

to appear in *Informational Processing in Medical Imaging (IPMI)*, 2003

# A View-Based Approach to Registration: Theory and Application to Vascular Image Registration

Charles V. Stewart, Chia-Ling Tsai, and Amitha Perera

Rensselaer Polytechnic Institute  
Troy, New York 12180-3590  
*stewart, tsaic, perera@cs.rpi.edu*

**Abstract.** This paper presents an approach to registration centered on the notion of a *view* — a combination of an image resolution, a transformation model, an image region over which the model currently applies, and a set of image primitives from this region. The registration process is divided into three stages: initialization, automatic view generation, and estimation. For a given initial estimate, the latter two alternate until convergence; several initial estimates may be explored. The estimation process uses a novel generalization of the Iterative Closest Point (ICP) technique that simultaneously considers multiple correspondences for each point. View-based registration is applied successfully to alignment of vascular and neuronal images in 2-d and 3-d using similarity, affine, and quadratic transformations.


## 1 Introduction

The view-based approach to registration is motivated by two common problems in medical image analysis, illustrated in the applications of aligning 2-d images of the retina and in aligning 3-d vascular and neuronal images [2] (Figs 1-6):


**Problem 1:** A single landmark correspondence, specified manually or detected automatically, is established between two images. Aligning the images based on this correspondence is reasonably accurate locally in the small region surrounding the landmarks, but quite poor image-wide (Fig 3a). Without requiring more landmark correspondences, is it possible to automatically “grow” an accurate registration from the initial, local alignment?

**Problem 2:** During an image-guided intervention procedure, a substantial jump occurs between two successive images. Possible causes include patient movement (respiration), instrument movement, and time-delays between acquisition of good-quality images. As a result, incremental registration (or tracking) could converge to an incorrect local minimum, especially for images containing repeated structures such as blood vessels or bronchial tubes (Fig. 4).

This paper presents a unified set of techniques to solve these two registration problems. Two primary innovations are described. The first is the introduction of a view-based framework for registration. The second is a generalization of the widely-used ICP algorithm [5, 9] to accommodate feature descriptions and multiple correspondences per point (feature). The main body of the paper describes



**Fig. 1.** Retinal images taken 3.5 years apart of a patient having a branch vein occlusion. Substantial differences in the non-vascular regions are apparent due to hemorrhaging and pathologies.



**Fig. 2.** The vascular features used in registering retinal images [7]. The left shows landmarks detected at the branching and cross-over points of the blood vessels. The right shows the extracted vessel centerline points.

these two innovations in detail and then combines them to form the algorithms solving Problems 1 and 2.

The presentation assumes vascular or neuronal features have already been segmented from the images. Although many different techniques could be used, the algorithms we employ extract elongated structures using two-sided (or multi-sided in 3-d) boundary following [1, 7]. The features are point locations along the vessel centerlines, akin to medial axis points. Each centerline point is described by a location, tangent direction, and width. Branching and cross-over points (“landmarks”), used in initialization, are extracted as well. The centerline points between a pair of landmarks are gathered into a “segment”.


## 2 Views and View Generation

Abstractly, a *view* is a definition of the registration problem together with a current transformation estimate. A view poses a problem for an iterative minimization technique, such as Iterative Closest Point (ICP) [5, 9] or Mutual Information [13, 20], to solve and provides the starting estimate for this minimization. A view specification includes an image resolution, a transformation model, a set of image primitives, and a minimization technique. By starting with a simplified view, such as a coarse image resolution or simplified transformation model, aligning the images based on this view, and then proceeding to a more complete and more detailed view, many local minima in the alignment can be avoided. In the view-based approach, registration is addressed as a three-stage process: initialization, view generation, and minimization. The latter two are alternated until convergence. Several different initial estimates may be evaluated.

Thus far, our notion of a view is simply a gathering and synthesis of current techniques that use multiple resolutions [4, 17], hierarchies of transformations models [8, 15], and hierarchies of image primitives [16]. In addition to this synthesis, however, we introduce two novel ideas. First, the view includes the image region over which the transformation estimate is considered accurate (Fig. 3). Second, instead of pre-specifying the transition between views (e.g. when to switch to higher resolutions or to new transformation models), some parts of the view are automatically generated.


Including the image region in our definition of a view allows the registration algorithm to start in a small region of each image (a “key-hole view” — Problem 1), align the images only in this small region (Fig. 3), and gradually increase the region size (a change in view). If these steps are done correctly, this can avoid errors due to initial image-wide misalignment of complicated structures such as networks of blood vessels or neuronal fibers. It does require, however, that initialization in the small region be reasonably accurate. Ensuring this requires manual specification of the initial region, or matching of small-scale features such as landmarks and their immediate surroundings (Fig. 2). Automatic matching can not be flawless, so several different initial matches and resulting initial views must be tested and evaluated based on their final alignment accuracies.

The second novel technique in our view-based approach to registration is automatic generation of views. For now, we only develop methods to generate the region (“region growth”) and the transformation model (“model selection”). The need for automatic region growth and model selection are most easily seen in Problem 1 (Fig. 3). The initial region and perhaps subsequent regions contain too few constraints to reliably estimate the 12 parameters of a quadratic model (Table 1) needed for accurate image-wide alignment [8]. Thus, a lower-order model must be used. The model should be chosen automatically based on the available constraints. A similar model selection problem arises in registering coarse resolution images. Region growth should also be data driven, with unstable estimates causing slow region growth, and stable estimates leading to rapid region growth.



**Fig. 3.** View-based solution to Problem 1. Panel (a) shows an initial alignment of two images (Fig. 1) based on a similarity transformation computed by matching a single landmark and its immediate surroundings from each image (small white rectangular region in the center). Vascular centerline points the two images are shown in black and in white. The global, image-wide alignment is quite poor. Panel (b) shows the alignment after three iterations of view generation (region growth and model selection) and estimation. No change has yet been made in the model, but in the next iteration (c) a reduced quadratic transformation is selected (Table 1). Panel (d) shows the final alignment using a quadratic transformation.

The remainder of this section describes in detail the model selection and region growth techniques of automatic region generation. Both are driven by the uncertainty in transformation estimates. We assume the transformation estimation process produces a covariance matrix and a discrete set of alignment errors. The covariance matrix may be approximated by the inverse Hessian of the objective function used in estimation, evaluated at the estimate. In the description of these steps,  $t$  denotes the iteration of the view selection and estimation loop,  $R_t$  denotes the transformation region, and  $M_t$  denotes the transformation model selected for the current view (Table 1).  $\hat{\theta}_t$  is the vector of estimated parameters of the transformation model, and  $\Sigma_t$  is the associated covariance matrix.



**Fig. 4.** Problem 2 — incremental registration. The left panel shows an initial alignment of two images of a healthy retina. As in Fig. 3, centerline points from the two different images are shown in black and in white. Some of the ICP correspondences are shown with white line segments between black and white points. The obvious mismatches lead to the poor alignment results shown on the right, even when using robust estimation.

Model	Equation	DoF	Accuracy
Similarity	$\mathbf{q} = \begin{pmatrix} \theta_{11} & \theta_{12} & \theta_{13} & 0 & 0 & 0 \\ \theta_{21} & -\theta_{13} & \theta_{12} & 0 & 0 & 0 \end{pmatrix} \mathbf{X}(\mathbf{p} - \mathbf{p}_0)$	4	5.05 pixels
Affine	$\mathbf{q} = \begin{pmatrix} \theta_{11} & \theta_{12} & \theta_{13} & 0 & 0 & 0 \\ \theta_{21} & \theta_{22} & \theta_{23} & 0 & 0 & 0 \end{pmatrix} \mathbf{X}(\mathbf{p} - \mathbf{p}_0)$	6	4.58 pixels
Reduced quadratic	$\mathbf{q} = \begin{pmatrix} \theta_{11} & \theta_{12} & \theta_{13} & \theta_{14} & 0 & \theta_{14} \\ \theta_{21} & -\theta_{13} & \theta_{12} & \theta_{24} & 0 & \theta_{24} \end{pmatrix} \mathbf{X}(\mathbf{p} - \mathbf{p}_0)$	6	2.41 pixels
Quadratic	$\mathbf{q} = \begin{pmatrix} \theta_{11} & \theta_{12} & \theta_{13} & \theta_{14} & \theta_{15} & \theta_{16} \\ \theta_{21} & \theta_{22} & \theta_{23} & \theta_{24} & \theta_{25} & \theta_{26} \end{pmatrix} \mathbf{X}(\mathbf{p} - \mathbf{p}_0)$	12	0.64 pixels

**Table 1.** The set of transformation models used in retinal image registration [8]. To clarify notation in the equations,  $\mathbf{p} = (x, y)^T$  is an image location in  $I_1$ ,  $\mathbf{q} = (u, v)^T$  is the transformed image location in  $I_2$ , and  $\mathbf{p}_0$  is the center of the registration region. In addition to the formulations, the table also shows the degrees of freedom (DoF) in each model and the average alignment error on  $1024 \times 1024$  images.

## 2.1 Region Expansion

For simplicity,  $R_t$  is rectilinear (Fig 3 shows transformed  $R_t$  in white). At each view generation step, each side of  $R_t$  is shifted outward. Let  $\mathbf{p}_c$  be the point in the center of a given side. The covariance matrix of the transformation can be converted to the covariance matrix of the mapping of this point  $\mathbf{p}'_c = M_t(\hat{\boldsymbol{\theta}}_t; \mathbf{p}_c)$  using standard covariance propagation techniques. This yields the transfer error covariance of  $\mathbf{p}_c$  [11, Ch. 4]. Let the component of this  $2 \times 2$  covariance in the outward direction from  $R_t$  be  $\sigma_c^2$ . Point  $\mathbf{p}_c$  is expanded outward in direct proportion to  $\beta / \max(\sigma_c, 1)$ , with the upper bound of 1 preventing growth that

is too fast and  $\beta$  controlling the overall rate of growth. The rectangle bounding the expanded center points from each side is the new region,  $R_{t+1}$ .

## 2.2 Model Selection

Region growth makes new constraints on the transformation estimate available for the next view. Determining if this warrants switching to a higher-order model requires the application of model selection techniques [6, 19]. These techniques trade-off the increased fitting accuracy of higher-order models against the increased stability of lower-order models. The stability is usually measured as a function of the covariance matrix  $\Sigma_t$ . The criteria we use is [6]:

$$\frac{d_t}{2} \log 2\pi - \sum_i w_i r_i^2 + \log \det(\Sigma_t), \quad (1)$$

where  $d$  is the degrees of freedom in the model  $M_t$ ,  $\sum_i w_i r_i^2$  is the sum of the robustly-weighted alignment errors (based on the estimate  $\hat{\theta}_t$ ), and  $\det(\Sigma_t)$  is the determinant of the parameter estimate covariance matrix. This equation is evaluated for the current model  $M_t$  and for other candidate models. For each other model  $M$ , the estimate  $\hat{\theta}_t$  for the current model  $M_t$  serves as the starting point to estimate a new set of parameters (for model  $M$ ). For simplicity, at each iteration  $t$  we usually just evaluate the model  $M$  which is the next more complicated model than  $M_t$ . Overall, the model that results in the greatest value of (1) is chosen as  $M_{t+1}$ .


## 3 Estimation Engine: ICP / IMCF

Although MI or ICP could be used as the estimation engine in view-based registration, here we introduce a generalization of ICP that allows multiple matches per point, and uses a novel, robust weighting technique to combine the match contributions. We'll start with a summary discussion of robust ICP.

Let  $\mathcal{P} = \{\mathbf{p}_i\}$  be the extracted set of blood (neuron) vessel centerline points in image  $I_1$ . Robust ICP (1) takes the point locations from  $\mathcal{P}$  that are in registration region  $R_t$  of the current view, (2) applies the current transformation estimate to each, and (3) finds the closest centerline point location  $\mathbf{q}_i$  in  $I_2$ . This generates correspondence set  $\mathcal{C}_t = \{(\mathbf{p}_i, \mathbf{q}_i)\}$ . The new transformation estimate is computed by minimizing the robust objective function

$$E(\theta_t) = \sum_{(\mathbf{p}_i, \mathbf{q}_i) \in \mathcal{C}_t} \rho(d(M_t(\theta_t; \mathbf{p}_i), \mathbf{q}_i)/\hat{\sigma}). \quad (2)$$

Here,  $M_t(\cdot; \cdot)$  maps  $\mathbf{p}_i$  into  $I_2$ , and  $d(\cdot, \cdot)$  is the distance of the mapped point to  $\mathbf{q}_i$ , measured as the perpendicular distance from  $M_t(\theta_t; \mathbf{p}_i)$  to the vessel contour tangent line through  $\mathbf{q}_i$ .  $\hat{\sigma}$  is an estimate of the standard deviation of the alignment errors which is robust to outliers; and  $\rho(\cdot)$  is a robust loss function.



**Fig. 5.** IMCF. The left panel illustrates the generation of multiple correspondences in IMCF. For transformed point  $\mathbf{p}'$ , the closest centerline point ( $\mathbf{q}_1$ ) is found and additional matches are sought from each centerline segment intersecting a circular region surrounding  $\mathbf{p}'$ . Using correspondences such as these, IMCF is able to correct the incorrect alignment shown in Fig. 4 and produce the alignment shown on the right (white centerline points aren't seen where they exactly overlap the black ones).

Equation 2 is minimized using iterative-reweighted least squares [12], with weight function  $w(u) = \rho'(u)/u$ . We use the Beaton-Tukey biweight function [3] which goes to 0 for errors above about  $4\hat{\sigma}$ , eliminating the effects of mismatches due to missing vessel structures. The ICP steps of matching and transformation estimation are repeated until the estimate converges for the current view.

### 3.1 IMCF

Even with a robust error norm, ICP fails when a significant fraction of the closest points are not correct correspondences (Fig. 4). To overcome this, we allow multiple matches per centerline point and then adjust the minimization to exploit the extra information provided by these matches. We also introduce a similarity measure between matches, based on comparison of feature properties. The resulting algorithm is called Iterative Multiple Close Features (IMCF).

In each matching iteration, IMCF looks for multiple matches for each  $\mathbf{p} \in \mathcal{P} \cap R_t$  (Fig. 5). Let  $\mathbf{p}' = M(\boldsymbol{\theta}; \mathbf{p})$  be the transformation of  $\mathbf{p}$  into  $I_2$ . Just as in ICP, the closest  $I_2$  centerline point to  $\mathbf{p}'$  is found ( $\mathbf{q}_1$  in Fig. 5). In addition, all other vessel segments intersecting a circular region surrounding  $\mathbf{p}'$  are identified. The closest point on each segment is used to form an additional match with  $\mathbf{p}$ , added to  $\mathcal{C}_t$ . The search radius is a small multiple of the alignment error  $\hat{\sigma}$ .

Like other multiple-match techniques ([10]), the influence of these correspondences is controlled through the weight function, but our weight function is novel. As background, in robust ICP the weight values are  $w_i = w(u_i) = \rho'(u_i)/u_i$ , where  $u_i = d(M(\boldsymbol{\theta}, \mathbf{p}_i), \mathbf{q}_i)/\hat{\sigma}$ . Thus, the ICP weighted least-squares objective

function is

$$\sum_{(\mathbf{p}_i, \mathbf{q}_i) \in \mathcal{C}} w_i \cdot (d(M(\boldsymbol{\theta}, \mathbf{p}_i), \mathbf{q}_i) / \hat{\sigma})^2. \quad (3)$$

The steps of weight calculation and updates to  $\boldsymbol{\theta}$  are alternated.

For each match in IMCF, we define three separate weights and then combine them into the cumulative weight,  $w_i$  in (3). The “geometric-error” weight is  $w_{i,r} = \rho'(u_i)/u_i$ , as above. The “signature-error” weight, based on comparison of other properties of the correspondence besides distance, is  $w_{i,s}$ . In aligning vascular images, these properties are vessel orientations and widths. The third weight is called the competitive weight. To define it for match  $(\mathbf{p}_i, \mathbf{q}_i)$  we need all other matches for feature  $\mathbf{p}_i$ . This is the match subset  $\mathcal{C}_{t,i} = \{(\mathbf{p}_j, \mathbf{q}_j) \in \mathcal{C}_t \mid \mathbf{p}_i = \mathbf{p}_j\}$ . Using this subset the competitive weight is

$$w_{i,c} = \frac{w_{i,r}w_{i,s}}{\sum_{(\mathbf{p}_j, \mathbf{q}_j) \in \mathcal{C}_{t,i}} w_{i,r}w_{i,s}}.$$

Intuitively,  $w_{i,c}$  is the fraction of the total matching weight for feature  $\mathbf{p}_i$  that accrues to this match. The final weight is  $w_i = w_{i,r}w_{i,s}w_{i,c}$ .

The main advantage of this technique is simplicity. No separate “outlier process” and associated tuning parameters are needed [10]. The use of competitive weights can not increase the final weight (over the combined geometric and similarity weights), so outliers remain outliers. In addition, weights of unambiguous matches remain unchanged. Hence, the competitive weight is mostly intended to downgrade the influence of ambiguously matched features, allowing less ambiguous matches to have greater influence on the alignment.

As a final comment, the robust, competitive weighting scheme just described requires accurate scale estimation. As IMCF iterates, this is the robustly-weighted rms  $\hat{\sigma}^2 = \sum_i w_i \cdot d(M(\hat{\boldsymbol{\theta}}; \mathbf{p}_i), \mathbf{q}_i)^2 / \sum_i w_i$ . Initially, however, it is a robust estimate called MUSE [14] based on the matches that are most similar (orientation and width) rather than closest.

## 4 Two Algorithms

Given the foregoing descriptions of the view-based approach, automatic view generation, and the ICP and IMCF algorithms, the descriptions of the solutions to Problems 1 and 2 are straightforward. In fact, at a software-level there is very little difference between the implementations (one of the goals of the work). For each algorithm, we briefly describe the initialization, view generation, and estimation stages. Results are presented in the next section.

For Problem 1, initial estimates and initial views are generated by placing one automatically detected landmark from each image in correspondence and computing an initial transformation by aligning the immediately surrounding vessels (Fig. 3). The correspondences may be generated manually, or they may be generated automatically by matching signatures computed from the widths and orientations of vessels meeting at the landmark. Multiple initial estimates




are generated in the latter case. The view includes the region and the transformation model. These are grown (bootstrapped) from the small starting regions. For estimation, we’ve used both ICP and IMCF. The more sophisticated IMCF algorithm does not substantially improve the results because the region growth and model selection reduce matching ambiguities; in practice we use ICP. A transformation is accepted as correct when the algorithm converges to an estimate having an image-wide, median error in the alignment of centerline pixels less than an empirical threshold (1.5 pixels in retina images [8]); when no initial estimate leads to such an alignment, the algorithm indicates that the images can’t be registered. The overall algorithm is called the “Dual-Bootstrap ICP” (see [18] for more details).

For Problem 2 (Fig. 4) the initial estimate is just the estimate for the previous image in the sequence. Alignment is against an image giving a map of the surgical region. The view includes the image resolution and the transformation model. Resolution changes follow a standard coarse-to-fine progression, switching image resolutions when estimation converges at a given each resolution. Transformation models are determined automatically for each resolution. IMCF is used to estimate the transformation once the view (resolution and model) is fixed. The overall algorithm is called M-IMCF (Multiresolution IMCF).

## 5 Results

The algorithms have been applied in four contexts: (1) registering pairs of 2-d retinal images, (2) incremental registration of sequences of retinal images with each registration starting from the previous estimate, (3) registering confocal stacks of neuronal images, and (4) registering confocal stacks of vascular (rat brain) images. Dual-Bootstrap ICP is used for (1), (3) and (4), which correspond to Problem 1 from the introduction, while M-ICF is used for (2). In the retinal registration problems, the final transformation model is a 12-parameter quadratic (Table 1), and the experimental results are extensive. For the confocal images, the final transformation is affine, and the results are preliminary.


Example alignments using Dual-Bootstrap ICP are shown in Fig. 3 for 2-d retinal images and in Fig. 6 for 3-d vascular images. In retinal image registration, tests were applied to over 4000 images pairs taken from healthy eyes and taken from diseased eyes (with some inter-image time intervals of several years). For image pairs that (a) overlapped by at least 35% of the image area and (b) had at least one common landmark detected by feature extraction, the algorithm always aligned the images to within a 1.5 pixel error threshold ( $1024 \times 1024$  images). Performance gradually degraded with less overlap. Moreover, no incorrect alignments were accepted as correct (i.e. none aligned to within the error threshold). The median number of initial estimates tried was 2, and the average registration time was about 2 seconds. Registration in 3-d takes about a minute, with most time occupied by feature extraction prior to registration. The pixel dimensions were  $0.7 \mu m$  within each plane and  $5 \mu m$  between planes. Median alignment errors were less than a pixel in each dimension.



**Fig. 6.** The Dual-Bootstrap ICP registration algorithm for 3-d vascular images of a rat brain taken using confocal microscopy. Panel (a) shows an max-intensity projection, with super-imposed 3-d vessel centerlines. Panels (b)-(d) show alignment of centerlines (solid lines for one image, dashed lines for the other) taken at two different times. Panel (b) gives the initial alignment in a small region. Panels (c) and (d) show intermediate and final alignments. Even though the figures are 2-d projections, the registration occurs in 3-d.

For M-IMCF, our primary experimental goal is evaluating the domain of convergence of initialized registration. We explore this with synthetic shifts of real image pairs, using a subset of the image pairs used to study Dual-Bootstrap ICP. For each pair, we have the “correct” (validated) quadratic transformation. We perturb this transformation by different translations and rotations to provide initial offsets, run the algorithm for each, and determine if it converges to within 1.5 pixels of the validated transformation. We count the initial transformation as a success if it does. Fig. 7 summarizes the success rates for four algorithms: ICP, Multiresolution ICP (M-ICP), IMCF, and M-ICMF. Clearly, both multiresolution and IMCF are important in increasing the domain of convergence.

To summarize the experimental results, the Dual-Bootstrap ICP algorithm overall is extremely successful at registering retinal image pairs, including pairs with relatively low overlap. M-IMCF works well for initialized registration, substantially outperforming ICP and multiresolution ICP. The difference between



**Fig. 7.** Comparing the domains of convergence of ICP, Multiresolution ICP, IMCF, and Multiresolution IMCF on a selection of image pairs of both healthy and unhealthy eyes. The horizontal axis is the radial initial offset from the correct transformation in pixels (all shift directions are combined) from the true alignment. All tested images are  $1024 \times 1024$ . The left plot shows results for no rotation, while the right plot shows results for  $10^\circ$  rotation.

Dual-Bootstrap ICP and M-IMCF in an application context is that the much faster M-IMCF algorithm will be used primarily for on-line registration during surgical procedures. Finally, the preliminary 3-d results demonstrate the applicability of the view-based approach to other, related problems.

## 6 Summary and Conclusions

This paper has introduced a view-based approach to registration. Each view includes an image resolution, a set of image primitives, a transformation model, an image region, a current transformation estimate, and an estimation technique. While the estimation technique could be a standard technique such as ICP or Mutual Information, this paper has introduced a new core estimation technique called IMCF, a generalization of ICP that allows multiple matches to be simultaneously considered. These techniques were used to solve two problems in retinal image registration, and the algorithms have been extended to 3-d registration of vascular and neuronal images as well. A C++ registration library is being built around the view-based approach.

Ongoing work is extending this research in a number of directions. Clearly, the fundamental theoretical extension to the view-based approach is incorporating deformable registration. In doing this the algorithm must start with small scale deformations for simple views and allow increasingly large-scale deformations as the view expands (especially in the image region). In addition to the theory, a variety of applications of view-based registration are being investigated. Currently, the most important of these is the use of Dual-Bootstrap ICP and M-IMCF in the diagnosis and treatment of retinal diseases.

## References

- [1] K. Al-Kofahi, S. Lasek, S. D., C. Pace, G. Nagy, J. N. Turner, and B. Roysam. Rapid automated three-dimensional tracing of neurons from confocal image stacks. *IEEE Trans. on Inf. Tech. in Biomedicine*, 6(2):171–187, 2002.
- [2] S. Aylward, J. Jomier, S. Weeks, and E. Bullitt. Registration of vascular images. *Int. J. of Computer Vision*, (to appear) 2003.
- [3] A. E. Beaton and J. W. Tukey. The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data. *Technometrics*, 16:147–185, 1974.
- [4] J. Bergen, P. Anandan, K. Hanna, and R. Hingorani. Hierarchical model-based motion estimation. In *Proc. 2nd ECCV*, pages 237–252, 1992.
- [5] P. Besl and N. McKay. A method for registration of 3-d shapes. *IEEE Trans. on PAMI*, 14(2):239–256, 1992.
- [6] K. Bubna and C. V. Stewart. Model selection techniques and merging rules for range data segmentation algorithms. *Computer Vision and Image Understanding*, 80:215–245, 2000.
- [7] 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 Info. Tech. for Biomedicine*, 3(2):125–138, 1999.
- [8] A. Can, C. Stewart, B. Roysam, and H. Tanenbaum. A feature-based, robust, hierarchical algorithm for registering pairs of images of the curved human retina. *IEEE Trans. on PAMI*, 24(3):347–364, 2002.
- [9] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. *Image and Vision Computing*, 10(3):145–155, 1992.
- [10] H. Chui and A. Rangarajan. A new algorithm for non-rigid point matching. In *Proc. CVPR*, pages II:44–51, 2000.
- [11] R. Hartley and A. Zisserman. *Multiple View Geometry*. Cambridge University Press, 2000.
- [12] P. W. Holland and R. E. Welsch. Robust regression using iteratively reweighted least-squares. *Commun. Statist.-Theor. Meth.*, A6:813–827, 1977.
- [13] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality image registration by maximization of mutual information. *IEEE Trans. on Medical Imaging*, 16(2):187–198, 1997.
- [14] J. V. Miller. *Regression-Base Surface Reconstruction: Coping with Noise, Outliers, and Discontinuities*. PhD thesis, Rensselaer Polytechnic Institute, 1997.
- [15] H. Sawhney, S. Hsu, and R. Kumar. Robust video mosaicing through topology inference and local to global alignment. In *Proc. 5th ECCV*, volume II, pages 103–119, 1998.
- [16] G. Sharp, S. Lee, and D. Wehe. ICP registration using invariant features. *IEEE Trans. on PAMI*, 24(1):90–102, 2002.
- [17] D. Shen and C. Davatzikos. HAMMER: Hierarchical attribute matching mechanism for elastic registration. *IEEE Trans. on Medical Imaging*, 21(11):1421–1439, 2002.
- [18] C. Stewart, C.-L. Tsai, and B. Roysam. The dual bootstrap iterative closest point algorithm with application to retinal image registration. Technical Report RPI-CS-TR 02-9, Department of Computer Science, 2002.
- [19] P. Torr. An assessment of information criteria for motion model selection. In *Proc. CVPR*, pages 47–52, 1997.
- [20] P. Viola and W. M. Wells III. Alignment by maximization of mutual information. *Int. J. of Computer Vision*, 24(2):137–154, 1997.