

An Uncertainty-Driven Hybrid of Intensity-Based and Feature-Based Registration with Application to Retinal and Lung CT Images^{*}

Charles V. Stewart, Ying-Lin Lee, and Chia-Ling Tsai

Rensselaer Polytechnic Institute
Troy, NY 12180-3590 USA

Abstract. A new hybrid of feature-based and intensity-based registration is presented. The algorithm reflects a new understanding of the role of alignment error in the generation of registration constraints. This leads to an iterative process where distinctive image locations from the moving image are matched against the intensity structure of the fixed image. The search range of this matching process is controlled by both the uncertainty in the current transformation estimate and the properties of the image locations to be matched. The resulting hybrid algorithm is applied to retinal image registration by incorporating it as the main estimation engine within our recently published Dual-Bootstrap ICP algorithm. The hybrid algorithm is used to align serial and 4d CT images of the lung using a B-spline based deformation model.

1 Introduction

Feature-based and intensity-based registration algorithms differ considerably in the image-based features that drive the alignment process [4]. Intensity-based techniques use all image pixels and do not require explicit feature extraction. They tend to be more stable around minima of the objective function because they don't rely on uncertain feature locations or on correspondences which may fluctuate with slight changes in the transformation. On the other hand, feature-based techniques tend to be faster, have a wider capture range, and allow alignment to be focused on only selected subsets of the image data. These strengths and weaknesses are studied experimentally in [9].

A growing set of papers has begun to address the issue of combining feature-based and intensity-based registration. As examples, in [5] registration is driven both by intensity similarity error and by errors in the positions of matched features. In Feldmar et al. [7] intensity is treated as a 4th dimension for ICP matching, while Sharp et al. [13] combines ICP with invariant features. In a much

^{*} The support of the GE Global Research Center and of the National Science Foundation through the Center for Subsurface Sensing and Imaging Systems is gratefully acknowledged.

different approach, Aylward et al. [1] locate tubular structures in one image and then align images using gradient-descent of a “tubularness” measure evaluated at the transformed locations of these structures in the other image. Ourselin, et al. [10] use block-matching of intensities together with robust regression. Shum and Szeliski [15] use block matching of intensities to refine video mosaics. The PASHA algorithm [4] combines correspondence and deformation processes in a global objective function. In the HAMMER algorithm [14] registration is driven by the alignment of feature positions based on a hierarchical description of the surrounding intensity distribution.

This paper presents a new hybrid of intensity-based and feature-based registration and applies the resulting algorithm in two contexts: aligning low-overlap retina images and spline-based alignment of lung CT volumes. Feature points found in the moving image are matched against the intensity structure of the fixed image. Importantly, the search range for the match is dictated by both the properties of the feature and the uncertainty in the current transformation estimate. This differs from common intensity-based techniques where constraints are driven by local changes in the similarity measure. It also differs from correspondence-based methods like ICP [2] where matching is purely a nearest-point search.

This core algorithm is built into two different overall registration algorithms. The first is an extension of our recent Dual-Bootstrap ICP [17] algorithm for aligning retinal images. Replacing the feature-to-feature ICP algorithm with the new feature-to-intensity similarity matching increases the effectiveness for extremely different cases. The second overall algorithm is a new technique for non-rigid, B-spline alignment of lung CT volumes. Small regions with sufficient intensity variation in the moving image are matched against the fixed image. These matches control the estimate of hierarchical B-spline deformations. The algorithm is applied to effectively align serial CT volumes.

2 The Role of Uncertainty in Registration

We motivate the new hybrid algorithm by considering the importance of uncertainty in registration. Let $S(\mathbf{p}, \mathbf{q})$ measure the region-based similarity error between moving image I_m at location \mathbf{p} and fixed (target) image I_f at location \mathbf{q} . Let $T(\mathbf{p}; \boldsymbol{\theta})$ be the transformation mapping \mathbf{p} from I_m onto I_f based on the (to be estimated) parameter vector $\boldsymbol{\theta}$. Finally, let \mathcal{P} be a set of locations in I_m where transformation constraints are applied.

Intensity-based algorithms search for the transformation minimizing the aggregate similarity error:

$$E(\boldsymbol{\theta}) = \sum_{\mathbf{p}_i \in \mathcal{P}} S(\mathbf{p}_i, T(\mathbf{p}_i, \boldsymbol{\theta})), \quad (1)$$

Regularizing constraints may be placed on $\boldsymbol{\theta}$, but we will ignore these for now. $E(\boldsymbol{\theta})$ is most frequently minimized through a gradient-descent technique [8]. The estimate, $\hat{\boldsymbol{\theta}}$, remains uncertain throughout the minimization process. This




Fig. 1. Two examples of a one-dimensional cross-section of the similarity error function surface surrounding a transformed point \mathbf{p}' . The interval over which the function is plotted is the uncertainty region. The arrow shows the direction of gradient descent based on the point. The location \mathbf{q} is the minimum of the similarity error function within the search region. In (a) the gradient descent direction is in the direction of the minimum, whereas in (b) the direction is away from the minimum.

in turn causes uncertainty in the location $T(\mathbf{p}_i, \hat{\boldsymbol{\theta}})$. We investigate the effect of this on the constraints used in registration.

Suppose we have a measure of the uncertainty on $\mathbf{p}'_i = T(\mathbf{p}_i, \hat{\boldsymbol{\theta}})$ and this can roughly be described in terms of a standard deviation σ . Consider a region of width $\pm c\sigma$ around \mathbf{p}'_i in I_f . If the uncertainty model is reasonably accurate, this region will likely contain the true homologous point in image I_f to \mathbf{p}_i . Call this point \mathbf{q}_i . Moreover, if the similarity error function is correct, this point will also minimize $S(\mathbf{p}_i, \mathbf{q})$ for all \mathbf{q} within this uncertainty region. If the gradient direction $\nabla S(\mathbf{p}_i, \mathbf{q})$ evaluated at \mathbf{p}'_i points in the direction of \mathbf{q}_i , then the gradient descent direction computed for point i is consistent with the direction the estimate needs to move (based on the local constraints). This happens when the search region is convex, as illustrated in Figure 1(a). On the other hand, if the gradient points away from \mathbf{q}_i then the constraint will pull the transformation estimate in the wrong direction (Figure 1(b)). This is likely to happen when the uncertainty region is large. When it happens at enough points, registration will descend into an incorrect local minimum. This problem can occur no matter at what resolution the images are represented. A similar problem occurs in ICP-like matching when the uncertainty region contains multiple points and the closer point drives registration in the wrong direction.

3 Iterative Most Similar Point (IMSP)

The natural alternative to gradient-based minimization is to search for the most similar location, \mathbf{q}_i , surrounding \mathbf{p}'_i . This leads to a conceptually straightforward registration algorithm: iteratively match the feature points by minimizing the

similarity error and estimate the transformation based on the matches. The trick (and primary novelty) is using uncertainty to dictate the search region for minimizing the error. A summary outline of this “IMSP” procedure is presented below, followed by a sketch of important details:

1. Given are initial transformation estimate $\hat{\theta}_0$, moving image I_m , fixed image I_f , and feature point set \mathcal{P} from I_m .
2. Initialize $t = 1$.
3. For each $\mathbf{p}_i \in \mathcal{P}$,
 - a) Use the properties of \mathbf{p}_i and the uncertainty in the current estimate to determine a search region $N(\mathbf{p}'_i)$ surrounding $\mathbf{p}'_i = T(\mathbf{p}, \hat{\theta}_{t-1})$ in I_f .
 - b) Find the location \mathbf{q}_i that minimizes $S(\mathbf{p}_i, \mathbf{q})$ over $N(\mathbf{p}'_i)$. Gather these image location pairs $(\mathbf{p}_i, \mathbf{q}_i)$ into a set of correspondences \mathcal{C}_t .
4. Estimate the new parameter set:

$$\hat{\theta}_t = \underset{\theta}{\operatorname{argmin}} \sum_{(\mathbf{p}_i, \mathbf{q}_i) \in \mathcal{C}_t} w_i D(T(\mathbf{p}_i, \theta), \mathbf{q}_i)^2 \quad (2)$$

Here $D(\cdot, \cdot)$ is an alignment error function, and w_i is a robust weight.

5. Compute the alignment error variance, σ_t^2 , and the covariance of $\hat{\theta}_t$.
6. $t = t + 1$
7. Repeat steps 3-6 until convergence.

3.1 Feature Points

Feature points may be extracted from I_m to form \mathcal{P} in several ways. When aligning retinal fundus images, these are automatically-detected centerline points along the retinal vessels [6]. In the lung CT application, the volume is divided into small regions, and the gradient structure of each region is analyzed (eigen analysis) to determine if it is “landmark-like” (significant variation of intensity in all directions) “vessel-like” (significant variation in two directions), “face-like” (significant variation in one direction), or homogeneous. Matching of the feature points is restricted to the directions along which their intensity structure varies. For example, the search for a match for a face-like feature is only along the normal to the face after it has been (approximately) transformed into I_f based on the current transformation estimate. Alignment error constraints, $D(T(\mathbf{p}_i, \theta), \mathbf{q}_i)$, are defined in the same way.

3.2 Uncertainties and Search Regions

The key to the new algorithm is calculating the uncertainty and the resulting search region, $N(\mathbf{p}'_i)$. The first issue is estimating the standard deviation of the alignment error, a difficult problem in robust estimation. We have found that a similarity-weighted average of the squared alignment errors (step 5 of the algorithm) is more effective than the usual robust estimation based on geometric distances. Intuitively, the reason is that the covariance-driven search already

eliminates the possibility of gross outliers, making the similarity score a much better indication of match correctness. The weight for a correspondence¹ is

$$w_i = \frac{S''(\mathbf{p}_i, \mathbf{q}_i)}{1 + S(\mathbf{p}_i, \mathbf{q}_i)}.$$

This favors matches with lower overall similarity error and sharper local minima of the error. The similarity error used throughout the paper is normalized SSD. The scale estimate is $\sigma_t^2 = \sum w_i D(T(\mathbf{p}_i, \boldsymbol{\theta}), \mathbf{q}_i)^2 / \sum w_i$. For the first iteration of matching a prior estimate σ_0 must be used.

Once we have the scale there are two ways to use it in defining the search region. The first is to define an isotropic volume of radius $c\sigma$ around each transformed feature location, \mathbf{p}'_i , in I_f . The second is based on computing the covariance matrix $\Sigma_{\boldsymbol{\theta}}$ of the transformation parameters at each stage of the algorithm (see below). If we compute the Jacobian, \mathbf{J} , of the transformation function $T(\mathbf{p}_i, \boldsymbol{\theta})$ with respect to point location \mathbf{p}_i , then standard covariance propagation techniques show that the covariance matrix of \mathbf{p}'_i , the mapped point, is approximately $\mathbf{J}\Sigma_{\boldsymbol{\theta}}\mathbf{J}^T$. This covariance matrix defines a search volume (error ellipsoid) surrounding \mathbf{p}'_i (in fact, the first way to define the search area approximates the covariance matrix as $\sigma^2\mathbf{I}$). Using either method, the final search region for each feature is then along the feature's search direction (mapped by the transformation estimation, as above) within this volume.

3.3 Robust Parameter Estimation

Parameter estimation in each iteration is based on weighted least-squares using the similarity weights above. Levenberg-Marquardt techniques are needed for estimating the spline parameters. The covariance matrix of the transformation parameters can be estimated approximately [11, Ch. 15] as $\Sigma_{\boldsymbol{\theta}} = \sigma_t^2 \mathbf{H}^{-1}(\hat{\boldsymbol{\theta}}_t)$, where $\mathbf{H}(\hat{\boldsymbol{\theta}}_t)$ is the Hessian matrix, evaluated at the parameter estimate. This explicit calculation is not practical for 3D B-spline deformation models. In ongoing work we are developing approximations.

4 Retinal Image Registration

The first application of the new IMSP algorithm is to use it in place of ICP within our Dual-Bootstrap ICP (DB-ICP) algorithm for retinal image registration [17]. DB-ICP generates a series of low-order initial transformation estimates based on automatic matching of small image regions. It then tests each initial transformation and associated region separately to see if it can be “grown” into an accurate, image-wide quadratic transformation (12 parameters). The growth process iteratively: (1) refines the transformation estimate, (2) expands the region, and (3) tests to see if a higher-order transformation model can be used, with steps (2) and (3) controlled by $\Sigma_{\boldsymbol{\theta}}$. IMSP can be used in place of ICP in

¹ See [16] for a different form of covariance-based weighting.




Fig. 2. An example retinal image alignment where DB-ICP failed but DB-IMSP succeeded. The small overlap region between the images contains little variation in the vessel contour directions, leading to instability in the feature-to-feature alignment but not in feature-to-intensity alignment.

step (1), matching vascular centerline points from the moving image against the intensities of the fixed image using normal distance constraints.

Although DB-ICP is extremely successful (100% success on image pairs that overlap by at least 35% and have at least a minima set of features), DB-IMSP is better. From our large data set we pulled 81 difficult image pairs having low image overlaps and poor contrast. DB-ICP aligned only 33% of these pairs (to a required average alignment error of less than 1.5 pixels on 1024x1024 images), whereas DB-IMSP aligned 86%. By not relying on consistency between extracted features in the two images, it was able to generate more constraints. An example is shown in Figure 2.

5 Spline-Based Lung CT Volume Registration

The second application of IMSP is in B-spline alignment [12] of lung CT volumes. As discussed above, feature set \mathcal{P} is determined by analyzing the distribution of intensity gradients in regularly-spaced blocks in each image. In each iteration, the neighborhood search volume is isotropic, with radius $c\sigma_t$; due to the large number of transformation parameters, the covariance matrix is not explicitly computed. The algorithm proceeds in a coarse-to-fine manner, starting at 64x64 (within each slice) and ending at 256x256. The uniform B-splines are also estimated in a hierarchical fashion, with a knot spacing of 360 millimeters at the 64x64 resolution and 90 millimeters at the 256x256 resolution. For each combination of




Fig. 3. Example registration results using IMSP and a B-spline deformation on serial lung CT images. The left shows checkerboard results (alternating blocks from the two images) in a single slice following alignment. The right shows the affine alignment using the final IMSP correspondences from B-spline registration.

resolution and B-spline knot spacing, the algorithm proceeds as described above. The correspondence constraints are powerful enough, at least in this application, that no additional smoothness constraints must be placed on the B-spline.

In preliminary studies, this algorithm has been applied to align the individual volumes in a 4d CT sequence and serial, intra-patient 3d CT volumes taken 2 to 6 months apart. The 4d CT alignment was accurate to significantly less than a voxel. In the serial alignment problem, using images taken with GE Lightspeed 16 and GE Lightspeed Ultra scanners, the deformations were quite large. The final alignment error for 3 different patients (4 alignments — one patient had 3 volumes) we’ve studied averaged approximately 1.5mm as compared to an affine alignment error of 11.7mm (using the final B-spline correspondences). Therefore, unlike rigid registration [3], the deformable results are approaching the accuracy needed for serial CAD studies. An example side-by-side comparison of the B-spline results and the affine alignment for a single slice is shown in Figure 3.

6 Discussion and Conclusions

The results on retinal and lung CT registration show the promise of the new IMSP algorithm. The computation and use of uncertainty effectively controls the search range for feature-to-intensity matching. The algorithm produces accurate matches while relying less on consistency between feature locations in the two images. The results raise several interesting issues for on-going work. First is an effective combination of ICP-style feature-to-feature correspondences and IMSP-style feature-to-intensity correspondences to improve both reliability and efficiency. Second is using IMSP constraints to build a variable resolution

B-spline model, introducing higher resolutions where the deformations are more substantial and the constraint set richer. Third is using the resulting algorithm for serial registration of lung CT volumes in computer-aided diagnosis.

References

- [1] S. Aylward, J. Jomier, S. Weeks, and E. Bullitt. Registration of vascular images. *IJCV*, 55(2-3):123–138, 2003.
- [2] P. Besl and N. McKay. A method for registration of 3-d shapes. *IEEE T. Pattern Anal.*, 14(2):239–256, 1992.
- [3] M. Betke, H. Hong, D. Thomas, C. Prince, and J. P. Ko. Landmark detection in the chest and registration of lung surfaces with an application to nodule registration. *Med. Image Anal.*, 7:265–281, 2003.
- [4] P. Cachier, E. Bardinet, D. Dormont, X. Pennec, and N. Ayache. Iconic feature based nonrigid registration: the PASHA algorithm. *CVIU*, 89:272–298, 2003.
- [5] P. Cachier, J.-F. Mangin, X. Pennec, D. Rivière, D. Papdopoulos-Orfanos, J. Régis, and N. Ayache. Multisubject non-rigid registration of brain MRI using intensity and geometric features. In *Proc. 4th MICCAI*, pages 734–742, 2001.
- [6] A. Can, H. Shen, J. N. Turner, H. L. Tanenbaum, and B. Roysam. Rapid automated tracing and feature extraction from live high-resolution retinal fundus images using direct exploratory algorithms. *IEEE Trans. on Inf. Tech. in Biomedicine*, 3(2):125–138, 1999.
- [7] J. Feldmar, J. Declerck, G. Malandain, and N. Ayache. Extension of the ICP algorithm to nonrigid intensity-based registration of 3d volumes. *CVIU*, 66(2):193–206, May 1997.
- [8] L. Ibáñez, L. Ng, J. Gee, and S. Aylward. Registration patterns: the generic framework for image registration of the Insight Toolkit. In *IEEE Int. Symp. on Biomedical Imaging*, pages 345–348, Washington, DC, USA, 7–10 July 2002.
- [9] R. A. McLaughlin, J. Hipwell, G. P. Penney, K. Rhode, A. Chung, J. A. Noble, and D. J. Hawkes. Intensity-based registration versus feature-based registration for neurointerventions. In *Proc. Med. Image. Understanding Analysis*, 2001.
- [10] S. Ourselin, A. Rocher, G. Subsol, X. Pennec, and N. Ayache. Reconstructing a 3D structure from serial histological sections. *IVC*, 19:25–31, 2000.
- [11] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1992.
- [12] D. Rueckert, I. Somoda, C. Hayes, D. Hill, M. Leach, and D. Hawkes. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Trans. Med. Imaging.*, 18:712–721, 1999.
- [13] G. Sharp, S. Lee, and D. Wehe. ICP registration using invariant features. *IEEE T. Pattern Anal.*, 24(1):90–102, 2002.
- [14] D. Shen and C. Davatzikos. Hammer: Hierarchical attribute matching mechanism for elastic registration. *IEEE Trans. Med. Imaging.*, 21(11):1421–1439, 2002.
- [15] H. Shum and R. Szeliski. Construction of panoramic image mosaics with global and local alignment. *IJCV*, 36(2):101–130, 2000.
- [16] A. Singh and P. Allen. Image-flow computation: An estimation-theoretic framework and a unified perspective. *CVGIP: Image Understanding*, 56(2):152–177, 1992.
- [17] C. Stewart, C.-L. Tsai, and B. Roysam. The dual-bootstrap iterative closest point algorithm with application to retinal image registration. *IEEE Trans. Med. Imaging.*, 22(11):1379–1394, 2003.