Part-Based Multi-Frame Registration for Estimation of the Growth Of Cellular Networks in Plant Roots

T. J. Roberts, S. J. McKenna
Applied Computing, University of Dundee
troberts@computing.dundee.ac.uk

J. Hans, T. A. Valentine, A. G. Bengough
Scottish Crop Research Institute (SCRI)
Dundee, Scotland *

Abstract

Motion estimation from confocal scanning laser microscope images of growing plant cell structures presents interesting challenges; motion exhibits multiple local discontinuities and noise is non-isotropic and non-Gaussian. A method is presented for estimating motion of cell networks based on a physically motivated, part-based model of cell boundary structure. Each part models the shape and appearance of a localised image region and can undergo constrained non-rigid deformation. This enables motion discontinuities between parts to be modelled. Parts are coupled in order to improve localisation and increase computational efficiency. Results from applying MCMC show accurate localisation of the structure across multiple frames. The form of the model assists biologists in interpreting growth.

1. Introduction

Confocal scanning laser microscopy (CSLM) imaging of cellular structures can be used to obtain high resolution images of live plant tissue such as Fig. 1. Furthermore, specific cell components can be investigated by using different chemical or genetic markers[6]. Automated quantitative analysis of such images is challenging due to multiple, complex neighbouring motions and highly non-Gaussian noise, amongst other factors. This paper develops a physically motivated, part-based registration method to describe the highly coupled motion of plant cell networks and provides a description of growth that is easily interpreted by biologists.

Previous work on automatic estimation of growth in plant tissue from CSLM images has focussed on optical flow methods [2]. It is well known that motion constraints are necessary in order to perform motion estimation. In particular, brightness constancy and motion smoothness constraints are often adopted to make the problem well posed. A significant body of work directly addresses the formulation of non-local constraints on optical flow fields. The performance of some of these methods has been directly compared [1, 5]. Whilst smoothness priors can help increase accuracy and reduce ambiguity, smoothing over motion discontinuities is undesirable. Methods based on robust estimation and piecewise smoothing can partly overcome such drawbacks [3] and the authors have applied such fig:robusta method to plant root image sequences (see Fig .2). These approaches can be extended further by allowing smoothing to depend upon the image structure. *However, it has proven difficult to obtain accurate cell growth measurements using such general purpose optical flow methods. Furthermore, the description of motion thus obtained is not at an appropriate level of abstraction for biologists*. Instead, a more application-specific model is proposed for accurate, efficient motion estimation. Cell networks are modelled by decomposition into multiple, coupled deformable templates. Joint inference of the motion of these parts is performed. This approach is similar to that of deformable templates as used for tracking by several authors, e.g. [4]. However, the method presented here is novel because of (i) the highly coupled nature of the cellular network model and (ii) the likelihood model which uses a non-Gaussian noise distribution and mixture model for appearance prediction.

Currently, the growth of cells within plant roots and the motion of organelles within cells is not well understood. An accurate, quantitative motion estimation method would facilitate control experiments to investigate the effects of genetic manipulation, imaging set-up and growth conditions [2]. The system described here is applied to *Arabidopsis thaliana* roots that have been altered to express green fluorescent protein (GFP) in either the endoplasmic reticulum (ER) or the plasma membrane (PM). Visually, both of these markers delineate cell 'boundaries'. Estimating cell motion and growth using the ER marker is challenging since other complex motions are also apparent. Automatic detection of cell boundaries is complicated by blurring in depth, leading to boundaries from other layers being visible, and signifi-

*Thanks to UK BBSRC for funding and Jim Haseloff for sourcing the GFP-PM line

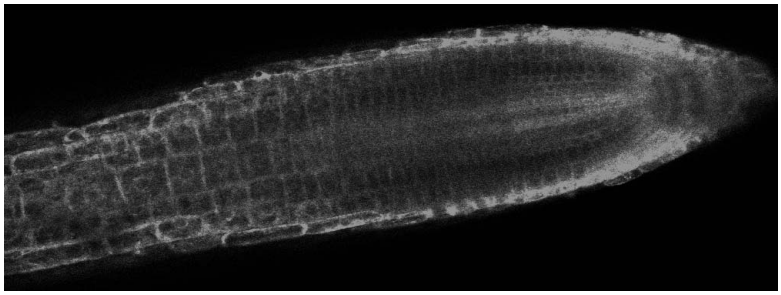

Figure 1. A confocal scan through a transgenic (GFP-ER) *Arabidopsis thaliana* **root.**

cant clutter in the case of ER datasets. This work therefore relies upon manual identification of trackable boundaries. In the future, better markers will make automatic initialisation easier.

Fig. 1 shows a CSLM image from a time-depth sequence. It was acquired on a Leica SP1 with a $20X$ objective at normal scan speed resulting in images with a resolution of 1024×1024. Confocal microscopes acquire slices in depth sequentially. Since the length of time to acquire adequate scans in depth is quite large, there is significant growth over the course of scanning in depth. Therefore, the data cannot be treated as instantaneously acquired volumes. *This paper focuses on estimating the in-plane 2D motion and defers the registration of slices to obtain 3D motion to later work*. Sequences used here had minimal motion in depth. The time interval between acquisition of images at the same depth is relatively large and inter-frame motion is correspondingly large ($\approx$ 7 pixels) making motion estimation challenging. In the portion of root being studied, expansion is $\approx 1/70$ pixels. At the resolution used, a visual boundary in the cell network actually corresponds to multiple cell boundaries and therefore multiple motions of interest are present at the single pixel level. The linear nature of cell boundaries and their lack of persistent texture mean that *local* motion can only be accurately measured normal to a boundary.

One of the difficulties in accurately establishing correspondence in CSLM images is the complex form of the image noise as shown in Fig. 3. The quantiles of the empirical distribution of noise in the intensity channel (from fast successive scans of the same root portion) are aligned with the quantiles of the normal distribution. If the noise was Gaussian this plot would be a straight line. In fact, three components can be identified: a central linear segment corresponding to approximately Gaussian *sensor* noise, a linear segment corresponding to approximately Gaussian *biological variation* (perhaps due to small scale motion of proteins) and large outliers in the tails.

Figure 2. Horizontal flow estimated from GFP-ER images using a robust method [3] (default parameters, smoothing level 2).

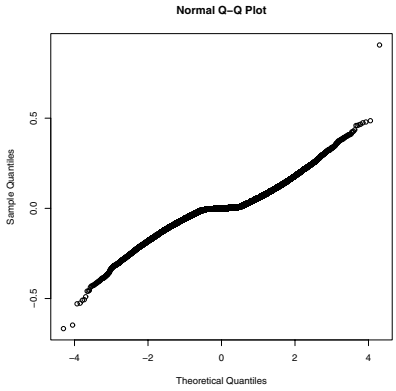

Figure 3. Normal quantile-quantile plot for pixel noise which is clearly non-Gaussian.

2. Method

Motion estimation is formulated here in a Bayesian generative probabilistic framework. This allows shape and motion priors to be incorporated explicitly and enables motion uncertainty, which is often significant and has a non-trivial form, to be represented and incorporated into subsequent stages of analysis. The following four subsections describe the method in terms of (i) a generative part-based model (ii) prior constraints on its parameters, (iii) a non-Gaussian likelihood model, and (iv) a Markov Chain Monte Carlo (MCMC) scheme for approximating the posterior over model parameters.

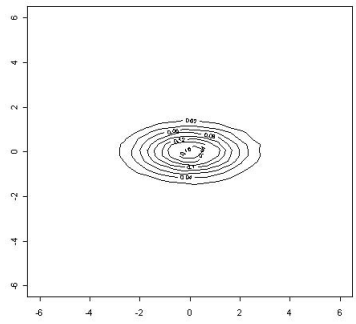

Figure 4. The distribution of acceleration parallel (horizontal) and orthogonal (vertical) to the direction of motion.

2.1. Coupled Part-based Model

In order to accurately model the complex motion and provide an easily interpretable description of growth, a generative part-based model is proposed. This paper only considers cell boundaries and defers modelling of other components such as cell nuclei to future work. Since individual cell boundary segments cannot be strongly localised (the aperture problem), a set of constraints must be adopted. To understand the motivation for the constraints adopted here consider Fig. 1 from which the following observations can be made. (i) The shape of cell boundaries is smooth and simple (often linear) between cell junctions. (ii) The boundaries are not smooth at junction points. (iii) There are no gaps between cells at this scale (due to internal cell pressure). Therefore, the overall boundary structure is represented in terms of a *deformable cellular network graph* with nodes encoding the positions, X, of junction points and (directed) edges encoding boundary segment shapes as Hermite cubics, H. Note that junction points and boundary segments are shared between cells. As Fig. 6 shows, this coupling of boundary segments enables good horizontal and vertical localisation. A Hermite cubic is fully specified by the positions of its two endpoints (nodes) and the derivatives at these endpoints (cell boundary segment shape). Due to lack of smoothness at junctions, the derivative parameters for the edges are independent. The model, W, is specified by the concatenation of the node parameters, indexed by i, and the edge parameters, indexed by j, over time, t: $W \equiv \{X_{i,t}, H_{j,t} : 0 \leq i < I; 0 \leq j < J\}$. The task of estimating growth (motion and shape change) is then that of recovering the posterior distribution over the model parameters $p(W|D) \propto p(D|W)p(W)$ given an ordered set of images $D \equiv \{I_t(x, y) : 0 \leq t < T\}$.

Figure 5. The 'un-warped' local appearance of a cell boundary segment

2.2. Motion Constraints

Independent dynamic priors on node parameters and boundary segment shape parameters were used. More specifically, in order to establish a suitable node motion prior, correspondence was manually established over time for sets of cell junction points. Fig. 4 shows the distribution of acceleration of these points in a co-ordinate system with accelerations parallel and orthogonal to the velocity. This distribution is well approximated as an uncorrelated zero-mean Gaussian with larger expected accelerations parallel to the velocity ($\sigma \approx 2pix.s^{-2}$) than orthogonal to it ($\sigma \approx 1pix.s^{-2}$). A zero-mean Gaussian prior on inter-frame boundary segment shape change was employed ($\sigma = 0.2$). Finally, to increase sampling efficiency, a uniform prior $[-5\%, 10\%]$ over segment length change was employed.

2.3. Likelihood

The likelihood was formed by comparing the local appearance around hypothesised boundary segments to the manually initialised network. Pixel-wise independence given the model was assumed. To determine the likelihood of a single pixel $I_t(x, y)$, the closest boundary segment, c, the distance along that segment, l, and the distance to that segment, d were determined (via optimisation). Using these parameters, the intensity in the manually initialised network, $I_0(c, l, d)$, was determined using nearest neighbour matching. Since we only want to specify the *local appearance*, the distance to the curve must be considered. Rather than matching all the points uniformly within some distance d_{max}, the distribution for each pixel was modelled as a mixture of an informative distribution (from the manually initialised appearance) and a non-informative uniform distribution, $U(0, 1)$:

$$p(I_t(x, y)|W) = \alpha(d)p(I_t(x, y)|I_0(c, l, d)) + (1-\alpha(d))U(0, 1)$$

The mixing parameter $\alpha(d)$ takes the form of a linear drop-off with distance up to a maximum of $d_{max} = 10$ pixels from the curve. For the informative distribution, $p(I_t(x, y)|I_0(l, d; x, y)$, a mixture of two Gaussians and a uniform distribution is used with parameters corresponding to the findings described in Fig. 3.

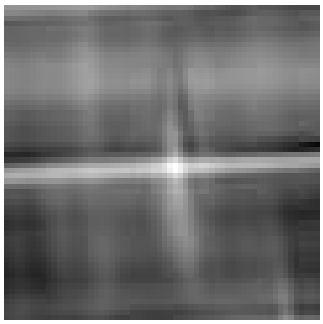

Figure 6. A log likelihood projection over translation (30×30) for a single cell model.

2.4. Estimation

MCMC was used to approximate the posterior. A component-wise Metropolis acceptance kernel was used to generate zero-mean Gaussian acceleration hypotheses ($\sigma = 1pix.s^{-2}$), node position changes ($\sigma = 1pix$) and shape changes ($\sigma = 0.2$). Sampling an acceleration variable that effects node positions at *all* times greatly increased the efficiency. The chain was run for 5000 iterations (taking ≈ 5 min. on a 2.4GHz PC). The acceptance rate was relatively low ($\approx 10\%$).

3. Results

Fig. 7 shows the difference between the most probable cell network synthesis and the last scan of a sequence with 6 frames (elapsed time ≈ 10 min). The manually initialised network is also shown. It can be seen that most boundaries are well localised. The poorest localisation was for cell walls for which the appearance had changed or whose appearance was less easily discriminated from the background (e.g. the blurred boundary segment at the bottom left of Fig. 7). If one inspects the expectation of cell boundary lengths over time a weak linear growth is observed. Furthermore, the uncertainty in cell length increases with time. A likely reason for this is appearance change.

4. Conclusions and Future Work

The coupled, part-based model described here allows for efficient, accurate estimation of root growth and provides a parameterisation that eases further analysis. The primary limitation of the method was the non-adaptive appearance model. The method will be extended to use a latent appearance model (allowing mosaicing and super-resolution/denoising). Extending this to a 3D latent appearance model would allow for modelling of the point spread convolution, attenuation, motion in depth and registration in depth.

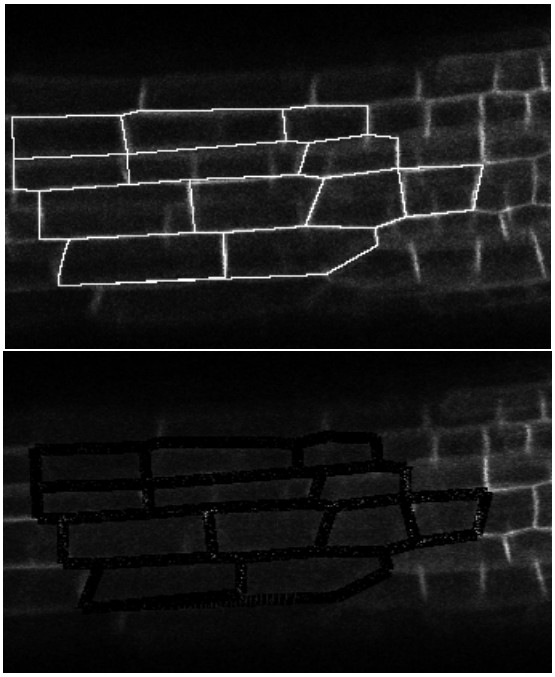

Figure 7. Top: manually initialised cell network. Bottom: difference between I_5 and synthesis of the most probable cell network.

References

[1] J. Barron, D. Fleet, and S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 1994.

[2] A. G. Bengough, M. F. Bransby, J. Hans, S. J. McKenna, T. J. Roberts, and T. A. Valentine. Root responses to soil physical conditions: growth dynamics from field to cell. *Journal of Experimental Botany, Plasticity Special Issue*, 2006.

[3] M. J. Black and P. Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer Vision and Image Understanding*, 63(1):75–104, Jan 1996.

[4] J. M. Coughlan, D. Snow, C. English, and A. Yuille. Efficient deformable template detection and localization without user initialization. *Computer Vision and Image Understanding*, 78:303–319, 2000.

[5] B. McCane, K. Novins, D. Crannitch, and B. Galvin. On benchmarking optical flow. *Computer Vision and Image Understanding*, 84(1):126–143, 2001.

[6] D. J. Stephens and V. J. Allan. Light microscopy techniques for live cell imaging. *Science*, 300(5616):82 – 86, 2003.