coordinates to cartesian coordinates, the Euler transformation $T(\alpha, \beta, \gamma)$ is applied and the result is transformed back to spherical coordinates. Note that, based on the definition of gradient intensity, $\rho = 1$ for spherical/cartesian conversions.

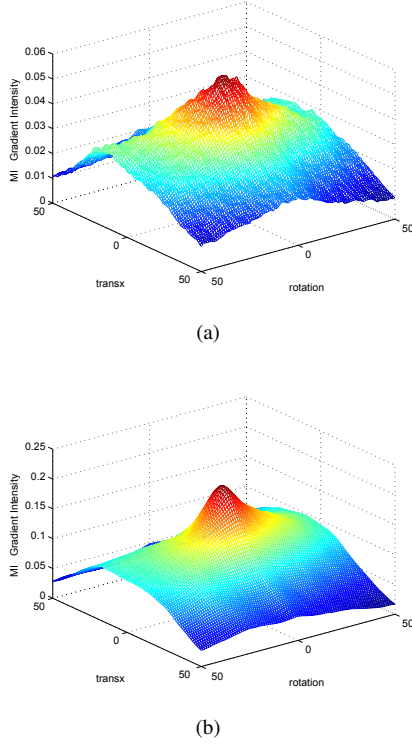


Figure 3. Comparison of 2D MI functions based on (a) standard and (b) uniform volume histogramming. Uniform volume histogram results in a better shaped and smoother cost function.




Figure 4. The process of transforming the GI matrix of a moving image by the three Euler angles α , β and γ .

The rotation parameters are then located by finding the maximum of the MI between the fixed GI matrix G_h^F , and the transformed GI matrix $G_h^{M(T)}$

$$[\alpha_t, \beta_t, \gamma_t] = \underset{\alpha, \beta, \gamma}{\operatorname{argmin}} -MI(G_h^F; G_h^{M(T)}). \quad (11)$$

This is a 3-parameter optimization problem over a 2D data-set which is much easier to solve than the original 9-parameter (for similarity registration) or 6-parameter problem (for rigid registration) over the much larger 3D data-set. The use of MI as a non-linear and statistical similarity measure is justified due to the multi-modal nature of the images which will result in non-linearly related GI matrices.

4.6. Soblex Optimization

The GI cost function attains a global minimum where the images are rotationally aligned. The cost function is relatively smooth, however, a standard optimization algorithm such as *Powell* or *simplex* [8] can easily be trapped by

local minima and fail to converge to the global minimum. We propose a robust optimization method based on simplex and sampling of the parameter space with the Sobol quasi-random sequence [8], called *Soblex* optimization.

The Soblex optimization is initially given a budget, in terms of time or number of cost function calls. Within the initial budget, Soblex evaluates the cost function using the Sobol sequence and initializes a simplex-shaped subspace, which is constructed from points with the lowest costs. The Sobol sequence ensures that we can progressively sample the parameter space in a virtually uniform fashion. Intuitively, if the budget is large enough, the simplex subspace can sufficiently close in on the global minimum to allow successful execution of the optimization algorithm.

Unlike most optimization methods that start with a single point in space, the simplex algorithm allows us to start from a region of space that can arbitrarily be made close to the global minimum by increasing the Soblex budget.

4.7. Estimating Translation Parameters

When images are rotationally aligned, translation (and scale) parameters can be estimated by optimizing the parameters on the reduced 2D images (see Fig. 5). Let $I(\mathbf{x})$ be the image intensity at voxel \mathbf{x} , and I_x, I_y, I_z reduced images in the x, y and z directions, respectively obtained by averaging the pixel intensities across the corresponding dimensions. Translation parameters x_t, y_t , and z_t can be found using

$$[y_t, z_t] = \underset{y, z}{\operatorname{argmin}} -MI(I_x^F; I_x^{M(T)}), \quad (12)$$

$$[z_t, x_t] = \underset{z, x}{\operatorname{argmin}} -MI(I_y^F; I_y^{M(T)}), \quad (13)$$


$$[x_t, y_t] = \underset{x, y}{\operatorname{argmin}} -MI(I_z^F; I_z^{M(T)}), \quad (14)$$

where I^F , and $I^{M(T)}$ are fixed and translated moving images, respectively and MI is the similarity function. We perform all three optimizations and average the results to get more robust estimations.

We use estimated rotation parameters to bring the images into rotational alignment, then reduce the volumes along principal axes and use Soblex to find the translation parameters. Note that the translation estimation shows some resistance to noise due to the averaging of the image layers. One could also include the isotropic scale parameter at this stage for similarity (rigid + scale) registration.

4.8. Final Optimization

We use estimations of rotations, translations and scale in order to initialize an optimization algorithm based on Powell's multi-dimensional direction set. The final optimization round is performed using the standard PI-based



(a) Reduced MR-T1 (b) Reduced CT

Figure 5. (a) MR-T1 and (b) CT images, rotationally aligned and averaged along the z axis. The 2D images are used to determine the x and y translation between the images. Note that the silhouettes are similar, which allows the MI-based optimization to quickly converge and return the translation parameters.

MI cost function on the full 3D volumetric data to achieve sub-voxel accuracy. Our implementation of Powell resembles the method described in [8] with one major difference that it takes advantage of the fact that it has been initialized not far from the minimum and as such, can quickly converge by refraining from checking the perimeter excessively. We define a minimum distance or resolution for the cost function. The cost function keeps track of each point in the N -dimensional space that it evaluates and will only evaluate a new point if it falls outside all previously evaluated points by a specified minimum distance. Obviously, this method is not suitable as a general purpose optimization tactic, and can only be used when the optimization algorithm can be properly initiated.

The choice of final optimization method is somewhat arbitrary. We chose Powell, primarily because it does not require calculation of the cost function gradient, only needs one point for initialization, and could be more easily adapted to our finite resolution method.

5. Results

We evaluated the performance and efficiency of our method using various images from the *Retrospective Image Registration Evaluation* project (RIRE) database [1], where a gold standard for registration using fiducial markers was known. We considered CT to MR-T1, T2, PD and PET to MR-T1, T2, PD and MR-T2 to T1 registration. The images were rather low quality and low resolution (refer to Table 1), which made the registration more challenging and allowed us to experiment with our method under a more difficult condition.

The images in the RIRE database were brought into alignment using the gold transformation, then 100 random rigid transformations were applied to one of the images for each pair and our method was used for registration. The dynamic range of transformations were $[-25^\circ, 25^\circ]$ for rotation parameters and $[-32 \text{ mm}, 32 \text{ mm}]$ for translation parameters. The conventional PI-based method was

| Image | Dimensions | Voxel Size (mm) |
|-------|----------------------------|--------------------------------|
| MR | $256 \times 256 \times 26$ | $1.25 \times 1.25 \times 4.00$ |
| CT | $512 \times 512 \times 29$ | $0.65 \times 0.65 \times 4.00$ |
| PET | $128 \times 128 \times 15$ | $2.59 \times 2.59 \times 8.00$ |

Table 1. Resolution and size of images used in the experiments.

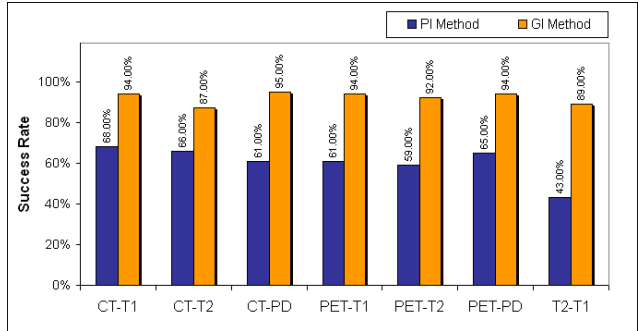


Figure 6. Superior performance of the GI-based method for various combinations of modalities.

used as the baseline for comparison. We used the standard implementation of Powell for the conventional PI-based method, which was initialized from the origin of the parameter space. The registration errors were calculated by averaging the registration errors of 10,000 equally spaced points within the brain volume. Soblex optimizations were given an initial budget of 1000 cost function calls for rotation estimation and converged with an average of 1100 calls (including the initial budget). For translation estimation Soblex with an initial budget of 100 function calls was used and converged with an average of 150 iterations. For translations a much smaller initial budget can be used as the cost function is much smoother for the pixel intensities compared to the gradient intensities.

5.1. Performance

We declared a registration failed if the average error exceeded the diagonal voxel size of the moving image, which was approximately 8.80 mm for PET images and 4.40 mm for CT and MR images. Fig. 6 compares the success rate of our method with the conventional method for various registration pairs and shows the superior performance of our method. Note that our method outperforms the PI-based method by a large margin, even if a lower threshold is chosen for registration errors. This can be better demonstrated by looking at the mean and median error graphs presented in the next section.

5.2. Accuracy

Fig. 7 shows average and median registration errors for each method. Registration errors are significantly reduced




Figure 7. Improved accuracy of the GI-based method with a lower mean and median error is demonstrated. PI mean, GI mean, PI median and GI median errors are shown for each combination of modalities from left to right.

using the GI-method, due to the ability of the method to bypass local minima. Also note that, median and mean errors are close for the GI method, which indicates the superior robustness of the method.

The improved statistical accuracy of our method compared to the standard PI-based method, as demonstrated in Fig. 7, is the result of improved robustness and success rate of our method, whereas the PI method fails to converge for large misalignments. However, we emphasize that where both methods converge, the accuracy of the standard method is similar to our method. This is because both methods use the same MI-based registration as the final stage.

Interpretation of errors should be treated with care. A common approach is to identify a number of volumes of interests (VOIs) and average the distance between centroids of VOIs, in registered images [14]. Obviously, error calculations will depend on the selection of VOIs and their distance from the center of transformation, as well as accuracy of segmentation of VOI pairs. The advantage of this method is that the error is meaningful for anatomical or pathological details of interest. The statistical-based method, on the other hand, does not require segmentation and definition of VOIs. However, it may overestimate the errors by inclusion of points that may not be of interest. Hence, one should not compare error calculations obtained by different methods.

5.3. Efficiency

Both the GI and PI-based methods finally register the images using a Powell based optimization over pixel intensities. The computational cost of the methods are directly comparable at this stage. We included the computational cost of pre-processing required by the GI method by increasing the number of Powell iterations. The Soblex optimization stage of the GI method was approximately $1/100^{th}$ of the cost of each iteration for the final Powell-based optimization or on average equal to 11 Powell optimizations. This has been factored in the end-results to allow

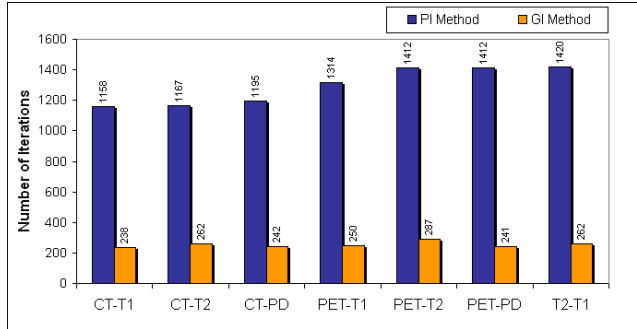


Figure 8. Superior efficiency of the GI-based method for various combination of modalities. The number of iterations are significantly reduced using the GI method. Computation times are proportional to the number of iterations.




Figure 9. A PET volume (a) and its noisy version (b) registered to MR-T1 (c) in our experiments. The original PET volume itself has a low resolution and SNR.

a fair comparison. Another 15 iterations were added to account for gradient intensity calculations. Fig. 8 compares registration efficiency of the GI and PI-based methods, and shows around 500% improvement as a result of using the GI method.

5.4. Robustness w.r.t. Noise

Gradient based methods are often more sensitive to noise. To demonstrate the robustness of the method w.r.t to noise, we added white Gaussian noise with $\sigma = 0.1$ to a PET image as shown in Fig. 9 and registered the image to MR-T1 100 times with random transformations applied to the PET image. The mean and median errors were 4.99 mm and 4.57 mm, up almost by only 1.00 mm compared to registration results without noise. The good results are to some extent due to the noise reduction feature of the uniform volume histogram (other noise resistant elements of our algorithm include the translation estimation process and the Gaussian gradient kernel itself).

5.5. Robustness w.r.t. Image resolution

Initial estimates using the GI-based method are invariant to scale. We conducted another experiment for PET to T1 registration where we scaled the size of T1 and PET images in each direction by 0.5. 100 random transformations were applied, the resulting mean and median registration er-

rors were 2.02 mm and 1.74 mm, respectively. The results outperformed registration results obtained at the higher resolution. This suggests that our method could be further improved if combined with a multi-resolution scheme. However, further validations on other modalities and at lower resolutions are required to draw a concrete conclusion.

6. Discussion

Gradient intensity is useful in determining the spatial orientation of 3D images and can be used to rotationally align the images, irrespective of their relative size and position. We experimented with various modalities of brain images and showed that our method is robust and performs well for all combinations, including the low resolution and conventionally challenging PET images.

We used an MI-based registration method in the final registration stage. However, the GI-based registration can be combined with other registration techniques and similarity measures. The registration performance using our method was on average 92%, up almost by 30% compared to the conventional PI-based method. The computational efficiency improved on average by 510% and median error was down from an average 4.13 mm to 2.14 mm across the whole range of our experiments comprised of 700 registrations for a relatively large range of misalignments.

In our experiments, we encountered a few cases where the final optimization algorithm had to recover from a completely incorrect initial estimate. This is not necessarily worse than starting from any other arbitrary position within the parameter space, such as the origin. However, the algorithm converges with performance and efficiency levels comparable to the PI-based method. Such cases can be further avoided and the performance can be improved by using a method that can estimate translation and scale independent of the rotation. Another avenue for getting better performance is to improve the GI-based cost function to obtain more accurate and more robust initial estimates. One such improvement is to consider a weight function for GI calculation to compensate for the non-linear size of the GI bins on the unit sphere. Another potential extension is to add a multi-resolution scheme for GI-based registration. We are also investigating the possibility of adapting our method for affine and non-rigid registration.

7. Acknowledgements

The authors would like to thank Dr. Sebastien Ourselin for reviewing the paper prior to submission and providing valuable feedback. The images and the standard transformation(s) were provided as part of the project, Retrospective Image Registration Evaluation, National Institutes of Health, Project Number 8R01EB002124-03, Principal Investigator, J. Michael Fitzpatrick, Vanderbilt University, Nashville, TN.

References

- [1] *Retrospective Image Registration Evaluation Project*. <http://insight-journal.org/rire/>, 2007.
- [2] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal. Automated multimodality medical image registration using information theory. In *Proc. Int. Conf. Information Processing in Med. Imaging: Computational Imaging and Vision 3*, pages 263–274, Apr. 1995.
- [3] B. K. P. Horn. *Robot Vision*. MIT Press, The Massachusetts Institute of Technology, 1986.
- [4] J. Liu, J. Tian, and Y. Dai. Multi-modal medical image registration based on adaptive combination of intensity and gradient field mutual information. In *IEEE Int. Conf. of Engineering in Medicine and Biology Society (EMBS)*, pages 1429–1432, Aug. 2006.
- [5] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality image registration by maximization of mutual information. *IEEE Trans. Med. Imaging*, 16(2):187–198, Apr. 1997.
- [6] 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. on Med. Imaging*, 19(8):809–814, Aug. 2000.
- [7] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever. Mutual-information-based registration of medical images: A survey. *IEEE Trans. on Med. Imaging*, 22(8):986–1004, Aug. 2003.
- [8] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes in C*. Cambridge University Press, Cambridge, second edition, 1992.
- [9] A. Rajwade, A. Banerjee, and A. Rangarajan. A new method for probability density estimation with application to mutual information based image registration. In *Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*, June 2006.
- [10] R. Shams, R. A. Kennedy, and P. Sadeghi. Efficient image registration by decoupled parameter estimation using gradient-based techniques and mutual information. In *Proc. IEEE Region 10 Conf. (TENCON)*, Taipei, Taiwan, Oct. 2007.
- [11] R. Shams, P. Sadeghi, and R. A. Kennedy. Gradient intensity: A new mutual information based registration method. In *Proc. IEEE Computer Vision and Pattern Recognition (CVPR) Workshop on Image Registration and Fusion*, Minneapolis, MN, June 2007.
- [12] C. Studholme, D. L. G. Hill, and D. J. Hawkes. An overlap invariant entropy measure of 3d medical image alignment. In *Pattern Recognit.*, volume 32, pages 71–86, 1999.
- [13] P. Viola and W. M. Wells III. Alignment by maximization of mutual information. In *Proc. Int. Conf. Computer Vision (ICCV)*, pages 16–23, June 1995.
- [14] 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. Pollock, C. A. Pelizzari, R. A. Robb, D. Hanson, and R. P. Woods. Comparison and evaluation of retrospective intermodality brain image registration techniques. *Journal of Computer Assisted Tomography*, 21(4):554–566, 1997.