Covariance-Driven Mosaic Formation from Sparsely-Overlapping Image Sets with Application to Retinal Image Mosaicing *

Gehua Yang Charles V. Stewart
Rensselaer Polytechnic Insitute
Troy, NY 12180-3590
{yangg2,stewart}@cs.rpi.edu

Abstract

A new technique is presented for mosaicing sparsely-overlapping image sets, with a target application of assisting the diagnosis and treatment of retinal diseases. The geometric image transformations required to construct the mosaics are estimated by (1) estimating the transformations between as many pairs of images as possible, (2) extracting sets of constraints (correspondences) from the successfully registered image pairs, and (3) using these constraint sets to simultaneously (jointly) estimate the final transformations. Unfortunately, this may not be sufficient to construct seamless mosaics when two images overlap but can not be successfully registered (step 1). This paper presents a new method to generate constraints between such image pairs, and use these constraints to estimate a more consistent set of transformations. For each pair, transformation parameter covariance matrices are computed and used to estimate the mapping error covariance matrices for individual features from one image. These features are matched in the second image by minimizing the resulting Mahalanobis distance. The generated correspondences are validated using robust estimation techniques and used to refine the estimates. The steps of covariance computation, matching, and transform estimation are repeated for all relevant image pairs until the final alignment converges. Results are presented and evaluated for several difficult image sets to illustrate the efficacy of the techniques.

1. Introduction

This paper addresses the problem of forming an image mosaic when the image set is sparse and the inter-image overlap can be quite low. The problem arises in the context of imaging the human retina (Figure 1). Mosaics must be constructed to build complete, non-redundant views of the retina as an aid to diagnosing a variety of blindness-causing diseases such as diabetic retinopathy and age-related macular degeneration. These diseases affect a large and growing percentage of the population, yet early detection and treatment, especially of diabetic retinopathy, can significantly improve long-term visual health [7]. For a variety of reasons, including cost, patient comfort and patient volume, images of the retina, acquired using a fundus camera, must be taken as quickly as possible and with as few shots as possible [6]. Combining this requirement with the need for a wide field of view in the constructed mosaic shows the importance of handling low overlaps between images.

The mosaicing problem has received a significant amount of attention both in the computer vision literature [2, 9, 12, 15, 16] and the ophthalmology literature [5, 19]. The most advanced techniques have focused on mosaicing video sequences, where the overlap between images tends to be relatively high. Some techniques have addressed the question of low image overlap, using the global inferencing of topological relationship between images to generate additional inter-image constraints [14].

The major steps of mosaic construction include registering pairs of image, extracting constraints from the registered pairs, and jointly estimating the transformations aligning all images using the extracted constraints. (A final step of constructing the actual mosaic by transforming and blending the images is not discussed here.) Let $\{I_0, \ldots, I_N\}$ be a set of images, and let A be an algorithm for registering pairs of images. We assume $A(I_i, I_j)$ produces two types of output if it successfully aligns I_i to I_j: $\hat{\mathbf{\Theta}}_{i,j}$ is the set of estimated transformation parameter values, and $\mathcal{C}_{i,j}$ is a set of inter-image constraints — corresponding points. For simplicity in this discussion, we assume A is applied to all possible image pairs. The result can be viewed as a connectivity / topology graph: each image forms a vertex, and an edge is formed between vertices i and j when $A(I_i, I_j)$ is successful (Figure 1). The final step, which we will call the "joint alignment" step, is to simultaneously estimate the final transformations using all the constraint sets $\mathcal{C}_{i,j}$[3, 16].

This technique will produce a final set of transformations for all images if the topology graph is connected. This does not ensure accurate alignment, however, especially when

*This work was supported by National Science Foundation Experimental Partnerships grant EIA-0000417.

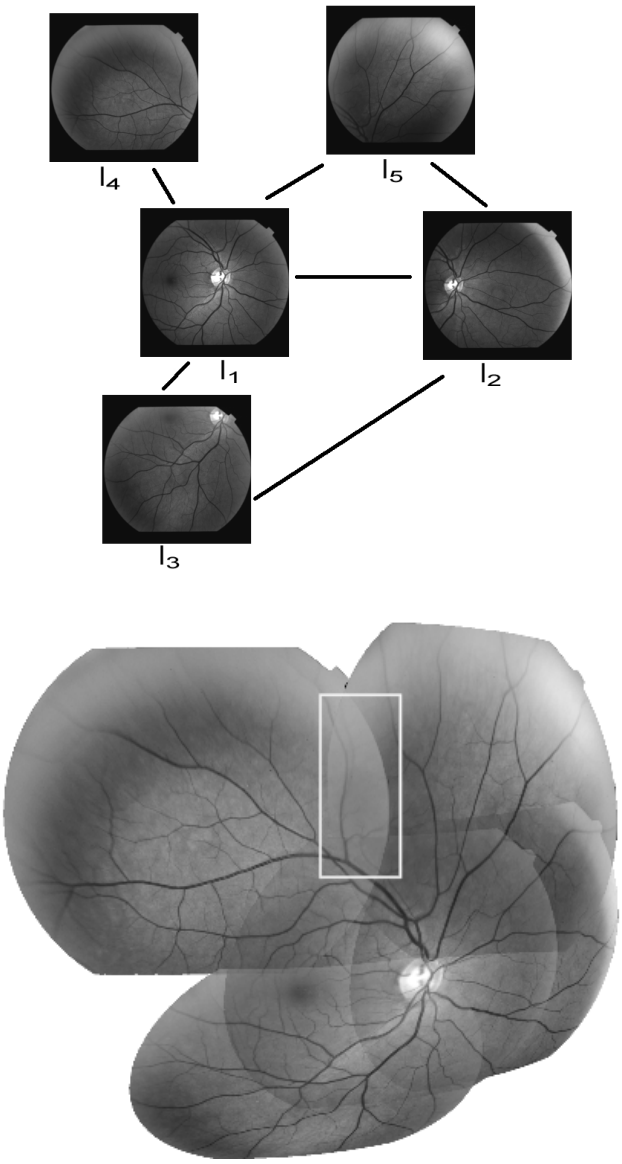

Figure 1: Building a mosaic of retinal images. The top shows a graph of retina images with edges between vertices if the associated images could be aligned via pairwise registration. Notice that image 4 and image 5 overlap (initial overlap 5.8%) but are not adjacent in the graph. This leads to inaccuracies in the mosaic (bottom), highlighted by the rectangular area. Figure 2 shows the rectangular region in more detail and illustrates the solution to the misalignment problem.

the graph is sparsely connected and there is relatively little overlap between some images. Two images, such as I_4 and I_5 of Figure 1, may overlap without there being enough information for pairwise registration to succeed. Constraints on the transformations from nearby images may not be enough to force consistency between these images during joint alignment. The result can be clear misregistration.[1]

The key problem addressed in this paper is how to generate additional constraints, beyond the results of pairwise registration, in order to avoid misregistrations in the final, joint alignment. Sometimes the misalignment may be small, which would allow a technique like incremental block matching [16] to succeed. At other times, the misalignment may be much greater. We handle this by using the alignment uncertainty to drive the generation of additional constraints. The covariance matrices of the transformation parameters estimated in the joint alignment are computed — even for transformations between image pairs that did not register pairwise. These are used to compute error covariances on the inter-image mapping of feature points, which guides the search for new correspondences. These correspondences are used in turn to refine the joint estimate. The effect of these techniques is illustrated in Figure 2. In short, we are using transformation parameter covariances to drive the generation of constraints between images even when pairwise registration fails. These additional constraints are not enough to make pairwise registration succeed, but they are enough to improve the results of joint alignment and eliminate most misregistrations in the final mosaic.

2. Background

In this paper, the pairwise registration algorithm A is treated as a black box. The actual algorithm [17] works automatically, either producing an accurate alignment between image pairs, or indicating that they can not be aligned. Experimental validation on an extensive data set has shown in [18] that this pairwise registration algorithm does *NOT* produces any incorrect alignment.

In the retina application, the alignment of two-dimensional images uses a 12-parameter quadratic transformation model [4, 17]. Given two images, I_m and I_n, let $\mathbf{p} = (x, y)^T$ be a pixel location in I_m, and $\mathbf{p}' = (x', y')^T$ be the transformed location in I_n. Define $\mathbf{X}(\mathbf{p}) = (x^2, xy, y^2, x, y, 1)^T$. Then the transformation of $\mathbf{p}$ onto $\mathbf{p}'$ is

$$\mathbf{p}' = \mathbf{T}(\mathbf{p}; \mathbf{\Theta}_{m,n}) = \mathbf{\Theta}_{m,n}\mathbf{X}(\mathbf{p}), \tag{1}$$

where $\mathbf{\Theta}_{m,n}$ is a 2×6 parameter matrix.[2] This model is accurate to less than a pixel on 1024×1024 retinal images.

The pairwise algorithm generates constraints based on the location of blood vessel landmarks — branching and

[1] We stress that the problem is **not** with pairwise registration. There aren't sufficient constraints in the images shown in Figure 1 for estimation of the 12-parameter transformation necessary to align retinal images.

[2] Abusing notation for the sake of convenience, in other contexts we may interpret $\mathbf{\Theta}_{m,n}$ as a 12×1 parameter vector as well.

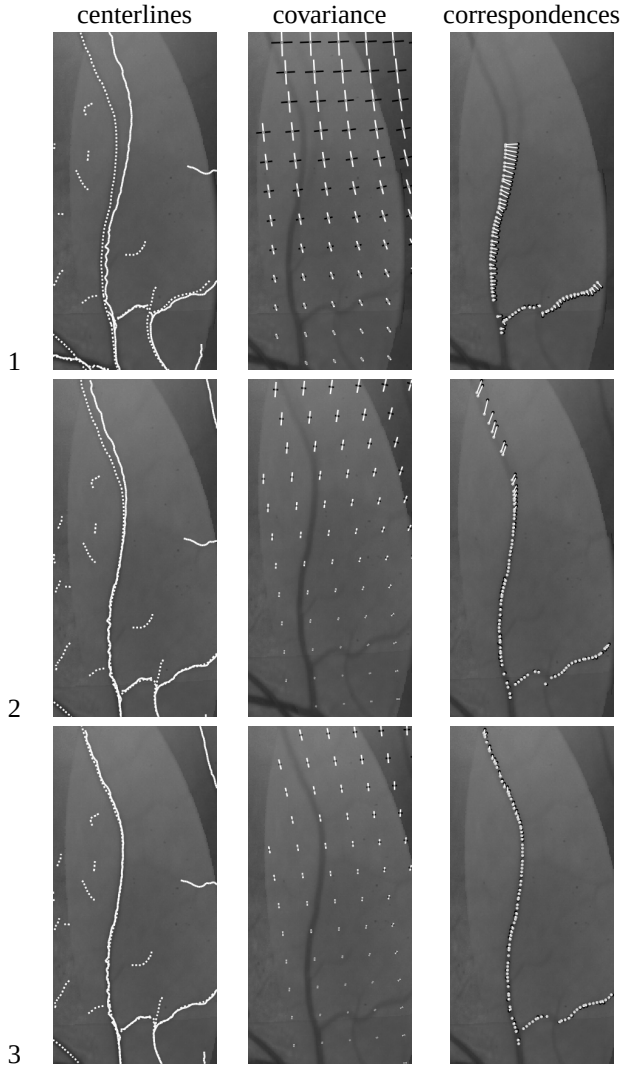

Figure 2: Three iterations of the covariance-driven process of adding constraints and refining the transformation estimates, shown for the rectangular region highlighted in Figure 1. The left column shows the centerline points of I_5 (solid segments) mapped onto I_4 and its centerline points (dotted segments). The misalignment at the top of the region is most prominent, especially in the first iteration. The second column shows sampled transfer error covariance matrices and their eigenvectors (scaled in proportion to their eigenvalues). The third column shows new generated correspondences. Clearly, as the iterations proceed, the alignment, the covariance matrices and the matching error improve dramatically.

cross-over points of the retinal vasculature — and, more importantly, points along the centerlines of the vasculature [4]. Vessels are used because they are prominent, static and easily detectable. (The appearance of the non-vascular background tends to vary with illumination, and background structures such as waste deposits and pathologies tend to change over time.) The centerline points are discrete samples along fairly straight contours, so that precisely matching and aligning them is unrealistic (the aperture problem for registration). Instead, the constraints for registration are alignment of a centerline point from one image with a linear approximation to the centerline contour in the other. Algebraically, if $\mathbf{p}_{m,i}$ is a centerline point location in image I_m and if $\mathbf{p}_{n,j}$ is a centerline point location in image I_n with local contour normal $\hat{\boldsymbol{\eta}}_{n,j}$, then the "normal distance" error of the transformation is

$$(\boldsymbol{\Theta}_{m,n}\mathbf{X}(\mathbf{p}_{m,i}) - \mathbf{p}_{n,j})^T \hat{\boldsymbol{\eta}}_{n,j}. \tag{2}$$

This measures the signed distance of the mapping of $\mathbf{p}_{m,i}$ to the line through $\mathbf{p}_{n,j}$ with normal $\hat{\boldsymbol{\eta}}_{n,j}$.

Based on this discussion, we can now define the correspondence sets that emerge from pairwise registration. The set of image pairs for which A succeeds is $\mathcal{P} = \{(m,n)\}$. Here, order matters because pairwise registration may in rare cases succeed from I_m to I_n, but not vice-versa. For each $(m,n) \in \mathcal{P}$, the constraint set is

$$\mathcal{C}_{m,n} = \{(\mathbf{p}_{m,i}, \hat{\boldsymbol{\eta}}_{m,i}; \mathbf{p}_{n,j}, \hat{\boldsymbol{\eta}}_{n,j})\}. \tag{3}$$

Finally, each constraint in each constraint set has a weight associated with it, $w_{m,n;i,j}$, computed using robust M-estimator weighting. It is important to note that this weight is the usual robust weight multiplied by $1/\sigma^2_{m,n}$, where $\sigma^2_{m,n}$ is the robustly computed variance of the pairwise alignment error between I_m and I_n. This normalizes the weights appropriately for different image pairs.

3. Joint Alignment

The first consideration in developing our technique is to present the joint alignment estimation equations. By specially designating one image I_0 as the "anchor image" on which to construct the mosaic (e.g. I_4 in Figure 1), our goal is to estimate the transformations $\boldsymbol{\Theta}_{1,0}, \ldots, \boldsymbol{\Theta}_{N-1,0}$ of the remaining $N-1$ images onto this anchor using the set of constraint sets $\{\mathcal{C}_{m,n}\}$. This generates $N-1$ of the $N(N-1)$ interimage transformations. We can choose each of the other images as the anchor in turn in order to estimate the remaining interimage transformations. This capability is important to the constraint generation procedure described below.

For a designated anchor, we divide the set of image pairs $\mathcal{P}$ in two: $\mathcal{P}_D = \{(m,0) \mid (m,0) \in \mathcal{P}\}$ and $\mathcal{P}_I = \{(m,n) \mid m \neq 0, n \neq 0, (m,n) \in \mathcal{P}\}$. In words, $\mathcal{P}_D$ is the set of pairs involving the anchor directly ("direct constraints"), while $\mathcal{P}_I$ is the set of pairs that do not involve the

anchor image ("indirect constraints"). These two constraint sets are treated separately in the joint alignment.

We have two choices in computing the joint alignment. The first, used in our previous work in matching vascular landmarks [3], ignores the normal directions in order to define a least-squares error norm based on Euclidean feature-point distances:

$$E(\mathbf{\Theta}_{1,0},\ldots,\mathbf{\Theta}_{N,0}) = \sum_{(m,0)\in\mathcal{P}_D}\sum_{(i,j)\in\mathcal{C}_{m,0}} w_{m,0;i,j}\|\mathbf{\Theta}_{m,0}\mathbf{X}(\mathbf{p}_{m,i}) - \mathbf{p}_{0,j}\|^2 + \sum_{(m,n)\in\mathcal{P}_I}\sum_{(i,j)\in\mathcal{C}_{m,n}} w_{m,n;i,j}\|\mathbf{\Theta}_{m,0}\mathbf{X}(\mathbf{p}_{m,i}) - \mathbf{\Theta}_{n,0}\mathbf{X}(\mathbf{p}_{n,j})\|^2 \quad (4)$$

The first term measures errors against feature locations in the anchor image, while the second term measures inconsistencies in the mapping of corresponding, but non-anchor features. This is quadratic in the transformation parameters and therefore can be solved non-iteratively.

The second choice, introduced here, defines a weighted least-squares error norm based on normal distances:

$$E_{\hat{\boldsymbol{\eta}}}(\mathbf{\Theta}_{1,0},\ldots,\mathbf{\Theta}_{N,0}) = \sum_{(m,0)\in\mathcal{P}_D}\sum_{(i,j)\in\mathcal{C}_{m,0}} w_{m,0;i,j}[(\mathbf{\Theta}_{m,0}\mathbf{X}(\mathbf{p}_{m,i}) - \mathbf{p}_{0,j})^T\hat{\boldsymbol{\eta}}_{0,j}]^2 + \sum_{(m,n)\in\mathcal{P}_I}\sum_{(i,j)\in\mathcal{C}_{m,n}} w_{m,n;i,j}[(\mathbf{\Theta}_{m,0}\mathbf{X}(\mathbf{p}_{m,i}) - \mathbf{\Theta}_{n,0}\mathbf{X}(\mathbf{p}_{n,j}))^T\hat{\boldsymbol{\eta}}'_{n,j}]^2. \quad (5)$$

The issue here is that the transformation of each centerline point normal direction, $\hat{\boldsymbol{\eta}}'_{n,j}$, depends on the transformation parameters. Thus, we no longer have a quadratic estimation problem and must resort to an iterative technique. We initialize the parameters using Euclidean distances as in (4). Then we alternate (a) estimating the normals using the Jacobian of the transformations and (b) re-computing the estimates using (5) and fixed normals. This converges quickly.

The algorithm in Can et al. [3] stops at this point. The new technique described in this paper then generates covariance matrices of all transformation parameter estimates and uses them to guide the generation of new constraints.

3.1. Covariance Matrices

The last step of our joint alignment estimation algorithm for a fixed set of constraints is to compute the covariance matrix of the estimated parameters. We can obtain an approximate covariance by inverting the Hessian matrix of (5) evaluated at the estimate [13, Ch. 15] [11, Ch. 7]:

$$\mathbf{\Sigma}_{\mathbf{\Theta}} = \mathbf{H}^{-1}(E_{\hat{\boldsymbol{\eta}}}(\hat{\mathbf{\Theta}})). \quad (6)$$

Normally, this would be multiplied by the variance of the alignment errors, but these values are already factored into the robust weights used in (5). The result is a $12N \times 12N$ matrix. It can be shown that the individual covariance matrices $\mathbf{\Sigma}_{i,0}$ are obtained by simply extracting the appropriate 12×12 subblock.

Obtaining this covariance matrix is the major reason why the more complicated objective function in (5) is used instead of (4): it gives a more reliable indication of the uncertainty. To demonstrate this intuitively, consider for example the alignment of two images with two parallel lines each, and a fixed set of correspondences along these lines. Using the Euclidean distance in (4), a shift of the transformation along the lines would increase the error proportionally, even though the lines themselves would still be registered. Thus, the transformation would appear much more stable that it truly is. Using (5) would correctly make this transformation appear unstable. This is crucial for correctly guiding the matching process.

4. Generating New Constraints

Our next step, and the most important innovation of the paper, is to use the joint alignment transformation estimates and covariance matrices to generate new correspondences. In doing so, we consider any pair of images I_m and I_n that overlap (based on the joint alignment) or nearly overlap and that were not aligned successfully by the pairwise registration algorithm. These image pairs generally have low overlap — in practice overlap is always lower than 35% of the image area and is generally much lower. We do not consider adding constraints for the image pairs that were aligned by pairwise registration because the pairwise constraints are always sufficient for accurate results in the joint alignment step.

Once the unregistered pairs are identified, each pair is tested in turn for the generation of new constraints. To simplify the process, pair (m, n) is only tested when I_n is the anchor image, i.e. $n = 0$. All pairs are eventually considered by varying the choice of anchor image. It is important to note that we obtain a covariance matrix for the mapping of I_m onto I_0 even though there is no pairwise result. This covariance depends on indirect constraints (see Equation 5).

We infer new constraints relative to the anchor, iteratively switching which image is chosen as the anchor, but always using all constraints in refinement. For pair $(m, 0)$, the following steps are repeated until the estimate converges

1. Initialize an empty constraint set $\mathcal{C}_{m,0}$.
2. Identify the centerline points in both images that fall in or near the apparent overlap region based on the transformation estimate $\hat{\mathbf{\Theta}}_{m,0}$. Denote these as $\mathcal{P}_m$ and $\mathcal{P}_0$, respectively (suppressing the dependence of each set on the other image in the notation).

3. For each centerline point $\mathbf{p}_{m,i} \in \mathcal{P}_m$:

   (a) Map $\mathbf{p}_{m,i}$ onto I_0 to compute location

   $$\hat{\mathbf{p}}'_{m,i} = \hat{\mathbf{\Theta}}_{m,0}\mathbf{X}(\mathbf{p}_{m,i})$$

   (b) Compute the uncertainty covariance matrix of this mapping $\mathbf{\Sigma}_{\hat{\mathbf{p}}'_{m,i}}$ (Figure 2).

   (c) Find the centerline point $\mathbf{p}_{0,j}$ from I_0 minimizing the square Mahalanobis distance to this transformed location. In particular,

   $$\mathbf{p}_{0,j} = \underset{\mathbf{p}\in\mathcal{P}_0}{\operatorname{argmin}} (\hat{\mathbf{p}}'_{m,i} - \mathbf{p})^T \mathbf{\Sigma}^{-1}_{\hat{\mathbf{p}}'_{m,i}} (\hat{\mathbf{p}}'_{m,i} - \mathbf{p}) \tag{7}$$

   (d) Add the constraint $(\mathbf{p}_{m,i}, \hat{\boldsymbol{\eta}}_{m,i}; \mathbf{p}_{0,j}, \hat{\boldsymbol{\eta}}_{0,j})$ to the constraint set $\mathcal{C}_{m,0}$ (Figure 2).

4. Estimate the transformation error scale $\sigma_{m,0}$ and compute the weights $w_{m,i;0;j}$ for the new correspondence set.

5. Temporarily add the image pair $(m, 0)$ to the set of direct constraint pairs, $\mathcal{P}_D$, and re-estimate the joint alignment transformations and covariances as described in Section 3 using all direct and indirect constraint sets.

After convergence, a test is made to verify the consistency of the final constraint set $\mathcal{C}_{m,0}$. If it passes, image pair $(m, 0)$ is added to the set of image pairs, $\mathcal{P}$, and $\mathcal{C}_{m,0}$ is retained. Thus, constraints between images I_m and I_0 are added, even though pairwise registration did not succeed.

The rest of this section describes a few of the steps in more detail.

4.1. Mapping Error and Matching

The uncertainty in point location mapping is computed from the covariance of the transformation parameter estimate. This uses the forward transfer error [8, Ch. 4]. The mapped point $\hat{\mathbf{p}}'_{m,i} = \hat{\mathbf{\Theta}}_{m,i}\mathbf{X}(\mathbf{p}_{m,i})$ is a random variable because it relies on the transformation, which is also a random variable [3] Its covariance matrix can be approximated from $\mathbf{\Sigma}_{\hat{\mathbf{\Theta}}_{m,0}}$ and the Jacobian, $\mathbf{J}$, of the transformation evaluated at $\mathbf{p}_{m,i}$:

$$\mathbf{\Sigma}_{\hat{\mathbf{p}}'_{m,i}} = \mathbf{J}\mathbf{\Sigma}_{m,0}\mathbf{J}^T \tag{8}$$

Examples of these transfer error covariance matrices are shown in Figure 1. Notice how these change spatially and, as expected, have their major uncertainty axis along the vessel directions.

[3] For simplicity, we do not treat $\mathbf{p}_{m,i}$ as a random variable because feature location error is generally much small than errors in the transformation.

When searching for the correspondence, minimizing the Mahalanobis distance for mapped point $\hat{\mathbf{p}}'_{m,i}$ is more challenging than minimizing Euclidean distance. This minimization is accomplished in two steps. First, with the centerline points organized into a spatial data structure, our algorithm first gathers all possible I_0 centerline points in a square region surrounding $\hat{\mathbf{p}}'_{m,i}$ whose width is determined by the maximum eigenvalue of $\mathbf{\Sigma}_{\hat{\mathbf{p}}'_{m,i}}$. Second, the correspondence is found by exhaustively searching this (generally small) set for the centerline point minimizing the Mahalanobis distance.

4.2. Scale Estimation and Weight Calculation

Denote by $r^2_{m,0;i,j}$ the square Mahalanobis distance between $\hat{\mathbf{p}}'_{m,i} = \hat{\mathbf{\Theta}}_{m,0}\mathbf{X}(\mathbf{p}_{m,i})$ and its matching point $\mathbf{p}_{0,j}$. Prior to re-estimating the joint alignment, we use the set of these distances to robustly estimate scale and to weight each individual correspondence. The scale, $\sigma_{m,0}$, is computed using a technique that automatically estimates and adjusts for the fraction of inliers [10]. The weights are computed using standard M-estimator weighting, multiplied as above by the inverse variance:

$$w_{m,0;i,j} = \frac{1}{\sigma^2_{m,0}} \frac{\rho'(r_{m,0;i,j}/\sigma_{m,0})}{r_{m,0;i,j}/\sigma_{m,0}}. \tag{9}$$

The Beaton-Tukey biweight function is used as ρ [1].

4.3. Verification of Constraint Sets

After the iterations of match generation, weight calculation, and transformation re-estimation have converged, the constraint set must still be validated. This decides if image pair $(m, 0)$ will be added to the constraint set. The test is simple. For each generated match, let $d_{m,i;0,j}$ be the normal distance, computed as in Equation 2. Compute the weighted average of these distances:

$$e_{m,0} = \sum w_{m,0;i,j} |d_{m,0;i,j}| \;/\; \sum w_{m,0;i,j}.$$

We call this measure the "Centerline Error Measure" or CEM. The constraint set is verified and retained permanently if $e_{m,0}$ is lower than a threshold, empirically determined in previous work to be 1.5 pixels. CEM, with the same threshold, is used to verify pairwise matching results. Thus, covariances and the Mahalanobis distance are used to "pull in" new correspondences for unregistered image pairs, but the final verification is as stringent as in pairwise registration.

5. Experiments

We have applied this technique to image sets taken from a large patient database of retinal images. We selected a few

Set	images	pairs	successes	failures	Max overlap
1	5	3	1	2	0.4%
2	8	5	3	2	2.8%
3	8	8	6	2	17.8%
4	9	4	2	2	1.1%

Table 1: Summary of the success rate for 4 image data sets. The column labeled "images" indicates the number of image used in the test. The column labeled "pairs" indicates the number of pairs that overlap but do not have a pairwise registration result; these are candidates for the addition of constraints. The column labeled "successes" indicates the number of pairs for which a constraint set was generated and verified. The column labeled "failures" indicates the number of failed pairs. Finally, the last column indicates the fraction of image overlap of the failed pairs.

to present here to illustrate the algorithm. In some cases, we've chosen a subset of the data sets to obtain a sparser set of image. In other cases, such as shown in Figure 3, the original image set has sparse overlap on the periphery of the retina. (Here we only show images from above and to the right of the optic disk.) These are representative of the images being acquired with newer imaging systems.

We can evaluate the results in several ways. The first is a simple visual evaluation of the mosaics constructed with and without the application of the new technique. This is perhaps the most important, but also the most subjective measure. Doing this (see for example Figures 2 and 3), show clear cases of misalignment before application of the new technique and none afterwards.

The other two ways of evaluating the results are numerical. Table 1 presents statistics on the successes and failures of the algorithm. These are evaluated for pairs that overlap but are not aligned by the pairwise algorithm. A success is declared if constraints are generated and verified by the new technique. A failure is declared otherwise. Combining the raw numbers, the algorithm has a 60% success rate on this preliminary test set. One of the eight failures has an image pair overlap of 17.8% of the image area, but examining these images shows no common constraints in the area of overlap. All of the other seven failures have image overlaps less than 3% of the image area. These clearly cause no significant misalignment in the mosaics.

The final numerical evaluation is to study the centerline error measure (CEM). Table 2 shows two important results. First, as expected, the CEM for previously unaligned image pairs improves substantially. Second, the CEM for previously aligned image pairs does not increase significantly. In other words, the addition of the new constraints corrects the bad results without biasing the previously good results.

Overall, these experiments have shown the efficacy of our new techniques both visually and numerically.

set	Pair		Others	
	before	after	before	after
1	5.48	0.93	0.69	0.71
2	1.92	1.30	0.83	0.83
3	1.16	1.10	0.97	0.97
4	0.82	0.48	0.58	0.58

Table 2: Effect of the addition of constraints on the average error in alignment (CEM). The second and third columns show the average CEM before and after application of the new technique for unregistered, but overlapping image pairs. The fourth and fifth columns show the average CEM before and after for registered image pairs. A table of the maximum CEM shows similar trends, although the errors are obviously somewhat higher.

6. Summary and Conclusions

We have presented a technique that constructs mosaics of sparsely-overlapping image sets, and demonstrated this technique on retinal fundus image sets. The primary technical contribution of this paper is the use of uncertainty to drive the formation of new matching constraints between unregistered image pairs in the context of multi-image alignment. This is particularly important for low overlap images and images with sparse feature sets such as the blood vessels of retinal images. We have shown that this technique corrects misalignments in sparsely overlapping image sets without sacrificing overall accuracy.

From the applications perspective this technique is crucial for moving retinal image registration and mosaicing techniques out of the laboratory and into the clinic. Mosaics for most image data sets can be constructed seamlessly without using the new techniques; a few can not. These few, however, can ruin the utility of a system. By adding the diagnostic and refinement capabilities of the covariance-driven techniques described here, we have developed a reliable, autonomous retinal image registration and mosaicing system.

References

[1] A. E. Beaton and J. W. Tukey. The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data. *Technometrics*, 16:147–185, 1974.

[2] M. Brown and D. Lowe. Recognising panoramas. In *Proc. ICCV*, 2003.

[3] A. Can, C. Stewart, B. Roysam, and H. Tanenbaum. A feature-based algorithm for joint, linear estimation of high-order image-to-mosaic transformations: Mosaicing the

curved human retina. *IEEE T. Pattern Anal.*, 24(3):412–419, 2002.

[4] 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 T. Pattern Anal.*, 24(3):347–364, 2002.

[5] M. DeGrezia and M. Robinson. Ophthalmic manifestations of HIV: an update. *J. Assoc. Nurses in AIDS Care*, 12(3):22–32, May-Jun 2001.

[6] Diabetic Retinopathy Research Group. Diabetic retinopathy study. Report Number 7. A modification of the Airlie House classification of diabetic retinopathy. *Inves. Ophth. & Vis. Sci.*, 21:210–226, Dec 1981.

[7] Early Treatment Diabetic Retinopathy Study Research Group. Results from the early treatment diabetic retinopathy study. *Ophthalmology*, 98:739–840, 1991.

[8] R. Hartley and A. Zisserman. *Multiple View Geometry*. Cambridge University Press, 2000.

[9] D. Hirvonen, B. Matei, R. Wildes, and S. Hsu. Video to reference image alignment in the presence of sparse features and appearance change. In *Proc. CVPR*, volume II, pages 366–373, 2001.

[10] J. V. Miller and C. V. Stewart. MUSE: Robust surface fitting using unbiased scale estimates. In *Proc. CVPR*, pages 300–306, 18–20 June 1996.

[11] J. Neter and W. Wasserman. *Applied Linear Statistical Models*. Richard D. Irwin Inc., 3rd edition, 1975.

[12] S. Peleg and J. Herman. Panoramic mosaics by manifold projection. In *Proc. CVPR*, pages 338–343, 1997.

[13] 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.

[14] 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.

[15] H. Sawhney and R. Kumar. True multi-image alignment and its application to mosaicing and lens distortion correction. *IEEE T. Pattern Anal.*, 21(3):235–243, 1999.

[16] H. Shum and R. Szeliski. Systems and experiment paper: Construction of panoramic image mosaics with global and local alignment. *IJCV*, 36(2):101–130, 2000.

[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.

[18] C.-L. Tsai, A. Majerovics, C. V. Stewart, and B. Roysam. Disease-oriented evaluation of dual-bootstrap retinal image registration. In *Proc. 6th Int. Conf. Med. Image Computing and Computer-Assisted Intervention*, volume II, pages 754–761, Montreal, Canada, 2003.

[19] R. Weleber and K. Gregory-Evans. *Retina*, chapter 18. Retinitis Pigmentosa and Allied Disorders, pages 362–460. Mosby, 2001.

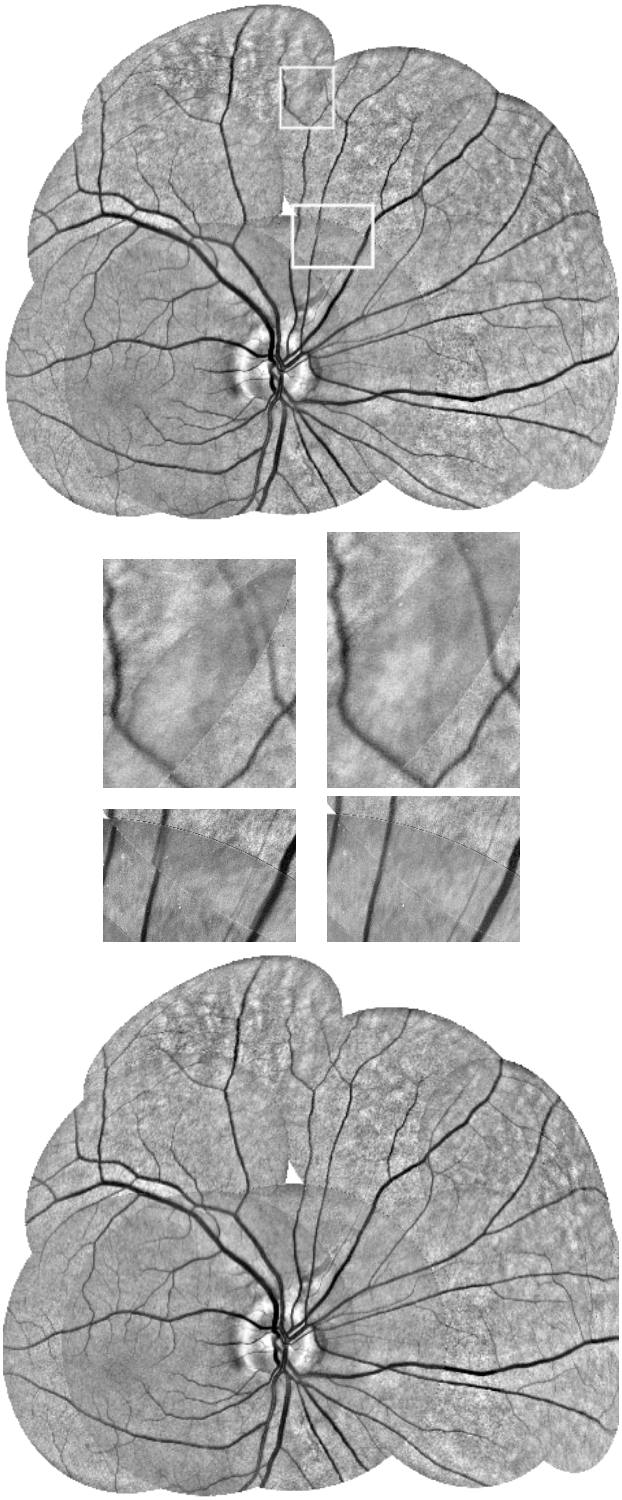

Figure 3: A subset of a 30-field image set illustrating misalignment before application of the new covariance-driven constraint-generation technique (top) and correction afterwards (bottom). The image sections shown in the center show regions of misalignment (left) and their correction (right) using the new technique.