

Simulating Nailfold Capillaroscopy Sequences to Evaluate Algorithms for Blood Flow Estimation

P. A. Tresadern, M. Berks, A. K. Murray, G. Dinsdale, C. J. Taylor and A. L. Herrick

Abstract—The effects of systemic sclerosis (SSc) – a disease of the connective tissue causing blood flow problems that can require amputation of the fingers – can be observed indirectly by imaging the capillaries at the nailfold, though taking quantitative measures such as blood flow to diagnose the disease and monitor its progression is not easy. Optical flow algorithms may be applied, though without ground truth (*i.e.* known blood flow) it is hard to evaluate their accuracy.

We propose an image model that generates realistic capillaroscopy videos with known flow, and use this model to quantify the effect of flow rate, cell density and contrast (among others) on estimated flow. This resource will help researchers to design systems that are robust under real-world conditions.

I. INTRODUCTION

Systemic sclerosis (SSc) is a disease of the connective tissue that causes fibrosis and vascular problems leading to digital ulcers or, in extreme cases, gangrene requiring amputation of the fingers. Prevalence in the US adult population is around 250 per million, with an incidence rate of approximately 20 new cases per million per year [1]. Due to limited long-term clinical data, however, there are few effective treatments.

There is, therefore, a need for systems that can diagnose disease early and measure its progression over time, for example in response to drug treatment. Of the techniques for assessing systemic sclerosis [2], the gold standard is *nailfold capillaroscopy*: taking images, through a microscope, of the capillaries that lie flush with the skin close to the base of the fingernail (Fig. 1a). From these images, we can measure capillary shape, size and density; detect abnormalities and haemorrhages; and observe blood flow in video sequences.

Capillaroscopy is attractive because it is non-invasive and relatively cheap. Capturing high quality images for quantitative assessment with current systems, however, requires training and expertise, and can be time-consuming. This limits research that depends on measuring capillary properties.

A. Related Work

Properties such as vessel density and morphology can be measured manually from capillaroscopy images [3] using software that combines individual images to create a panorama of the whole nailbed and enhances the visual appearance of the vessels [4], [5]. Recent software has

This work was generously supported by The Wellcome Trust. P. A. Tresadern, M. Berks and C. J. Taylor are with the Centre for Imaging Sciences, University of Manchester, United Kingdom M13 9PT. A. K. Murray, G. Dinsdale and A. L. Herrick are with the Institute of Inflammation and Repair, University of Manchester, United Kingdom M13 9PT. Corresponding author: philip.tresadern@manchester.ac.uk




Fig. 1. Capillaroscopy images: (a) Frame from a real capillaroscopy sequence; (b) capillary loop from the sequence; (c) synthetic capillary generated by our system. The capillary loops in (b) and (c) have been rescaled to enhance contrast.

automated the measurement process for images of a given quality [6], [7].

One measurement of interest is blood flow velocity, which can vary rapidly over time. Though efforts have been made to estimate the accuracy of image-based measurements with respect to other computerized solutions [8] or between different image-based methods [9], such evaluations are limited by a lack of ground truth flow values. Although using simulated images to evaluate flow estimation has recently been proposed, sequences generated currently lack realism [10].

Image-based flow methods are typically based on one of two paradigms: optical flow [11] exploits differential properties of images of a dynamic scene captured close together in time, whereas region matching methods [12] use local search over a neighbourhood and are good for dealing with large motions. The two approaches are complementary, though effort has been made to combine their strengths [13]. All can be evaluated using synthetic sequences [14] or real data, captured under highly controlled conditions with near-perfect ground truth [15].

B. Our Objectives

Our aim is to develop a system that will offer three main benefits to the clinic: images will be quick and easy to capture, using modern cameras with automated motion control; blood vessels in the image will be detected, segmented and measured automatically to give objective and quantitative measures of abnormality; and blood flow will be measured using image processing to provide complementary, dynamic data. These features will allow our capillaroscopy system to be rolled out more widely within healthcare systems.

Here we assess the efficacy of algorithms for measuring blood flow in capillaroscopy sequences by developing a

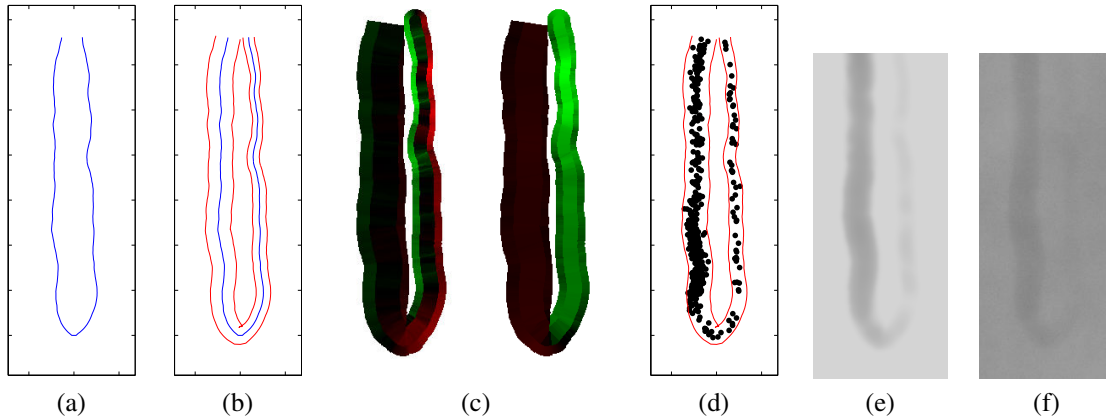


Fig. 2. Image synthesis process: (a) path of the vessel centre; (b) edges of the venous and arterial limb; (c) horizontal and vertical flow fields, respectively; (d) blood cell positions along the vessel centre; (e) synthetic image, before adding noise artefacts; (f) final synthetic image. For the flow fields in (c), hue indicates the direction of flow over the sequence (green is right/down, red is left/up) and intensity is proportional to flow rate (black indicates zero flow).

model that generates synthetic sequences of capillary blood flow, complete with exact flow values at every pixel.¹ These values can then be compared with those estimated from the image data to quantify error in estimated flow with respect to parameters such as flow velocity, cell density and noise.

II. METHODS

A. Generating Images

In generating an image sequence, we first define the capillary shape and its boundaries. Healthy capillaries have a characteristic ‘hairpin’ shape: the capillary is largely straight with a ‘u-turn’ at its apex. We therefore generate each side of the vessel path by tracing the locus of a point that travels with a constant vertical velocity and a horizontal velocity that falls to zero, plus a small random component so that the vessel is not perfectly straight (Fig. 2a).

A capillary has two limbs: a thin, arterial limb carrying blood from the heart; and a thicker, venous limb that carries blood back to the heart. We generate the edges of the two limbs using a stochastic process whereby we define the width at the apex and allow it either to grow or shrink to a constant value, with a small random component to model natural variation in width along the length of the vessel (Fig. 2b).

Next, we map the vessel centre and edges to a pixel grid where we can define a flow direction and magnitude at every pixel. Pixels that lie between the vessel walls are assigned a flow vector whose direction is parallel to the vessel path at its closest point, and whose magnitude is inversely proportional to the cross sectional area of the vessel at that point (such that flow is conserved along the length of the vessel). Pixels that lie just outside of the vessel wall are assigned a similar flow vector, plus a component toward the centre of the vessel to simulate collisions with the vessel wall. All other pixels are assigned a zero flow vector (Fig. 2c).

Now that we have the vessel structure and flow at every point, we can simulate the motion of individual blood cells

over time. To do so, we choose a fixed number of blood cells and initialize their positions at sampled points along the centre of the vessel. At each time instant for every blood cell, we interpolate the flow vector from the pixel grid at the cell’s position and apply the corresponding motion (plus a small random motion component) to the blood cell to compute its position at the following time step (Fig. 2d). By iterating this process, we generate a sequence of blood cell positions for any number of frames.

In a capillaroscopy image, an individual blood cell appears as a dark, blurry blob that we approximate with a Gaussian in the image. To simulate the appearance of plasma (gaps in the blood flow), we darken the image by a random amount for every cell. Next, we mask out pixels that lie outside the vessel boundary, and smooth the image. We then apply an offset and scaling to this ‘clean’ image to approximate the brightness and contrast of a typical capillaroscopy image (Fig. 2e).

Finally, we add artefacts to the image to simulate common causes of image degradation in capillaroscopy. First, we add a random but spatially correlated background texture to the image to simulate slight intensity variation in the underlying tissue.² We then modulate the brightness and contrast of the image randomly over the sequence to approximate illumination changes. Last, we add a small random translation to every image to simulate ‘shake’ that is almost unavoidable when viewing live structures at a microscopic scale. The final result is an image sequence that is qualitatively similar to a real capillaroscopy sequence (Fig. 2f).

B. Estimating Flow

Given a sequence of images, real or synthetic, we can apply optical flow methods to estimate blood flow for every pixel at every point in time. We assume, for now, that the images in the sequence have been registered to remove any small translations between frames; synthetic sequences can be generated such that this is so, whereas real sequences must go through a registration process [16]. We also assume that

¹Example sequences are available at the time of writing from http://personalpages.manchester.ac.uk/staff/philip.tresadern/proj_nailfold.htm

²This texture was generated using Matlab code provided by Peter Kovesi. <http://www.csse.uwa.edu.au/pk/Research/MatlabFns/Misc/noiseonf.m>

brightness (mean intensity) and contrast (intensity variance) are largely constant over the sequence, which also can be achieved by design or image post-processing.

In its simplest form, optical flow assumes that a pixel associated with a moving object will maintain its brightness over time such that, by a first order Taylor expansion,

$$I(x, y, t) = I(x + dx, y + dy, t + dt) \quad (1)$$

$$\approx I(x, y, t) + \frac{\partial I}{\partial x} \frac{dx}{dt} + \frac{\partial I}{\partial y} \frac{dy}{dt} + \frac{\partial I}{\partial t} \quad (2)$$

$$\Rightarrow 0 \approx I_x u + I_y v + I_t, \quad (3)$$

where I_x is the partial derivative $\partial I / \partial x$ (and likewise for y and t) that can be estimated by finite differencing; the flow vector components $u = dx/dt$ and $v = dy/dt$ are the quantities we want to know.

This one equation with two unknowns has an infinite number of solutions, since any motion parallel to an edge is unobservable (known as the ‘aperture effect’). Though global constraints can be applied to ensure a smoothly varying flow field [11], we instead pool measurements over a small window (*e.g.* 3×3 pixels) to increase the number of constraints on u and v . Because this is essentially a linear regression, we can also use an eigenanalysis to identify patches where the solution is poorly constrained and even quantify the uncertainty in the estimated flow directions.

Estimating flow from a single pair of noisy images, however, is error-prone. We therefore average over every consecutive pair of frames in the sequence to obtain the mean flow at every pixel. Alternatively, we could compute an approximate distribution of flow over time for every pixel and examine the properties of the distribution, such as variance which reflects uncertainty.

C. Experiments

The purpose of creating synthetic capillaroscopy sequences is to assess the accuracy of the estimated blood flow under different conditions. In particular, we want to know when the optical flow algorithm will work well or break down, and to quantify the effect of different parameters. We therefore generated sequences for various values of flow rate, cell density, brightness variation, contrast variation, translational ‘shake’ and additive Gaussian noise. For the purposes of this quantitative study, we did not add a background texture.

To quantify error, we computed flow at every pixel in every consecutive pair of images over the sequence. We then computed statistics for every pixel, such as the mean flow vector and its variance (independently in u and v) over the sequence. Finally, we computed the mean absolute error (MAE) in u and v between the estimated flow and the ground truth over the whole image. All sequences contained 180 frames, equivalent to three seconds of video at 60 frames per second (the frame rate of our hardware). Unless stated otherwise, baseline background brightness is 120 grey levels and contrast (*i.e.* the difference between the lightest and darkest pixels) is 16 grey levels.




Fig. 3. Qualitative comparison of true flow field with that estimated from optical flow in the (a) horizontal and (b) vertical directions. Hue indicates the direction of flow over the sequence (green is right/down, red is left/up) and intensity is proportional to flow rate (black indicates zero flow).




Fig. 4. Real image (Fig. 1) overlaid with estimated image flow in the (a) horizontal and (b) vertical directions. Because flow magnitude is not available for this real sequence, the colouring instead indicates the direction of flow over the sequence: green is right/down, red is left/up. The intensity denotes the proportion of the sequence in the given direction. The regions of flow in the two vessels are clear.

III. RESULTS

Qualitatively, the flow in a slow-moving column of blood was recovered well, particularly in the vertical direction in which most of the flow took place (Fig. 3). We also applied the algorithm to registered sets of real images, though we have no ground truth and can present only the estimated *direction* of flow; it would seem, however, that the algorithm detected the regions of flow in the sequence (Fig. 4).

When comparing quantitative error rates with respect to vertical flow magnitude (since this is where most flow takes place), mean absolute error varied approximately linearly with respect to flow rate (Fig. 5a). With respect to the number of cells (or cell density), error increased with the number of cells simulated (Fig. 5b).

Small random variations in brightness (Fig. 5c), ‘shake’ and added Gaussian noise (not shown) increased the error rapidly initially before reaching a plateau of approximately 0.5 to 0.6 pixels per frame, where further increases in the size of the variations had a much smaller effect.

Contrast variation, being a multiplicative (rather than additive) parameter has wider influence such that error continues to increase linearly within the range tested (Fig. 5d).




Fig. 5. Mean Absolute Error (MAE) in the vertical direction with respect to (a) relative flow rate, (b) number of simulated cells, (c) brightness variation and (d) contrast variation.

IV. DISCUSSION

Our results suggest that flow rate is the most critical parameter that affects error. This is not surprising because optical flow estimation algorithms consider only a small neighbourhood around any pixel at a given time, such that algorithms will be ineffective for flows greater than some value. This can be addressed in two ways: increase the frame rate of capture hardware; or apply a hierarchical pyramid scheme whereby the flow that is estimated for each spatially subsampled image is used to initialize flow estimation at the next level of the pyramid.

Flow is most accurately measured in regions of high image gradient (*i.e.* edges) so the *cell density* in the image also has an effect: as density increases, cells overlap and obscure each other's boundaries. Cell density, however, is largely out of our control which could affect the feasibility of estimating blood flow by capillaroscopy alone. Other imaging methods (*e.g.* OCT) may achieve accurate estimates, though this makes the system more complex and expensive, and requires a difficult data fusion step.

Random variation in brightness introduces errors in flow estimates via the temporal derivatives, whereas variation in contrast, translational noise ('shake') and Gaussian noise affect both spatial and temporal derivatives. Brightness and contrast changes can, however, be normalized to a large degree by applying a global transformation to every image that matches the brightness and contrast to a reference image (*e.g.* the first) for every frame. Shake, or 'jitter', in the sequence can be reduced by applying image registration [16] to the sequence before computing optical flow.

V. CONCLUSION

We have presented a model that simulates capillaroscopy sequences, from which we evaluate algorithms for estimating blood flow under various conditions. The results suggest that

apparent flow rate (*i.e.* distance travelled per frame) has the greatest effect, whereas the increase in error is limited for other parameters such as brightness variation.

Though the model has promise in designing and developing algorithms for estimating flow, it may have uses beyond this task. Because the model includes a parameter to add random translations (that may also be correlated over time by smoothing), we can simulate the effect of errors in registration to evaluate the accuracy of registration algorithms. Our model of image synthesis is also well-suited to generating annotated training data from which optical flow parameters can be learnt rather than prescribed [17].

ACKNOWLEDGMENTS

This work was supported by The Wellcome Trust.

REFERENCES

- [1] M. Mayes, J. Lacey, J. Beebe-Dimmer, B. Gillespie, B. Cooper, T. Laing, and D. Schottenfeld, "Prevalence, incidence, survival, and disease characteristics of systemic sclerosis in a large us population," *Arthritis Rheum*, vol. 48, pp. 2246–55, 2003.
- [2] A. K. Murray, T. L. Moore, J. B. Manning, C. Taylor, C. E. M. Griffiths, and A. L. Herrick, "Noninvasive imaging techniques in the assessment of scleroderma spectrum disorders," *Arthritis Rheum*, vol. 61, no. 8, pp. 1103–11, 2009.
- [3] M. E. Anderson, P. D. Allen, T. Moore, V. Hillier, C. J. Taylor, and A. L. Herrick, "Computerized nailfold video capillaroscopy – A new tool for assessment of raynauds phenomenon," *J Rheumatology*, pp. 841–848, 2005.
- [4] P. D. Allen, C. J. Taylor, A. L. Herrick, and T. Moore, "Image analysis of nailfold capillary patterns from video sequences," in *Proc. Int. Conf. on Medical Imaging Computing and Computer Assisted Intervention*, 1999.
- [5] —, "Enhancement of temporally variable features in nailfold capillary patterns," in *Proc. British Machine Vision Conf.*, 1998.
- [6] K. Feng, "Automatic nailfold capillary measurement and analysis," Ph.D. dissertation, 2009.
- [7] M. Paradowski, H. Kwasnicka, and K. Borysewicz, "Capillary blood vessel tortuosity measurement using graph analysis," in *KES*, vol. 2, 2009, pp. 135–142.
- [8] D. M. Mawson and A. C. Shaw, "Comparison of CapiFlow and frame by frame analysis for the assessment of capillary red blood cell velocity," *J Medica Eng Tech*, vol. 22, no. 2, pp. 53–63, 1998.
- [9] C.-C. Wu, G. Zhang, T.-C. Huang, and K.-P. Lin, "Red blood cell velocity measurements of complete capillary in finger nail-fold using optical flow estimation," *Microvascular Research*, vol. 78, pp. 319–324, 2009.
- [10] T.-C. Huang, W.-C. Lin, C.-C. Wu, G. Zhang, and K.-P. Lin, "Experimental estimation of blood flow velocity through simulation of intravital microscopic imaging in micro-vessels by different image processing methods," *Microvascular Research*, vol. 80, pp. 477–483, 2010.
- [11] B. K. P. Horn and B. G. Schunk, "Determining optical flow," *Artif. Intell.*, vol. 17, pp. 185–203, Aug. 1981.
- [12] B. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proc. Image Understanding Workshop*, 1981.
- [13] A. Bruhn, J. Weickert, and C. Schnorr, "Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods," *Int. J. Comput. Vis.*, vol. 61, no. 3, pp. 211–231, 2005.
- [14] J. L. Barron, D. J. Fleet, and S. S. Beauchemin, "Performance of optical flow techniques," *Int. J. Comput. Vis.*, vol. 12, no. 1, pp. 43–77, 1994.
- [15] S. Baker, D. Scharstein, J. P. Lewis, S. Roth, M. J. Black, and R. Szeliski, "A database and evaluation methodology for optical flow," *Int. J. Comput. Vis.*, vol. 92, pp. 1–31, 2011.
- [16] B. Zitova and J. Flusser, "Image registration methods: a survey," *Image Vision Comput.*, vol. 21, pp. 977–1000, 2003.
- [17] D. Sun, S. Roth, J. P. Lewis, and M. J. Black, "Learning optical flow," in *Proc. European Conf. on Computer Vision*, vol. 3, 2008, pp. 83–97.