Optical flow for Validating Medical Image Registration

J. R. Cooper (principal author)
Department of Computer Science
Curtin University of Technology
Perth, WA, Australia, 6845.
email: jc@cs.curtin.edu.au

N. Ritter
School of Information Technology
Murdoch University
Perth, WA, Australia, 6150.
email: N.Ritter@Murdoch.edu.au

ABSTRACT
Many approaches to the task of computing medical image registration have been presented, but there is little knowledge of how to evaluate the quality of the models of transformation these approaches use, or accuracy of the computed parameters. There is the well known technique of comparing the reference image to a transformed secondary image [1], but evidence to support or deny accuracy of the transformation is still hard to obtain.

This paper presents a technique to give researchers and clinicians clear visual evidence to validate the accuracy of 2D and 3D registration of medical images. It works by borrowing the computer vision technique of optical flow to compute disparities between the reference image and the transformed secondary image. The resulting disparity information may be presented as a needle diagram to assist with communication of results via paper, or used in subsequent steps of a registration algorithm.

KEY WORDS
Image Registration Validation, Optical Flow Application, Image Disparity, Medical Imaging

1 Introduction

Registration of medical images has application in topography analysis from multiple images and in comparing images taken at different times to track the progress of disease and aid diagnosis. This registration always involves developing a transformation model that explains the changes in the data and computing the parameters for that transformation for particular inputs. Once the transformation is computed, however, we need a way to validate it so that we may either reject the results, or develop confidence in them.

Many algorithms have been developed to determine such parameters from image data, but the difficult issue of deciding whether the computed transformation model and associated parameters are correct remains relatively unexplored [2]. A range of approaches to validation are cataloged by Woods [1]. These include visual inspection, comparison with results of other techniques and warping one image (the secondary image) via the transformation so that it matches the other image (the reference image). This enables direct comparison. The latter approach is the most convincing strategy, but we need a way to objectively compare the transformed image with the reference image in order to convince ourselves of the accuracy.

Disparities between these two images indicate errors in the transformation. These errors may result from problems with the transformation model itself, or inaccuracies in the parameters used to apply the model to a particular pair of images.

A number of techniques for comparing the reference image with the transformed secondary image have been used. The simplest is image differencing, where validation is obtained from examination of the pixel-wise brightness difference between the reference image and the transformed secondary image. This is susceptible to changes in lighting conditions and thus of limited applicability. A second approach is to construct a composite of the reference image and the transformed secondary by taking parts of one image and writing them over the other in a checker board pattern [3]. A third method is image switching, where the pair of images is alternated on a computer screen to allow the user to notice changes between the two.

Unfortunately, none of these techniques compute the required information. In validation we want to know whether the positions of the image features match in the transformed and the reference image. Image differencing reports the wrong information, since one is interested in the geometry of the transformation rather than the pixel brightnesses. The checker board approach allows one to find out if some features match up, but only where appropriate features can be found on checker board boundaries. Furthermore there is no objective number one can use to indicate how good the match is. Alternating images on the screen allows one to perceive disparities, but not to measure how large the they are and it is hard to see small disparities. All of these techniques are adversely affected by changes in lighting conditions between the images, so they make comparison of images taken at different times difficult.

This paper describes an image registration validation tool that uses an optical flow technique to measure disparities between a reference image and a transformed secondary image. The tool is demonstrated by applying it to the registration of human eye images taken at different times under different lighting conditions. The particular transformation is a combination of an affine transformation and a model of how the iris changes as the lighting conditions change [4].

Using optical flow to compare these images has some advantages. Firstly, the output of an optical flow algorithm is the computed disparity information—the differences in positions of image features between the two images. This is exactly what we need for validation. Secondly, this information is computed at many places in the scene so we can use patterns in the output to evaluate the transformation.

2 Optical Flow

A comparison of 9 optical flow techniques is given by Barron *et al.* [5]. They classify optical flow techniques into: differential techniques; Region-based matching; Energy-based models and Phase-based techniques. Of these, all but the phase-based techniques are susceptible to changes in lighting conditions because they compare gray levels between images. The phase-based techniques reprocess the image data by convolving it with a quadrature convolution kernel with zero DC component and then work with the phase of the filter response. The phase information, which is not affected by average image brightness, can be used via a differential technique to obtain a single constraint on the two-valued optical flow vector at each point. All the phase-based techniques then integrate flow information across the scene to obtain the second constraint on disparity.

2.1 The Algorithm

We use complex Gabor convolution kernels computed from

$$G_x(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{\frac{-x^2}{2\sigma^2}} e^{-ikx}, \qquad (1)$$

where k and σ are the peak frequency and the spread of the Gabor. Under convolution, this kernel responds to a range of image frequencies centered on k. The phase of the response is used to compute one component (v_x or v_y) of the disparity using

$$v_x = -\frac{\delta\phi/\delta t}{\delta\phi/\delta x}, \qquad (2)$$

where ϕ is the phase of the filter response; $\delta\phi/\delta t$ is the partial derivative of phase with respect to time; and $\delta\phi/\delta x$ is the partial derivative of phase with respect to position. These derivatives are estimated from the image data. We can put the two components together to obtain

$$\tilde{v} = \begin{pmatrix} v_x \\ v_y \end{pmatrix}.$$

This vector provides only a single constraint on disparity.

The image data contains noise and we can improve the results by taking this noise into account. We use the technique described by Hastings and Cooper [6] to compute the uncertainty in the disparity estimate $\tilde{v}$.

Using the computed single constraint flow values and associated uncertainties we apply the Kalman filter to combine flow values across a small region of the scene. The Kalman filter is essentially an average that takes into account the uncertainty in the data. In this case the Kalman filter integrates constraint information with different orientations to produce fully constrained disparity vectors. It also estimates uncertainties in those vectors.

2.2 Behavior of the phase-based Flow Algorithm

When interpreting the computed disparity output we need to know what image features the algorithm will respond to. The filters defined in (1) are complex-valued functions. The real part of the function consists of a cosine wave of frequency k multiplied by a Gaussian whose spread is σ. The imaginary part of the function is a sine wave of frequency k and also multiplied by the Gaussian. The wavelength of the function is $\lambda = 2\pi w/k$ pixels, where w is the width of the image. When convolved with the image data, these filters respond to image features of size approximately $\lambda/2$ pixels.

We also need to know what range of image disparities the system can measure. The algorithm computes disparity by dividing the difference in phase of the filter response in the two images at one point by the rate at which phase changes across the scene in one image as shown in (2). However, the phase of the filter response is an angle, which 'rotates' as the filter is moved past the image. This means that if the disparity between a pair of images is large enough to produce a between-image phase difference of more than about 90°, the output will be incorrect. For a kernel with wave length λ pixels, this means the disparity must be less than $\lambda/4$ to be measured correctly. With larger disparities the phase differences and hence computed disparities will be incorrect. The resulting disparity field will be essentially random, but it is improbable that all the disparity vectors coming from the Kalman filter will be zero, so a false impression of accurate registration is unlikely.

At places in the scene where the feature size is much larger or much smaller than the kernel wavelength, the filter response will be small. This would happen if a large filter is used when there are only small scale variations over the region, or if a small filter is used when the image is blurred or has almost uniform brightness in the region of interest.

At such places, the filter response becomes small and the uncertainty in the phase becomes large [6]. This is not a problem with the algorithm. It is a problem with the structure of the image features from which the disparity is to be computed. The algorithm merely records the fact that at some locations the image data does not provide disparity information. Algorithms that use this data can use the uncertainty information to ignore such uncertain values.

The resulting disparities can be used as an objective measure of the accuracy of the image transformation model and associated parameters. This can be done either by using needle diagrams to assist in the visual assessment of the transformation or by computing measures of 'quality of

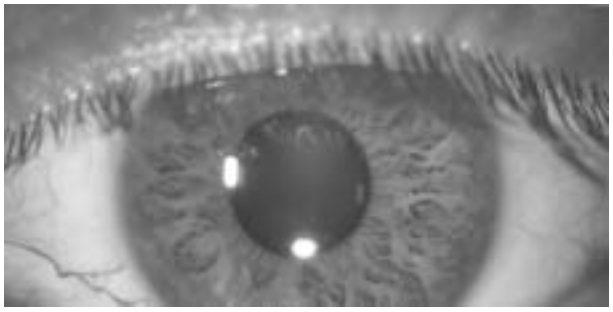

Figure 1. A reference image.

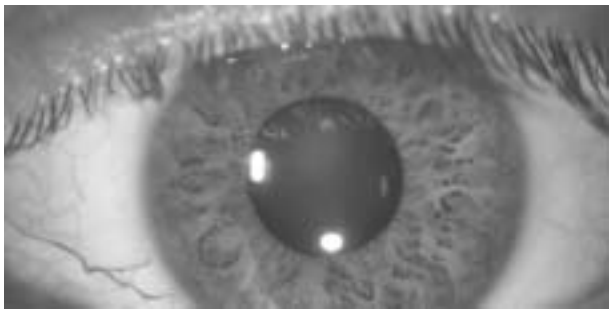

Figure 2. A second image to be registered with the reference image shown in Fig. 1. Note that the eye has moved as well and the pupil has changed its size and position within the iris.

fit' of the transformed image. When used for visual assessment, the needle diagrams can quickly show errors in the transformation model or its parameters because such errors lead to some parts of the transformed image matching the reference while others do not.

3 Application to a Registration Problem

The technique was used to assess the quality of registration of pairs of photographs of human eyes. To register such images, one must find the translation, scaling and rotation parameters as the eye and camera move with respect to each other. Since the iris is a large part of the image, one must also take into account changes in the pupil size and position— and associated iris distortions—that result from changes in light levels. Thus, the transformation that maps one image to the next has seven parameters [7]: rotation, scaling, x translation, y translation, pupil radius, and the x and y positions of the pupil center relative to the iris. A pair of input images can be seen in Figs. 1 and 2. Note that the pupil size is different as are the position, size and orientation of the eye.

A third image was created by transforming the image of Fig. 2 to match that of Fig. 1 using a transformation and other techniques given by Ritter [7]. This transformed image is shown in Fig. 3.

Comparing the reference image in Fig. 1 and the

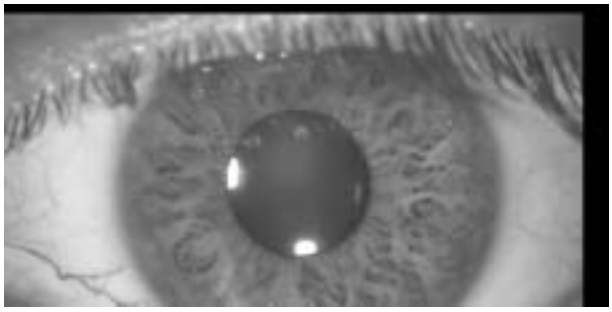

Figure 3. The secondary image shown in Fig. 2 transformed using techniques given by Ritter [7] to match the reference image shown in Fig. 1.

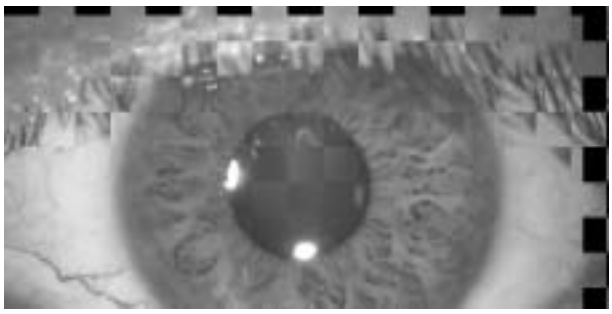

Figure 4. A checker board image constructed from alternate squares of the reference image shown in Fig. 1 and the transformed secondary image of Fig. 3. It takes careful on-screen study of this image to find disparities.

transformed image of Fig. 3 is difficult. Figure 4 shows a checker board image constructed by taking alternate squares from the two images [7]. At first glance this appears to show an almost perfect registration: there appears to be no discontinuity across the iris and the borders of the iris appear relatively smooth. It takes careful on-screen examination of this image to notice a miss-match in the vein position in the bottom left corner of the image. An enlargement of the relevant section is shown in Fig. 5. Other problems with the registration are not shown at all by the checker board approach.

Alternating the transformed image with the reference image on screen does indicate other problems with the registration, as well as more clearly highlighting the dislocation in the vein. However, we want a way to quantify the errors and measuring the disparities manually is tedious and error prone.

A needle diagram showing the disparities computed using the optical flow technique is shown in Fig. 6. In the needle diagram the disparity between the position of the vein in the images shown in Figs. 1 and 3 is immediately obvious. The image also shows small disparities in parts of the iris and clearly indicates the lack of information over much of the sclera. Disparities in the iris-sclera border are not recognized by the algorithm because the width of this feature—20 pixels—is much greater than half the wave-

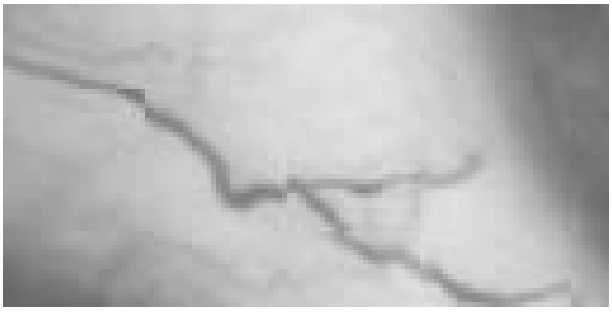

Figure 5. The vein from the left of the checker board image of Fig. 4. In this enlargement the disparity between the reference and transformed images becomes clearer.

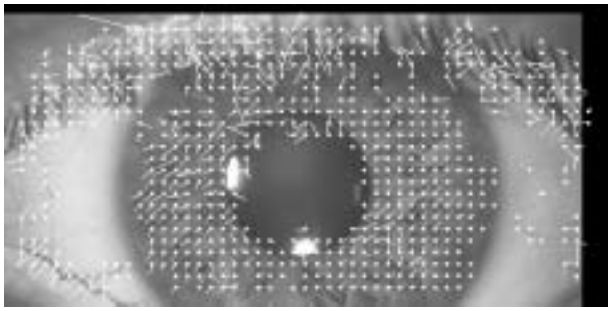

Figure 6. A needle diagram showing the disparities computed between the transformed secondary image (Fig. 3) and the reference image (Fig. 1) using the optical flow algorithm. The filter wavelength is 10 pixels. The needles are drawn on a 10×10 pixel grid and scaled in length by a factor of 5. Comparing this image with the checker board image shown in 4 one can see that the disparity in the vein on the left—shown more clearly in the enlargement of Fig. 5—is far more obvious in the needle image. Notice also that the needle image highlights other disparities not made obvious in the checker board image.

length of the filter.

Since the needles are drawn at 10 pixel intervals and scaled by a factor of 5, one can use the diagram to estimate the disparity in the position of the vein mentioned above as between two and three pixels. Manual measurements agree with this estimate.

3.1 Interpreting the Needle Diagram

The needle diagram provides much more information that used so far. The direction and extent of disparity is shown at many places in the image. This enables more detailed analysis of errors in the transformation. In subsequent paragraphs we present an example of transformation error analysis, using the computed needle diagram.

The seven parameter transformation [7] is an affine transformation combined with a particular model [4] of how the iris changes shape as the pupil moves, dilates and contracts as lighting levels change. If this model and the parameters calculated for this pair of images were completely accurate, one would expect the transformed image to match the reference image exactly. In that case the disparities between the transformed secondary image and the reference image would be zero and the disparity needles would have zero length. If the transformation parameters were not accurate, the needle diagram would show characteristic systematic disparity patterns.

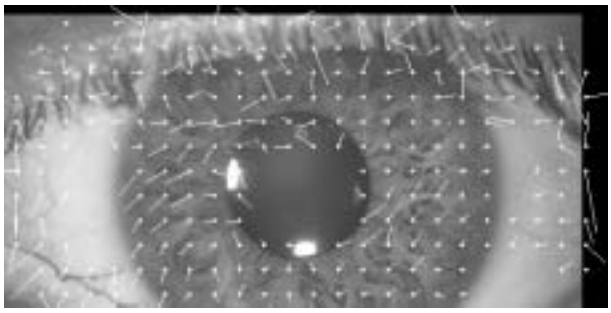

Figure 7. The same as Fig. 6 with the needles scaled by a factor of 10 and drawn on a 20×20 pixel grid.

There is such a pattern of disparities shown in Fig. 6. Parts of the sclera on the left appear moved down and further to the left, while parts on the right appear moved to the right. This is shown more clearly in Fig. 7, which shows the same disparities, but scaled by 10 pixels and printed on a 20×20 grid. We could thus conclude that the scale factor and the rotation used in the transformation was inaccurate.

The needle diagrams in Figs. 6 and 7 also show that the transformation model is inaccurate in a subtle way. Disparity needles in the lower left of the iris show a pattern that appears like a 'swirl' when the pair of images is alternated on screen. The lower right side of the iris shows a large 60×60 pixel patch of iris with 'northwestern' disparities between the reference image and the transformed secondary image. Using the scale information provided by the grid in Figs. 6 and 7 we can estimate from the needle diagram that the disparity has a magnitude of between 1 and 1.5 pixels in both areas. These disparity patterns are not consistent with errors in any combination of the 7 parameters of the transformation model. We may therefore conclude that they represent deviations between the transformation model and the image data the model is designed to explain.

4 Potential Applications and Future Work

The dramatic way the output shows errors suggests that the technique could be used to improve on the accuracy of the transformation. One could use the computed disparities from strategically located image regions to refine the estimates of the transformation parameters.

The output can also be used to compute a measure-of-fit of the transformed image to assist with automating the quality assurance process. A simple function like the average magnitude of the disparities would enable a fast quality check of the results of another registration algorithm.

The optical flow algorithm used in this work also

computes the level of uncertainty in the disparity measurements. In the analysis presented here, uncertainty information was used to avoid showing spurious disparities in a displayed disparity field. In a parameter refinement process one could use the uncertainty information as an indication of the amount of evidence provided by each disparity vector. This could be done either by using the disparity field, or by feeding phase differences to a Kalman filter designed to estimate the transformation parameters directly.

The optical flow technique can be applied in three dimensions in the same way it is used here. In such domains one might need to show the disparity field in slices or use transparency to enable visualization.

The technique has one weakness. The flow algorithm is affected by the scale of the features, and only works well when the disparity is small. However this should not be a problem with images that have already been approximately registered.

The results shown here indicate that the method is sensitive to disparities of the order of one pixel. Registration errors of this magnitude are difficult to detect or analyse using other techniques.

5 Conclusion

The optical flow approach given in this paper provides a rich source of information to assist the visual validation of two dimensional or three dimensional medical image registration algorithms. The use of needle diagrams allows us to display disparities in such a way that problems with the transformation model as well as errors in the transformation parameters can be discovered and evaluated.

The approach also has potential for automating the evaluation of other registration algorithms by providing information to be used in registration accuracy checks or for registration refinement. Its ability to give a visual display of disparities and highlight small disparities makes this tool a powerful improvement over previous validation techniques.

References

[1] R. P. Woods, "Validation of registration accuracy," in *Handbook of Madical Imaging: Processing and Analysis* (I. N. Bankman, ed.), ch. 30, pp. 491–497, Syndey: Academic Press, 2000.

[2] A. Guéziec, K. Wu, A. Lalvin, B. Williamson, P. Kazanzides, and R. van Vorhis, "Providing visual information to validate 2-d to 3-d registation," *Medical Image Analysis*, vol. 3, pp. 357–374, 2000.

[3] N. Ritter, R. Owens, J. Cooper, R. Eikelboom, and P. P. van Saarloos, "Registration of stereo and temporal images of the retina," *IEEE Transactions on Medical Imaging*, vol. 18, pp. 404–418, May 1999.

[4] J. G. Daugman, "Biometric personal identification system based on iris analysis." Patent Office, 1994.

[5] J. Barron, D. J. Fleet, and S. S. Beauchemin, "Performance of optical flow techniques," *International Journal of Computer Vision*, vol. 12, no. 1, pp. 43–77, 1994.

[6] R. Hastings and J. Cooper, "Multi-scale optical flow using phase-based velocity tuned flow operators," in *Tenth Australian Joint Conference on Artificial Intelligence Proceedings-Poster Papers* (A. Sattar, ed.), (Perth, Western Australia), pp. 156–161, The ACS National Committee on Artificial Intelligence and Expert Systems, 2-4 DECEMBER 1997.

[7] N. Ritter, *Registration of Images of the Retina and Cornea*. PhD thesis, University of Western Australia, Department of Computer Science, Crawley WA, 6009, Australia, June 2001.