

# Fully Automated Montaging of Laser Scanning In Vivo Confocal Microscopy Images of the Human Corneal Subbasal Nerve Plexus

Jason T. Turuwhenua,<sup>1,2</sup> Dipika V. Patel,<sup>3</sup> and Charles N. J. McGhee<sup>3</sup>

**PURPOSE.** Creating wide-field montages of the human corneal subbasal nerve plexus using laser scanning in vivo confocal microscopy (IVCM) requires considerable expertise and remains highly labor intensive. A typical montage contains several hundred images to be quality checked and manually arranged. The purpose of this study was to develop and validate software for off-line montaging of IVCM images of the living human cornea.

**METHODS.** The software was developed and tested using four large data sets of IVCM images from normal human corneas. Two of the data sets were used for calibration purposes, the remaining images served as a validation set. Techniques utilized included image binarization, clustering, key-point generation, and feature-based stitching. A range of tests involving computer processing and visual inspection were applied to audit and compare the automated montages with manually constructed montages.

**RESULTS.** The original IVCM images ( $N = 2565$ ) from four corneas were processed into image groups, reducing the number of effective images by 68% to 86%. Each data set contained a large primary grouping. A clustering strategy was used to reduce the total potential workload by 57%. Both programmatic and visual inspection confirmed the method was robust to errors, with a specificity of 100% (i.e., no falsely matched images). The time taken to complete the montage varied from 1.5 to 3 hours.

**CONCLUSIONS.** Computer-driven image stitching is a useful, effective, and time-saving tool for studies involving IVCM corneal nerve imaging. Further research will extend and optimize these methods. (*Invest Ophthalmol Vis Sci.* 2012;53:2235–2242) DOI:10.1167/iops.11-8454

The cornea is one of the most densely innervated tissues in the human body, its innervation originating from the ophthalmic branch of the trigeminal nerve. Nerve bundles enter the peripheral midstroma in a radial pattern and course

anteriorly, giving rise to multiple branches innervating the anterior and midstroma. Branches subsequently form the subepithelial nerve plexus that lies at the interface between Bowman's layer and the anterior stroma. Nerve bundles then penetrate Bowman's layer throughout the central and peripheral cornea. Bundles then divide and run parallel to the corneal surface between Bowman's layer and the basal epithelium, forming the subbasal nerve plexus. Nerve fibers subsequently form branches that enter the corneal epithelium, where they terminate.<sup>1,2</sup>

In recent years in vivo confocal microscopy (IVCM) has increasingly been used to enhance our knowledge of corneal nerves in health and disease. Subbasal nerves are the most readily imaged and easily quantifiable of corneal nerves. Epithelial nerves cannot be imaged using IVCM because they are beyond the resolution of the instrument. Although stromal nerves are visible, their variable orientations and depths make accurate and repeatable quantitative analysis difficult.<sup>3</sup>

IVCM has enabled elucidation of the two-dimensional structure of the normal, human, corneal subbasal nerve plexus, demonstrating a radiating pattern of nerve fiber bundles converging toward an area approximately 1 to 2 mm inferior to the corneal apex in a whorl-like pattern.<sup>4,5</sup> However, manually creating a full montage of the cornea is a time-consuming process, which requires judging the quality of many hundreds of images and piecing them together manually. Automation of this process is therefore attractive, offering efficiency and enabling larger scale studies to be performed.

The purpose of this study was to develop and test a system for semi-automated montaging of IVCM images of the living human corneal subbasal nerve plexus. The study also considers the problem of montaging large sets of acquired images, assuming only some ordering of images within a given set.

## METHODS

### Image Acquisition

Laser scanning IVCM was performed using the Heidelberg Retina Tomograph II Rostock Corneal Module (RCM; Heidelberg Engineering GmbH, Heidelberg, Germany). Four large sets of 8-bit grayscale images (384 by 384 pixels) were obtained from three subjects with normal corneas (containing 373, 640, 706, and 846 images, respectively). For each image set, a completed montage was available as an encapsulated postscript (EPS) file. This file contained a manually created montage by an expert user (DVP) (using Macromedia Freehand 10.0, Townsend, San Francisco, CA). The identities, dimensions, and positions of the component images were extracted by computer processing codes written using MATLAB (The MathWorks, Inc., Natick, MA).

From the <sup>1</sup>Auckland Bioengineering Institute, University of Auckland, Auckland, New Zealand; <sup>2</sup>Department of Optometry and Vision Science, New Zealand National Eye Centre, Faculty of Science, University of Auckland, Auckland, New Zealand; and <sup>3</sup>Department of Ophthalmology, New Zealand National Eye Centre, Faculty of Medical and Health Sciences, University of Auckland, Auckland, New Zealand.

Submitted for publication August 21, 2011; revised November 30, 2011; accepted February 26, 2012.

Disclosure: **J.T. Turuwhenua**, None; **D.V. Patel**, None; **C.N.J. McGhee**, None

Corresponding author: Jason Turuwhenua, Auckland Bioengineering Institute, University of Auckland, Private Bag 92019, Auckland, New Zealand; j.turuwhenua@auckland.ac.nz.

## Description of System for IVCM Image Montaging

**Feature-Based Image Stitching.** Feature-based image stitching identifies common landmarks or keypoints across images that define the transformation needed to successfully overlap two images. Feature based montaging is generally characterized by the following six basic steps<sup>6</sup>: (1) Keypoints are automatically extracted from (typically) feature-rich images. (2) Putative or initial matches between keypoints (across images) are established. (3) A suitable matching algorithm estimates in-lying (“true”) matches from within the set of putative matches. The resulting matches give transformations that map pairs of images onto each other (yielding an image “stitch”). (4) The resulting stitch is then validated, after which (5) an overall montage is prepared. (6) Finally the montages can be postprocessed to blend the component images together, yielding a seamless final result.

The key modifications to the standard workflow are (a) an initial step before keypoint extraction, to remove images with unwanted artifacts, (b) a custom keypoint generating method based on the structure of corneal nerve branches, (c) clustering of images to reduce the number of image pairs tested, and (d) specific methods for assessing the quality of a particular stitch between an image pair. The final post-processing step (step 6) was not included in this pilot study but could be implemented in further studies.

**Image Acceptance/Rejection.** Many recorded images contained artifacts or had insufficient detail for stitching purposes. Such images had been removed from the manually stitched montages, but required automated detection and removal in this study. Epithelium routinely confounded processing by inducing unwanted matches across unrelated images. Simple automated processing was applied that detected epithelium and other features not typically associated with nerves. Thus images containing significant artifacts were excluded from further processing.

Images were first normalized by background estimation and subtraction, thereby reducing the influence of uneven illumination. This was followed by hysteresis thresholding, to produce a binary image. Standard binary image routines provided by MATLAB's Image Processing Toolbox, identified and grouped connected pixels into bounded boxes. Additional operations were applied to this information for the purpose of image rejection (vide supra) and Figure 1 summarizes the rejection process, showing the binary image, the original grayscale image, and a typical corneal nerve, detected by its low density (pixels/bounding box area).

## Clustering of Images

Accepted images were paired and tested to determine whether they could be matched. However, the very large number of possible pairings—theoretically  $[N \times (N - 1)]/2$  for  $N$  images—needed to be reduced for maximum efficiency. This reduction was achieved by grouping images into smaller clusters based on gross nerve orientation and the natural ordering of images within the data set. The nerve orientation measure was determined by Hough transform,<sup>7</sup> which detected the angle of the largest straight lines in the image (see Fig. 2). The actual sorting into clusters was done using standard means.<sup>8</sup> For each cluster member, the neighboring images (i.e., in order of acquisition) were also included. These images were often related to the member (thereby improving clustering) but eye movements disrupted this relationship order in many cases. Pairs were created within each resulting cluster. In this testing 10 clusters were empirically specified.


## Keypoint Generation

Standard keypoint generators create unique identifiers based on local image structures, that can be compared across images for similarity.<sup>6</sup> However, in corneal nerve images useful information is contained in slender nerve structures superimposed on background fluctuations. Therefore a standard keypoint generator,<sup>9,10</sup> in the absence of this information, may readily generate keypoints in empty space (e.g., due to background noise).


The high quality of the binarization procedure offered a simple solution: all non-zero pixels were taken to be segments of corneal nerve, and hence were potential keypoints. *Actual* keypoints were generated from these pixels with the aid of simple heuristics. Firstly, a random sample of potential keypoints was selected from the entire set of nonzero pixels. Next, a keypoint identifier (a feature) was calculated by placing a circle (of specified diameter) around each candidate keypoint. The binary intensity was measured around the perimeter, but only identifiers with three “branches” were actually retained. The particular pattern of (binary) intensity (keypoint ID) was used for the purpose of matching between image pairs (see Fig. 3).

## Image Stitching

The images within a pair to be matched were termed “base” and “input” images. A particular keypoint in the base image can be compared with all keypoints in the input image for a best match. If a



**FIGURE 1.** Two simple tests were applied to detect “bad” images. Firstly, a “cell” detector was used to detect the “honeycomb” structure of (primarily) epithelium. The density of the heaviest bounding box was then determined and thresholded: allowing detection of other unwanted structures (e.g., epithelial cells visible on non-tangential images). (A) A binarized image indicating a low density bounding box typical of corneal nerve (the green box), a high density bounding box (cyan box) and holes typical of lighting and epithelium (red boxes). (B) The same information overlaid on the original grayscale images.



**FIGURE 2.** Example of clustering algorithm, using gross orientation detection. The green bars are potential candidates for gross orientation, the long red bar is the longest candidate found using the nerve detector.

particular keypoint (in the input image) also best matches the original keypoint in the base image then it is “double-matched.” Putative matches were established by this double-matching procedure, which was facilitated by standard means.<sup>11</sup>

Random consensus sampling (RANSAC) was used to robustly determine “true” matches from within the set of all putative matches.<sup>12</sup> Figure 4a shows the result of the RANSAC method, whereby outliers (mismatches) have been separated from inliers (valid matches). The inliers can be used to yield parameters for the transformation that stitches the input image onto the base image. The stitch suggested by Figure 4a is shown in Figure 4b.


### Assessing the Quality of Stitching

Two incorrectly matched images will routinely yield a highly distorted stitch, primarily including large shears and scaling. Sanity testing the geometrical shear and scaling inferred by a stitch was used to identify distorted transformations. If a stitch was not highly distorted, hence a probable valid match, an additional test was applied and overlapping areas of the stitched images were checked for similarity.

In this work we used a custom measure for this purpose. A pixel in the overlapping part of the base image will have a nearest pixel in the overlapping part of the input image. Conversely, a pixel in the input image will have a closest pixel in the base image. A histogram of these pixel vicinities provides an image “similarity” measure. The graph itself summarizes the number of pixels (on the vertical axis) that fall within a certain distance of each other (the horizontal axis). Thus if the area under the histogram lying within a prescribed distance (e.g., 2 pixels), exceeds a given threshold (e.g., 70% of total area), then the images are similar and hence connected.

### Final Grouping of Overlapped Images

The resulting list of connected pairs of images (the adjacency list) and its transformation can be used to determine image groupings that comprise the final montage. A breadth first walk<sup>13</sup> was used to form the final stitch—a connected grouping of images that could be checked both visually and programmatically against the manually created



**FIGURE 3.** The keypoint generation module, with keypoints shown as cyan triangles. The keypoint ID is shown in the inset, which is essentially the intensity around the perimeter of a circle centered on a particular keypoint. The keypoint generating strategy tends to retain keypoints near branches; at the same time it retains many points that are not branches, but are still useful in obtaining matches between image pairs.

montages. This method was implemented for simplicity, there are alternative approaches that could be used.<sup>6</sup>

### Validation

Montaging software was developed in MATLAB with the aid of the Image Processing Toolbox. The toolbox provided binary image functions that significantly simplified processing throughout all stages. The program was tested on the four stacks of images. Results from the first two sets of images were reported previously<sup>14</sup> and were used here to calibrate the method before testing the two remaining validation sets.

Each image in a set was classified as “accepted” or “rejected.” The resulting classification was assessed by comparison with the manually created montages. Those images contained or not contained in a particular manually created montage defined the gold standard set of “accepted” and “rejected” images, respectively. The corresponding confusion matrices, sensitivities, and specificities of the method were determined by comparing with the automated results.

In order to assess the clustering procedure, the full list of clustered image pairs (over all 10 clusters), was compared against the known list of overlapped images (present in the manually created montages). To create this list, a computer program checked whether the bounding boxes of two given images in the manually produced montages overlapped. The resulting adjacency list of known overlaps was compared with the list of all clustered pairs of images determined automatically.

The accepted images were montaged into image groupings using the methods described. Thus a full list of matched pairs was generated (see example in Fig. 4). Two users experienced in interpreting IVCM images (DVP and JT) inspected each pair and visually assessed the quality of the in-lying matches. The in-lying matches were checked to determine whether they appeared to match corresponding features correctly.

The quality of the positioning of the images was then determined programmatically. For each group of images, the arrangement of images in that group was compared with the arrangement in the manually created montages. A given group was overlaid on the corresponding images in the manually created montages automatically, by aligning their geometric centers. The relative displacement (normalized by the




FIGURE 4. (A) An example showing putative matches between two related images. The inlying matches are highlighted in yellow, whilst the remaining outlying matches are shown in green. (B) A resulting successful stitch. It is noted that translation, rotation and some shear of the images appear to be evident.

corner to corner image distance) was calculated between the automatically and manually positioned images, if possible. In some cases, the automated montage contained images that did not appear in the manually created images, so these images were not included in the comparison. The resulting normalized displacements were collected into a histogram.

Finally, it is noted that the range of transformations allowed for generating pairwise stitches were not actually used in creating the final groups. Instead, only translations were allowed, so that comparisons could be made more appropriately with the manually created montages since they are created by shifting images into final position without rotation, scaling, or skew.

## RESULTS

### Image Acceptance or Rejection

The results for the rejection process performance, over all four image sets of 2565 images, providing the raw number and proportion of total images accepted or rejected by human or computer, are summarized in Table 1. The high sensitivity (80.6%) indicates that the automated method accepted the

TABLE 1. Confusion Matrix for the Image Acceptance/Rejection Step of the Montaging Method Aggregated over All Four Data Sets

| Image Rejection Step (Total Images = 2565)                 |                   |                  |
|--|-------------------|------------------|
| Automated  | Human             |                  |
|  | Accepted          | Rejected         |
| Accepted   | TP<br>1040<br>41% | FP<br>774<br>30% |
| Rejected   | FN<br>250<br>10%  | TN<br>501<br>20% |
| Total sensitivity = $\frac{100 \times TP}{TP + FN} = 81\%$ |                   |                  |
| Total specificity = $\frac{100 \times TN}{TN + FP} = 39\%$ |                   |                  |

TP, true positives; FP, false positives; FN, false negatives; TN, true negatives.

**TABLE 2.** Confusion Matrix and Performance Metrics for the Clustering Step of the Montaging Method Aggregated over All Four Data Sets

| Clustering Performance (Total Possible Pairings = 426,290)  |             |                 |
|---|-------------|-----------------|
| Automated   | Human       |                 |
|   | Overlapping | Not Overlapping |
| Paired  | TP          | FP              |
|   | 5239<br>1%  | 180,195<br>42%  |
| Not Paired  | FN          | TN              |
|   | 4413<br>1%  | 236,443<br>55%  |
| Pairs reduction = $\frac{\text{Total Not Paired (FN+TP)}}{\text{Total Possible Pairings}} = 54\%$ |             |                 |
| Total sensitivity = $\frac{100 \times \text{TP}}{\text{TP} + \text{FN}} = 54\%$                   |             |                 |
| Total specificity = $\frac{100 \times \text{TN}}{\text{TN} + \text{FP}} = 57\%$                   |             |                 |

majority of images that the expert human user also selected. The low specificity (30.9%) shows that the human was more aggressive in rejecting images. Further investigation revealed this latter result to be due to the fact that near duplicate images were frequently discarded from the final manually stitched montage. This was originally done to reduce the burden on computer resources, but in this automated application near duplicates resulted in better grouping results and so were retained. Roughly 71% of images were accepted for further processing.

**Clustering**

The clustering strategy reduced the number of pairs to test by over half (57%) of total possible pairings (Table 2). The clustered image pairs were compared with the known overlapped image pairs in the manually created montages. Around half of overlaps were present in the clustered image pairs (sensitivity of 54%). The high false positive (FP) ratio (42%) indicated that many of the clustered image pairs were not overlapping in the human created montage, pointing toward a certain amount of redundancy.

**Image Stitching**

Table 3 summarizes the matching performance of the automated method compared to the manual montage method.

**TABLE 3.** Confusion Matrix and Performance Metrics for the Matching Phase of the Montaging Method Aggregated over All Four Data Sets

| Image Matching Performance<br>(Total Clustered Pairs = 185,434)                  |             |            |                |
|--|-------------|------------|----------------|
| Automated  |             | Human      |                |
|  |             | Matched    | Not Matched    |
| Automated  | Matched     | TP         | FP             |
|  |             | 2277<br>1% | 0<br>0%        |
|  | Not Matched | FN         | TN             |
|  |             | 7675<br>4% | 175,482<br>95% |
| Total sensitivity = $\frac{100 \times \text{TP}}{\text{TP} + \text{FN}} = 23\%$  |             |            |                |
| Total specificity = $\frac{100 \times \text{TN}}{\text{TN} + \text{FP}} = 100\%$ |             |            |                |

TP = True Positives, FP = False Positives, FN = False Negatives, TN = True Negatives.

The table breaks down the raw number and percentages (of total possible pairs) classified as matching or not matching. The most important value is the zero FP rate, or equivalently, the maximum specificity of 100%. This is highly desirable; even a single falsely matched image (FP > 0) will ultimately create an incorrectly stitched montage. Visual inspection confirmed that the image pairings were acceptable, and no falsely matched images were detected.

The accepted images were sorted into image groupings, reducing the effective number of images by 68% to 86% over the data sets tested. For each data set, the montaging produced a single large primary image group, followed by smaller and increasingly more numerous groups as summarized by Figure 5. Figure 6 shows the results of the automated displacement tests. The tested groupings have good agreement with their manually created counterparts, with >90% of images within a 10% normalized displacement. Finally, an example of a single image grouping is shown in Figure 7.

In terms of efficiency it was observed that the range of times taken for the automated stitching process ranged from 1.5 to 3 hours, using a 2.2 GHz Intel Core Duo Mac-Book Pro with 2 GB of RAM.

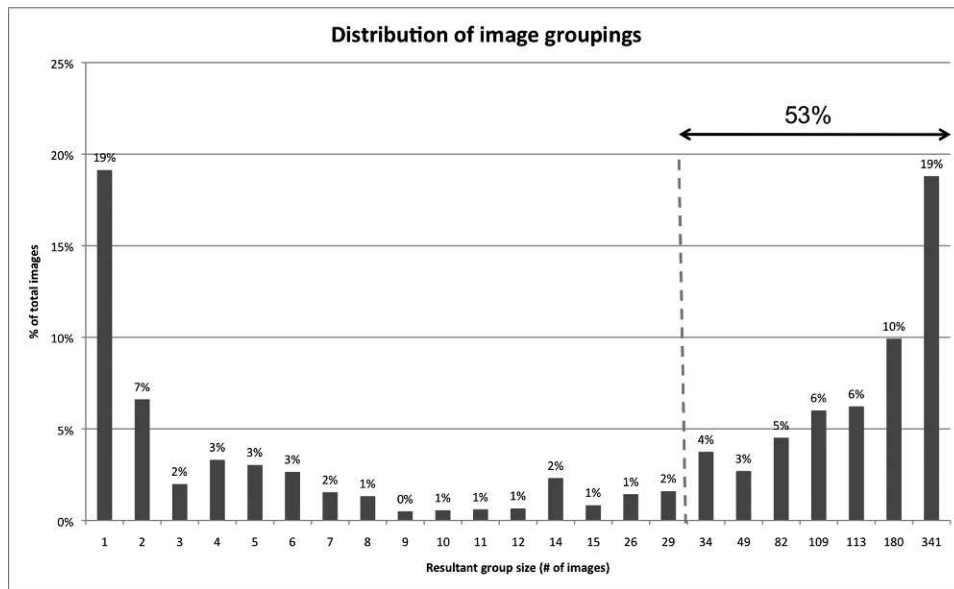
**DISCUSSION**

The use of immunohistochemically stained, thick anterior-cornea whole mounts have enabled visualization of the entire corneal innervation.<sup>1,2</sup> Such techniques have revealed important three-dimensional relationships and provided detailed and accurate information regarding nerve dimensions and density. Although IVCM cannot provide images with the same degree of detail (IVCM cannot image nerve branches and terminals of less than 0.5 μm in diameter), it does have the advantage of allowing imaging of the living human cornea. The noninvasive nature of IVCM therefore enables serial imaging of the same cornea over time.<sup>5</sup>

IVCM has increasingly revealed the complexity of the human corneal subbasal nerve plexus but manual, wide-field, montage-type reconstruction of the plexus is extremely time consuming (taking approximately 10-20 hours depending on the number of images).<sup>4,5,15</sup> This report describes software developed specifically for more automated montaging of large data sets of IVCM-acquired images of the corneal subbasal nerves plexus.

The methods developed here are relevant when images have been obtained and processed off-line, with some assumption as to ordering of the images in the data. This approach successfully reduces the original number of single accepted images to a number of smaller groups, the reduction ranging from 68% to 86%. For each data set, a single primary group was identified, accounting for 41% of all accepted images. This group was followed by additional groups decreasing in size, but becoming more numerous. We identified good agreement between the positioning of images in both automated and manually created montages and >90% of images were within a 10% displacement, indicating good positional agreement. This agreement was also confirmed by visual inspection.

These encouraging results suggest that automated montaging of IVCM data is robust and will significantly reduce the workload required to create a full corneal montage. Indeed, the computerized method completed single eye IVCM montages in 1.5 to 3.0 hours. Although further, but limited, human intervention is required to finalize the montaging task, the authors nevertheless believe this intervention represents a significant time savings compared to manual techniques.



**FIGURE 5.** The distribution of image groupings. The proportion of images (vertical axis) belonging to a group of a particular size (horizontal axis). Over half of the accepted images were assigned into seven large groups; the four largest groups being the primary group of each image set (35% - 61% of total accepted images in each set). At the lower end of the scale, 19% of images remained unmatched (i.e., were not paired into a group at all).

Notably, the computer codes used were not fully optimized, so with further refinement time costs will be reduced further.

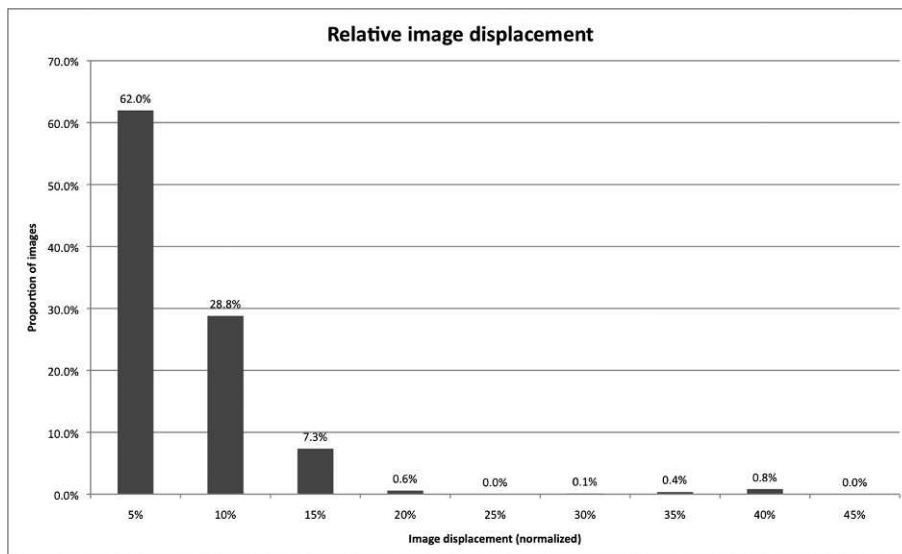
The clustering step allowed a significant reduction (by 57%) in the number of image pairs needed to be processed. Although the method in its present form performed well, the specificity and sensitivity of the clustering was relatively modest at 57% and 54%, respectively. Therefore, there is scope to improve this step further and hence improve upon the efficiency of the overall method.

The overall method was calibrated on two of the four image sets. The number of FP matches was zero on the validation set (i.e., the remaining two image sets) as well as the calibration set. The validation set result is highly desirable since it highlights that errors on the part of the computer system in this pilot study were nonexistent and therefore should be rare

in practical applications. The corollary is that human interventions in this process could be kept to a minimum. However, we believe that manual inspection is still beneficial and that remediation for possible errors should be part of the montaging procedure.

The fundamental key to the entire montaging method was the binarization step, which is a significant departure from standard montaging methods. It formed the basis of the image rejection step, feature detection, key-point identifier generation, clustering, and the image overlap similarity measure. To this end, the MATLAB Image Processing toolbox was found to be particularly convenient because it facilitated many of the binary image operations used in this work.

Recent work<sup>16,17</sup> has introduced real-time montaging of the subbasal nerve plexus; the key benefit of this approach is a



**FIGURE 6.** Relative displacement of images. The graph shows the proportion of images (vertical axis) lying within a band of relative displacement (horizontal axis).




FIGURE 7. Example of an automated grouping.

significant reduction in the time required to produce a montage. Whilst the stitching methods used in Zhivov et al.<sup>16</sup> (in particular) were not described in detail, the encouraging results we have encountered could also be integrated into that real-time approach. It may be possible to assist in improving upon the reported field-of-view results and overall quality of those montages.

Efron et al.<sup>15</sup> recently reported an alternative technique for mapping the corneal subbasal nerve plexus using the video capture facility (sequence mode) of the RCM. Images are captured while the subject tracks a moving target on a large computer screen. This procedure, which takes about 20 seconds and captures 100 contiguous images, is repeated along 13 radial meridians. The second stage of montaging is performed with Image-Pro Plus 7 software (MediaCybernetics, Bethesda, MD) to align and blend the radial image strips together. Although this is a relatively fast, semi-automated method, it suffers the disadvantages of reduced montage quality with fewer subbasal nerve branching details, particularly between imaged strips, due to the blending process.

Corneal nerves are of great interest to both clinicians and scientists due to their important roles in regulating corneal epithelial integrity, proliferation, and wound healing in addition to their protective functions.<sup>18</sup>

The ability to produce wide-field montages provides a more global view of subbasal nerve density and architecture rather than the limited information provided by single localized images. It is envisaged that the montaging technique described in this study will aid our ability to elucidate the effects of trauma, surgery, and disease on corneal innervations, and allow noninvasive testing of the effects of novel therapeutic

intervention strategies for diseases affecting corneal nerves. It may also be used in combination with function tests such as corneal sensitivity testing to determine structure-function relationships. A limitation of imaging to produce montages is the requirement for extended contact between the cornea and the microscope, precluding its use in patients with epithelial defects. In cases where subbasal nerves are sparse or fragmented, acquisition of multiple overlapping images is difficult and therefore reduces the likelihood of successful montaging (either manually or using the software described in this study). An observation from this study is that sparse nerve images tend to yield poor matches (resulting for example, from reduced numbers of valid keypoint matches). However, we expect that abnormal corneas with tortuous and distorted nerves could be successfully matched, provided adequate numbers of keypoints across an image can be generated.

This study demonstrates that creation of a wide-field IVCM montage image of the corneal subbasal nerve plexus can be achieved using an automated custom image stitching approach. Such automated montaging provides a useful, time-effective means of manipulating large data sets of images taken over the entire cornea. It is anticipated that these methods will help facilitate large studies of subbasal nerve structure in corneal health and disease, including investigating changes in nerve structure over time.

## References

1. He J, Bazan NG, Bazan HEP Mapping the entire human corneal nerve architecture. *Exp Eye Res.* 2010;91:513-523.

2. Marfurt CF, Cox J, Deek S, Dvorscak L. Anatomy of the human corneal innervation. *Exp Eye Res.* 2010;90:478-492.
3. Patel DV, McGhee CNJ. In vivo confocal microscopy of human corneal nerves in health, in ocular and systemic disease, and following corneal surgery: a review. *Br J Ophthalmol.* 2009;93:853-860.
4. Patel DV, McGhee CNJ. Mapping of the normal human corneal sub-basal nerve plexus by in vivo laser scanning confocal microscopy. *Invest Ophthalmol Vis Sci.* 2005;46:4485-4488.
5. Patel DV, McGhee CNJ. In vivo laser scanning confocal microscopy confirms that the human corneal sub-basal nerve plexus is a highly dynamic structure. *Invest Ophthalmol Vis Sci.* 2008;49:3409-3412.
6. Szeliski R. Image alignment and stitching: a tutorial. *Found Trends Comput Graph Vis.* 2006;2:1-104.
7. Duda RO, Hart PE. Use of the Hough transformation to detect lines and curves in pictures. *Commun ACM.* 1972;15:11-15.
8. Teknomo K. K-means clustering tutorials. Available at: <http://people.revoledu.com/kardi/tutorial/kMean/index.html>. Accessed August 10, 2011.
9. Harris C, Stephens MJ. A combined corner and edge detector. In: Proceedings of the Fourth Alvey Vision Conference; Manchester, UK. 1988:147-152.
10. Lowe DG. Object recognition from local scale-invariant features. In: Proceedings of the Fourth IEEE International Conference on Computer Vision; Corfu, Greece. 1999;2:1150-1157.
11. Kovese PD. MATLAB and OCTAVE functions for computer vision and image processing. Available at: <http://www.csse.uwa.edu.au/~pk/research/matlabfns/citesite.html>. Accessed August 10, 2011.
12. Hartley R, Zisserman A. Multiple View Geometry in Computer Vision. 2nd ed. Cambridge University Press; 2003.
13. Heineman GT, Pollice G, Selkow S. Algorithms in a Nutshell. 1st ed. Sebastopol, CA: O'Reilly Media, Inc.; 2008.
14. Turuwhenua J, Patel DV, McGhee CNJ. Semi-automated montaging of the entire corneal sub-basal plexus. *Invest Ophthalmol Vis Sci.* 2009;E-Abstract 3696.
15. Efron N. The Glenn A. Fry Award Lecture 2010: Ophthalmic Markers of Diabetic Neuropathy. *Optom Vis Sci.* 2011;88:661.
16. Zhivov A, Blum M, Guthoff R, Stachs O. Real-time mapping of the subepithelial nerve plexus by in vivo confocal laser scanning microscopy. *Br J Ophthalmol.* 94:1133-1135.
17. Allgeier S, Zhivov A, Eberle F, et al. Image reconstruction of the subbasal nerve plexus with in vivo confocal microscopy. *Invest Ophthalmol Vis Sci.* 2011;52(9):5022-5028.
18. Oliveira-Soto L, Efron N. Morphology of corneal nerves using confocal microscopy. *Cornea.* 2001;20:374.