

Estimating the Motion of Plant Root Cells from *in vivo* Confocal Laser Scanning Microscopy Images

Timothy J. Roberts, Stephen J. McKenna, Cheng-Jin Du
School of Computing, University of Dundee, Dundee DD1 4HN, UK

Nathalie Wuyts, Tracy A. Valentine, A. Glyn Bengough
Scottish Crop Research Institute, Dundee DD2 5DA, UK

June 24, 2009

This is an author-created version of a paper to be published in the Springer journal Machine Vision and Applications. The published version will be available at www.springerlink.com)

Abstract

Images of cellular structures in growing plant roots acquired using confocal laser scanning microscopy (CLSM) have some unusual properties that make motion estimation challenging. These include multiple motions, non-Gaussian noise and large regions with little spatial structure. In this paper, a method for motion estimation is described that uses a robust multi-frame likelihood model and a technique for estimating uncertainty. An efficient region-based matching approach was used followed by a forward projection method. Over small timescales the dynamics are simple (approximately locally constant) and the change in appearance small. Therefore a constant local velocity model is used and the MAP estimate of the joint probability over a set of frames is recovered. Occurrences of multiple modes in the posterior are detected, and in the case of a single dominant mode, motion is inferred using Laplace's method. The method was applied to several *Arabidopsis thaliana* root growth sequences with varying levels of success. In addition, comparative results are given for three alternative motion estimation approaches, the Kanade-Lucas-Tomasi tracker, Black and Anandan's robust smoothing method, and Markov random field based methods.

1 Introduction

The dynamic processes of cell growth and expansion are of fundamental importance to our understanding of plant function and morphogenesis. However, the

growth of cells within plants and the motion of organelles within cells is still not well understood. Tracking in time-lapse microscopy images has been employed in the identification of morphological traits in many plant cell cultures, such as barley [25], single suspension cells of carrot (*Daucus carota* L.) [36], and leaf explants from *Dactylis glomerata* L. [32]. Moreover, images have usually been captured by traditional microscopy and cell tracking has been performed manually. Some commercial image analysis tools are available such as Volocity from Improvision [15] and Metamorph from Universal Imaging Corporation [37], but their capabilities for automatic plant cell tracking are limited.

The innovations of confocal laser scanning microscopy (CLSM) and fluorescent proteins have enabled new research perspectives on the study of cellular dynamics within the living plant at submicron resolutions [10]. CLSM bridges the gap between lower resolution wide field/bright field microscopy which allows the imaging of live tissue but with little cellular detail and electron microscopy, where cellular structure can be imaged but with the disadvantage that samples cannot be imaged alive. CLSM combined with fluorescence labelling can provide rich, high resolution spatiotemporal data on the structure and dynamics of live biological specimens. When combined with targeted, fluorescent markers such as green fluorescent protein (GFP), specific cell structures can be isolated, providing rich data for studying the morphology and physiology of living plant systems, e.g. [11, 12].

There have been investigations into estimating growth using other (external) imaging modalities. van der Weele *et al.* [38] applied low-level motion estimation methods to non-fluorescence microscope images of the external surface of *A. thaliana* roots (without the application of graphite particles). Walter *et al.* [39] tracked maize root surfaces using infra-red illumination. Barron and Liptay [3] used differential flow from multi-view optical imaging and near infra-red imaging to recover the growth of maize seedlings. Recently, machine vision has also been applied to detecting and measuring roots in minirhizotron images captured by a miniaturized color camera on a telescopic handle [41].

This paper focuses on automatic analysis of the motion of cellular structure in *Arabidopsis thaliana* roots (a model system in plant sciences) captured by confocal laser scanning microscopy (CLSM). An example of an *A. thaliana* root image produced using CLSM is shown in Fig. 1. This specimen expresses a single fluorescence marker targeted to the plasma membrane (LTI-eGFP) [18].

Success for this application would allow large quantities of confocal data for a variety of configurations to be analysed thus improving our ability to understand plant biological processes. An accurate, automatic, and quantitative motion estimation method would facilitate controlled experiments to investigate the effects of environmental conditions on root growth responses [4]. One long-term goal is to study the coordination of growth between cells in the *A. thaliana* root tissue, and specifically which cell types control root elongation. This biological question could be linked to cellular mechanisms controlling the direction of cell expansion, via the cytoskeleton and specifically microtubular orientation.

The image data of growing *A. thaliana* roots acquired using CLSM have some unusual properties that make automated low-level estimation of the mo-




Figure 1: An example of an *Arabidopsis thaliana* root image captured using CLSM at 10 \times magnification.

tion of cellular structures challenging. Firstly, the images are formed from the interaction of laser illumination with sparse, 3D cellular markers. The resulting sparse images often contain little stable texture in many regions and large amounts of impulsive noise, and can be affected by illumination changes and thresholding effects. Secondly, cells form a complex structure, where neighboring cells are similar to each other while having small but important differences. Thirdly, automatic detection of cell boundaries is complicated by blurring in depth, leading to boundaries from other layers being visible, and significant clutter in the datasets. Fourthly, the apparent motion is the result of a combination of non-rigid cell growth and its interaction with the growth medium, motion of fluorescent protein not yet at its targeted position within the cell, and occasionally global translation of the whole specimen.

There are broadly two main approaches to generic motion estimation in computer vision: local approaches where motion estimates are based upon data close to the point to be estimated, and global approaches whereby motion estimates are based upon data from the image as a whole. Previous work on automatic analysis of plant cells from confocal data has focussed mainly on local optical flow methods [4]. 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. 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 [5]. However, due to the lack of persistent local features, it is difficult to obtain accurate cell growth measurements using such general purpose optical flow methods. Therefore, it is important, especially in applications such as that considered in this paper, to estimate some measure of certainty of the motion estimates.

The aim of this paper is to develop a local region-based algorithm incorporating a robust likelihood, a constrained temporal model, and a suitable technique

for estimating uncertainty. The likelihood function is multi-modal. Similarly to the method proposed by Singh [31], modes were found using efficient region-based matching followed by a forward projection method, the use of which is novel in this context. Since the dynamics of the root, over the small timescales investigated here, are constrained and simple (compared to the variability in appearance), estimation is based on multiple frames to reduce gross errors and improve accuracy. By re-weighting the modes of the likelihood function according to a temporal prior, a more accurate estimation of uncertainty can be obtained.

The performance of the proposed motion estimation method is compared to that of the Kanade-Lucas-Tomasi (KLT) tracker [1, 23], Black and Anandan’s robust smoothing method [5], and Markov random field (MRF) based methods. Empirical evidence presented here suggests that the proposed method improves upon these methods for this application.

The remainder of this paper is organized as follows. In Section 2, the proposed method is presented in detail followed by brief descriptions of the other methods to which it is compared in Section 3. The method is then evaluated by applying it to various sequences and comparative results are reported in Section 4. Finally, conclusions and recommendations for future work are given in Section 5.

2 Robust Motion Estimation with Uncertainty

2.1 Overview

An algorithm with three steps is presented for motion estimation of plant root cells. In the first step, the probabilities of displacements whose horizontal and vertical components are integers are considered for each pixel location. In other words, displacements on a pixel grid are explored. Secondly, those pixels at which local motion might be estimated with reasonable certainty are selected. These correspond to the pixels at which one specific motion has high probability. At these locations, a forward projection method is applied to robustly recover subpixel motion. Thirdly, 2D motion estimates are computed at these locations, along with their uncertainty. This is achieved using the Laplace method with centred difference approximation.

2.2 Motion Model

Inspection of the data showed that without accounting for significant changes in appearance, it is only possible to track over a few frames, between two and five in typical datasets. As mentioned earlier, the dynamics of the root over these small timescales are constrained and simple so we can make use of multiple frames to reduce gross errors and improve accuracy. Motion estimates are not propagated forward sequentially as in many traditional methods such as KLT [1]. Rather, the goal of the algorithm is inferring a joint posterior over a sequence of

frames. The motion is modelled as locally translational with a constant velocity prior. This is a representative model of motion encountered in biology [29]. For such time periods, a uniform prior over a predefined range of velocities was used. Cells grow continuously over time with a finite maximum velocity. A reasonable upper bound, N , on the displacement of a pixel from one frame to the next can be estimated visually fairly easily.

2.3 Likelihood

Once a region-based approach has been adopted, one important aspect in accurately establishing correspondence across a plant root sequence is how to measure the similarity between image regions at different times. The most widely used dissimilarity measurement between regions is the sum of squared difference (SSD). This arises from the use of a Gaussian likelihood function. However, a Gaussian likelihood model is not appropriate for this application.

The complex form of CLSM pixel noise is shown in Fig. 2. The quantiles of the empirical distribution of noise (from fast successive scans of the same root portion) were aligned with the quantiles of the normal distribution. If the noise was Gaussian the plot in Fig. 2(b) 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.

Such empirical observation of the form of the image noise motivated the use of a non-Gaussian likelihood function of a motion \mathbf{v} . We begin by defining the likelihood of an integer displacement and subsequently describe a method for handling sub-pixel displacements. Assuming independence between local pixels given the motion \mathbf{v} at pixel \mathbf{x} , the negative log-likelihood function is defined (up to a constant scaling factor) as

$$-\log(p(I_{1:T}|\mathbf{v})) \propto \sum_{t=1}^T \sum_{(i,j) \in \Omega} K(i,j)D(t, \mathbf{v}, i, j) \quad (1)$$

where the dissimilarity measurement is a truncated quadratic:

$$D(t, \mathbf{v}, i, j) = \begin{cases} [I_1(\mathbf{x} + (i, j)) - I_t(\mathbf{x} + (t-1)\mathbf{v} + (i, j))]^2 & \text{if } D(t, \mathbf{v}, i, j) \leq \gamma \\ \gamma & \text{otherwise} \end{cases} \quad (2)$$


which is based on the computation of the distance $D(t, \mathbf{v}, i, j)$ between the neighbourhood of pixel $\mathbf{x} = (x, y)$ in the first frame I_1 and the neighbourhood of pixel $\mathbf{x} + (t-1)\mathbf{v}$ from the t^{th} frame I_t . The cut-off γ was set empirically to 3 times the standard deviation. Note that no normalising constant needs to be specified for this 'likelihood' for what follows. Various other robust norms were tried with similar results.

The sum in Eq. (1) is taken over several consecutive frames (typically 3) and a circular region $\Omega = \{(i, j) | \sqrt{i^2 + j^2} < R\}$. Kernel K is linearly weighted and



(a)

Normal Q-Q Plot



(b)

Figure 2: (a) The (scaled) absolute difference between two scans taken in rapid succession so that growth is negligible. (b) A normal probability plot showing the non-Gaussian nature of the noise (formed from the difference between two close scans). The form of the curve suggests three components which are demarcated by the vertical lines on the plot: approximately Gaussian background noise (central linear section, corresponding to the majority of the pixels), approximately Gaussian ‘biological noise’ (larger linear sections) and outliers.

has maximum value at the centre of the region and zero value at its boundary. R is the radius.

A forward projection approach [13] was used to determine the likelihood of subpixel displacements. In contrast to the usual approach of interpolating the image, this more principled approach uses the correct amount of data from the image and reduces the bias issues associated with interpolation based methods which make subpixel motions less probable than whole pixel motions [30]. For sub-pixel estimation, interpolation between the robust functionals with weights proportional to the overlapping pixel regions was performed (see Fig. 3). The final weight of each pixel $\omega(m, n)$ accumulates the sum of the weights of subpixels

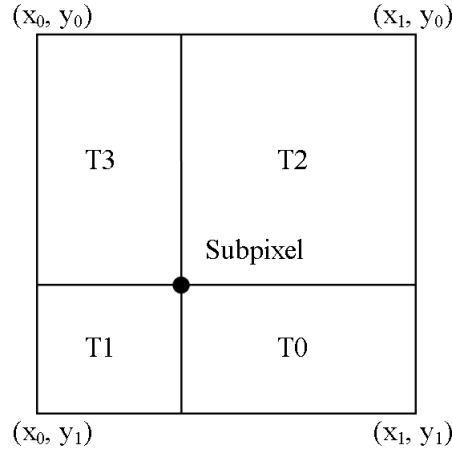


Figure 3: Illustration for the calculation of weight $A(i, j)$, where $x_{0/1}$ and $y_{0/1}$ denote image pixel coordinates, $A(x_0, y_0) \propto \text{area of T0}$, $A(x_1, y_0) \propto \text{area of T1}$, $A(x_0, y_1) \propto \text{area of T2}$, and $A(x_1, y_1) \propto \text{area of T3}$.

$$\omega(m, n) = \sum_i \sum_j A(m, n; i, j) \quad (3)$$

In computing forward projection, both the kernel value $\widehat{K}(m, n)$ and the intensity value $\widehat{I}_1(\mathbf{x} + (i, j))$ of a pixel within the circular region Ω are weighted sums as follows:

$$\widehat{K}(m, n) = \sum_i \sum_j A(m, n; i, j) K(i, j) \quad (4)$$

$$\widehat{I}_1(\mathbf{x} + (m, n)) = \sum_i \sum_j A(m, n; i, j) I_1(\mathbf{x} + (i, j)) \quad (5)$$

In order to consider subpixel accuracy, the log-likelihood function in Eq. (1) is re-written in the following form:

$$-\log(p(I_{1:T}|\widehat{\mathbf{v}})) \propto \sum_t \sum_m \sum_n \omega(m, n) \widehat{K}(m, n) \widehat{D}(t, \mathbf{v}, m, n) \quad (6)$$

where $\widehat{D}(t, \mathbf{v}, m, n)$ is obtained by substituting \widehat{I}_1 of Eq. (5) in place of I_1 in Eq. (2).

2.4 Inference

For computational efficiency, we begin by evaluating an integer displacement grid to quickly determine a set of interesting points (points with high probability of a specific motion), where displacements are assumed to be a wide range of integer pixels/frame, with $\sqrt{v_x^2 + v_y^2} \leq N$. The fact that the negative log-likelihood never increases as terms are included in its sum is used to speed up the search by terminating likelihood evaluations early when they reach a sufficiently low value for a particular displacement. This value is an empirically determined constant that determines the minimum likelihood value needed for a motion estimation to proceed. If two significantly different motions have a high likelihood this indicates multi-modality, in which case that point is labelled as having ambiguous visual motion.

For all the interesting points found, a second step is conducted to robustly recover subpixel motion. As detailed in Section 2.3, a forward projection method similar to that presented by Harauz and Ottensmeyer [13] is applied for estimating accurate subpixel motion. The sampled motion \mathbf{v}_0 with maximum posterior is found for each pixel not labelled as having ambiguous visual motion. These estimates are refined by computing a local expected value:

$$\mathbf{v} = \frac{\sum_{\Delta_x} \sum_{\Delta_y} p(\mathbf{v}_0 + h(\Delta_x, \Delta_y) | I_{1:T})(\mathbf{v}_0 + h(\Delta_x, \Delta_y))}{\sum_{\Delta_x} \sum_{\Delta_y} p(\mathbf{v}_0 + h(\Delta_x, \Delta_y) | I_{1:T})} \quad (7)$$

where h is a subpixel displacement step, and was set to 0.1 in the experiments reported here.

The Laplace method [24] is applied to estimate uncertainties in these motion estimates, approximating the posterior $p(\mathbf{v} | I_{1:T})$ as Gaussian. The posterior can be expressed as

$$p(\mathbf{v} | I_{1:T}) = p^*(\mathbf{v}) / C \quad (8)$$

where $C \equiv \int p^*(\mathbf{v}) d\mathbf{v}$ is a normalization constant. Using a Taylor expansion of the logarithm of $p^*(\mathbf{v})$ centred around the maximum at \mathbf{v}_0 , we have

$$\log p^*(\mathbf{v}) \simeq \log p^*(\mathbf{v}_0) - \frac{1}{2}(\mathbf{v} - \mathbf{v}_0)^T \mathbf{A}(\mathbf{v} - \mathbf{v}_0) \quad (9)$$

where \mathbf{A} is a 2×2 Hessian matrix defined by

$$A_{ij} = -\frac{\partial^2}{\partial v_i \partial v_j} \log p^*(\mathbf{v}) |_{\mathbf{v}=\mathbf{v}_0} \quad (10)$$

Thus, the posterior can be approximated by the following Gaussian

$$p(\mathbf{v}) = \frac{1}{2\pi |\mathbf{A}|^{-1/2}} \exp\left(-\frac{1}{2}(\mathbf{v} - \mathbf{v}_0)^T \mathbf{A}(\mathbf{v} - \mathbf{v}_0)\right) \quad (11)$$

The centred difference approximation with fourth order error [19] is employed to compute the Hessian matrix \mathbf{A} . For example, the second derivatives A_{11} and A_{22} can be obtained by

$$\begin{cases} A_{11} = (-p(\mathbf{v}_0 + (2h, 0)|I_{1:T}) + 16p(\mathbf{v}_0 + (h, 0)|I_{1:T}) - 30p(\mathbf{v}_0|I_{1:T}) + \\ \quad + 16p(\mathbf{v}_0 + (-h, 0)|I_{1:T}) - p(\mathbf{v}_0 + (-2h, 0)|I_{1:T}))/12h^2 \\ A_{22} = (-p(\mathbf{v}_0 + (0, 2h)|I_{1:T}) + 16p(\mathbf{v}_0 + (0, h)|I_{1:T}) - 30p(\mathbf{v}_0|I_{1:T}) + \\ \quad + 16p(\mathbf{v}_0 + (0, -h)|I_{1:T}) - p(\mathbf{v}_0 + (0, -2h)|I_{1:T}))/12h^2 \end{cases} \quad (12)$$

The covariance matrix $\Sigma = \begin{bmatrix} \sigma_x^2 & \sigma_{xy} \\ \sigma_{xy} & \sigma_y^2 \end{bmatrix}$ is given by \mathbf{A}^{-1} .

2.5 Implementation

The overall pseudo-code for the motion estimation algorithm for each pixel is given in Table 1. The algorithm was implemented in Java under Windows XP on an *Intel*^R Core2 processor, 2.4 GHz machine, and interfaced to ImageJ, the Java image processing program developed by Wayne Rasband at the United States National Institutes of Health and available at <http://rsb.info.nih.gov/ij>. An ultimate aim of this work is to deliver an open-source workbench application capable of accurate analysis of large amounts of confocal data. In this application speed is not critical but the method must be feasible for deployment to standard PC workstations.

3 Other Methods

The proposed method was compared to three other algorithms widely used for motion estimation in the literature. These methods are now outlined for completeness.

As a local feature tracking method, KLT [1, 23] is a straightforward and popular approach to motion estimation. Since this method, in the form investigated here, uses only spatiotemporally localised data, the method is quite simple, popular and general. It was used here as a baseline for comparing the methods. Birchfield's publicly available implementation of the KLT method (v1.3.2) was used in this study.

The use of global spatial smoothing to remove local ambiguities and improve accuracy has a long history in optical flow estimation [14]. The second method, which is due to Black and Anandan [5], is based upon robust spatial smoothing that preserves motion discontinuities. In this method the equations resulting from the minimisation of both a robust function on the image residuals and a robust spatial smoothing function are solved using an iterative, multi-scale, graduated non-convexity algorithm, a key idea in which is approximation of the non-convex robust functions with convex functions.

Markov random fields can be used to elegantly model many low-level vision tasks that require global spatiotemporal context such as motion estimation, stereo disparity and segmentation. They emphasise a discrete and probabilistic

<ol style="list-style-type: none"> 1. Initialise the parameters for likelihood calculation <ol style="list-style-type: none"> (a) Extract a circular region centered at the pixel and obtain the intensity value of each pixel within the region (b) Calculate a linearly weighted kernel K having maximum value at the center and zero at the boundary 2. Explore the posterior over integer displacements <ol style="list-style-type: none"> (a) Calculate the motion likelihoods for the sequence (b) Find the maximum motion likelihood in the region 3. If the maximum likelihood $>$ given likelihood threshold value <ol style="list-style-type: none"> (a) Introduce a weight parameter ω to account for subpixel accuracy using a forward projection method (b) Calculate the weighted intensity values and kernel K of the region (c) Repeat Step 2 to find the maximum likelihood and its corresponding motion with subpixel displacement 4. Estimate the motion as local expectations and their corresponding covariance matrices (uncertainties) <ol style="list-style-type: none"> (a) Estimate the pixel motions by the sum of neighbourhood motions weighted by their posteriors (b) Calculate the corresponding covariance matrix based on the Laplace method and the centered difference approximation

Table 1: The overall structure of the proposed algorithm for motion estimation.

approach. MRFs are discussed in detail in Li’s book [21]. The key problem in employing MRFs, which model the overall probability of a configuration using a cyclic network of local functions, is solving the resulting global and high dimensional inference problem. A selection of techniques are discussed for this purpose, in particular graph cuts (GC) and loopy belief propagation (BP), which have been reported to give robust and high performance inference for MRF models in many recent papers [6, 17, 9, 35]. We have investigated MRFs with a range of data and smoothing energy functions. In particular, the following data energy functions were tried: quadratic (r^2), McClure Geman ($r^2/(s^2 + r^2)$), and truncated quadratic (r^2 or c^2 if $r > c$ with c constant). The following smoothing functions were tried: linear, truncated linear and quadratic, and McClure Geman with a range of different smoothing strengths. GC [17] and loopy BP [9] were used to solve the resulting inference problems. These methods were evaluated based upon the publicly available Middlebury MRF energy minimisation code (v1.6) [34], which contains algorithms contributed by various authors. Discrete (i.e. not subpixel) velocity labels were used.

4 Experimental Results

4.1 Plant Preparation and Image Acquisition

Arabidopsis thaliana roots transformed to express eGFP targeted to the plasma membrane (PM) were grown *in vitro* in $0.5\times$ Murashige and Skoog nutrient media, solidified with 0.7 phytoagar [18]. Images of their roots were acquired on a Leica TCS SP1 confocal microscope using a $10\times$ (NA 0.3) objective lens at medium scan speed (400MHz). Excitation at 488nm of the fluorochrome probe was achieved with a 20mW Argon laser light source. Fluorescence emission was detected using a 500 – 530nm filter. Two-dimensional image layers at fixed depth were captured as time series at high resolution of 1024×1024 with 8 bits per pixel.

The time interval between acquisition of images was relatively large (60s) and inter-frame motion was correspondingly large making motion estimation challenging. An upper bound on velocity for all test sequences was manually estimated as 8 pixels per frame in any direction. In the portion of root being studied, cell expansion was of the order of 1 pixel per 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.

4.2 Visualisation

A colormap with 256 entries is defined to visualize the motion estimation results in Figs. 4 and 7. The data values are mapped to colors linearly between the minimum and the maximum values, which represent the first and last color in the colormap. For the horizontal and vertical components of motions, signed data

are used, which are mapped to colors from blue through aqua and yellow to red. Blue and aqua represent negative movements, while yellow and red represent positive movements. Regions that have a speed close to zero are in dark aqua or yellow, i.e. black appearance. Blue and red display the fastest negative and positive movements, respectively. In the case that there are less than one quarter of motion points having either negative or positive movement, only one color is used to show that movement. A similar scheme is used to display the covariance, where black corresponds to no correlation. The standard deviations of motion σ_x, σ_y are always positive, and are displayed in the same four colors. Colors from blue through aqua and yellow to red indicate increasing variance of the motions estimated. Points where velocity is not estimated are shown in white. In order to make the results comparable, grey scale maps are used for the figures presented in Section 4.4.

4.3 Motion Estimation Using the Proposed Method

The algorithm described in Section 2 was applied to a large variety of plant root data sets with varying levels of success. Figure 4 shows an example of a plant root (8 days after seed germination), its estimated horizontal and vertical components of motion, their standard deviations and covariance. Cell expansion is one of the key parameters that determines the ultimate form of plant organs [28]. How far a cell expands and in what direction are the two primary determinants of a cell’s final shape. In Fig. 4(c), it is shown that the root tip moves at the largest speed (around 2.4 pixels per frame). Note that without sub-pixel motion estimation, only three values (0, 1, and 2 pixels) would be present in Fig. 4(c). Regions located at positions progressively distant from the tip move at progressively lower speeds, until the end of the growth zone is reached, where the speed becomes zero. The cumulative elongation of cells in the growth zone results in the root tip being propelled at the maximum speed. This is consistent with the reports of other researchers [38, 4]. The transverse expansion is relatively slow with maximum velocity around 0.78pixels per frame (Fig. 4(b)). The use of Laplace’s method enables uncertainty to be quantified. In Figs. 4(d) and (e), higher uncertainty for horizontal motion is obtained as compared with that for vertical motion. Fig. 4(f) shows the covariance of horizontal and vertical motion estimated. The small values of covariance suggest that the correlation between horizontal and vertical motion is weak.

A quantitative comparison with manual motion estimation was performed. A plant biologist selected 31 points on a root image at which he was confident he could provide motion estimates. These are shown in Fig. 5. The same biologist then marked up corresponding points on subsequent images in the sequence. This root was growing in a vertical direction so the vertical component of motion was considered here. Fig. 6 plots the estimates obtained using this manual method against those obtained using the proposed automatic method. The manual estimates were more heavily quantised because the locations were annotated to the nearest pixel. Manual estimates were

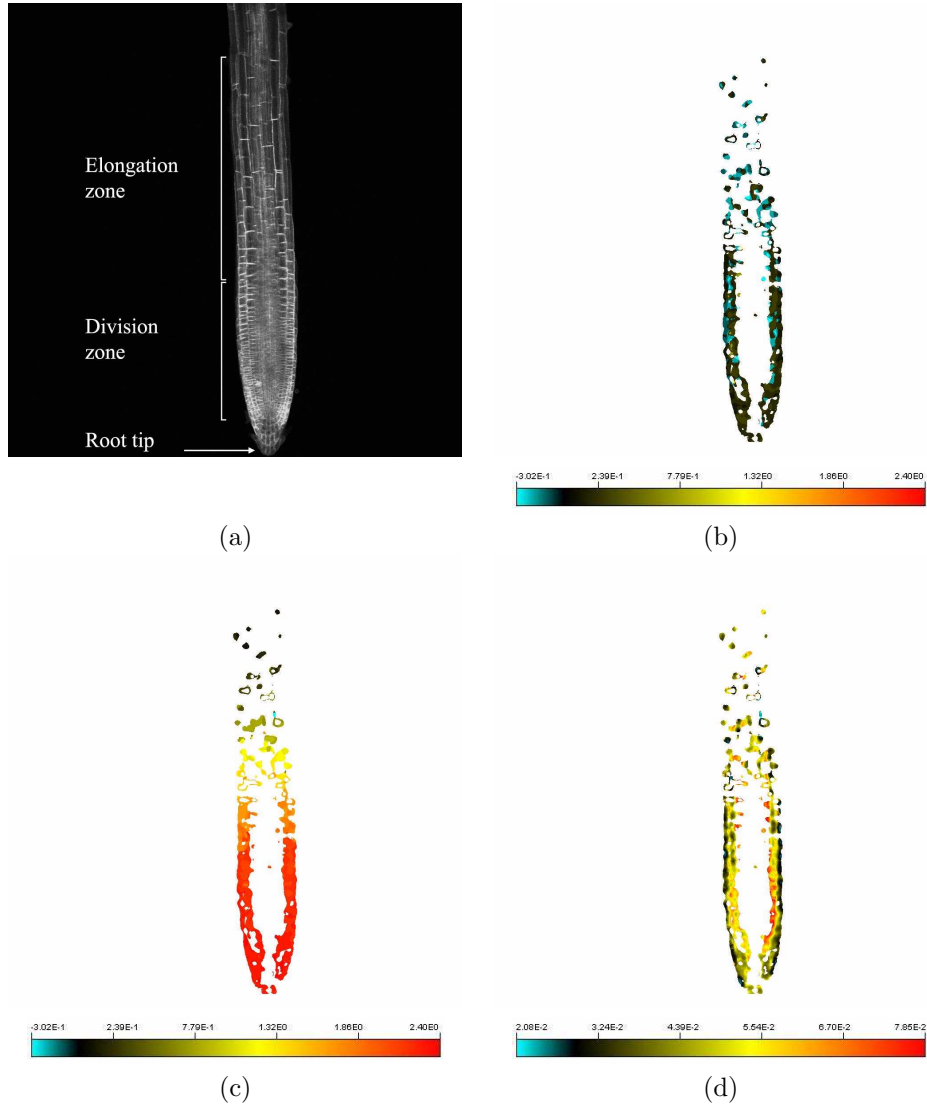


Figure 4: Estimated motion in a time-lapse sequence of CLSM images of a growing *Arabidopsis thaliana* root (8 days after seed germination): (a) confocal image (first frame in the sequence); (b) estimated horizontal motion (pixels/frame); (c) estimated vertical motion (pixels/frame); (d) standard deviation of horizontal motion component; (e) standard deviation of vertical motion component; (f) covariance of horizontal and vertical motion components. (This figure should be viewed in colour).

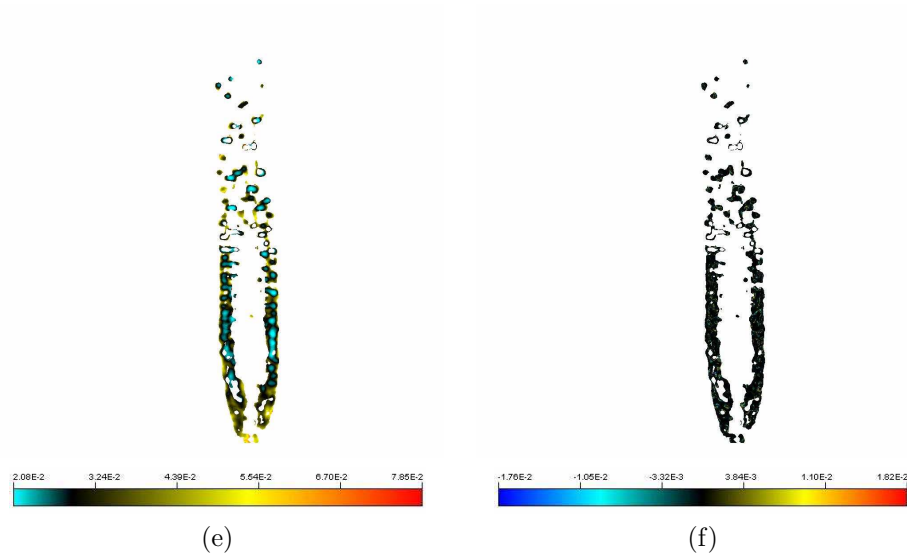


Fig. 4 (continued)




Figure 5: The first image in a sequence used to compare manual and automatic estimation. Estimates were compared at the annotated points.


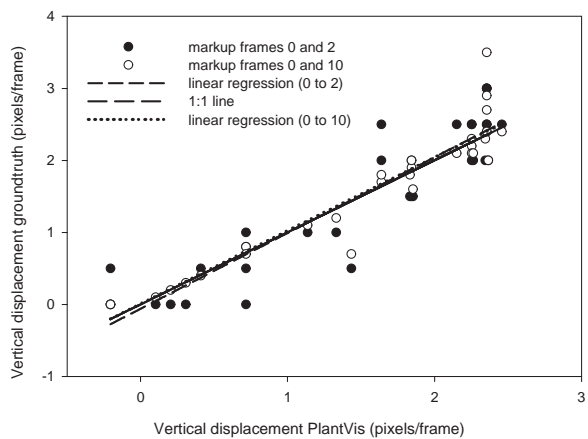


Figure 6: Top: Estimates of vertical motion obtained at the points annotated in Fig. 5 using the manual and automatic methods. Bottom: A Bland-Altman plot comparing the two methods.

obtained over temporal intervals of two frames and ten frames. The former was the same interval used by the automated method. The ten-frame interval provided less heavily quantised estimates but made a stronger assumption about constancy of velocity. Fig. 6 shows linear regression fits and demonstrates a strong linear relationship between the methods; in fact the gradients of these lines are within one standard error of 1.0. Figure 6 also shows a Bland-Altman plot comparing the manual estimates with the automated estimates (both estimated over a 2-frame interval). It plots the average of the two methods against the difference between them. The central line is the bias (-0.05). The upper and lower lines are the *limits of agreement* (i.e. the bias plus and minus two standard deviations of the difference). The standard error was 0.077. The methods always disagreed by less than one pixel per frame.

The motion estimation results for three other time-lapse image sequences of growing *A. thaliana* roots are shown in Fig. 7. Besides the growing direction, the ages of these roots are different from the root in Fig. 4. The image sequences used in the first, second, and third columns of Fig. 7 are of roots 5, 6 and 12 days after seed germination, respectively. Under standardized growing conditions, root growth rate increases during seedling development. Tropic reactions which make the root bend from its normal direction of growth may negatively affect root growth rate. The motion estimation results obtained are quite satisfactory as shown in the second, third, fourth, and fifth rows of Fig. 7, which indicate that the algorithm reflects root growth at different times during development and different growing directions.

However, some regions contain errors for various reasons. In Fig. 4(c), it can be observed that the motion of one region near the end of the elongation zone is wrongly estimated as being negative. This error appears to be caused by out-of-plane motion due to nutation or gravitropic reaction which has resulted in cell membranes from other layers being visible. One of the most important factors affecting the motion estimation performance is the image quality. Because of the inherent imperfections of the image formation process, CLSM images are often corrupted by intensity variations manifesting themselves via spurious intensity variations not present in the original scene [22], usually referred to as shading or intensity inhomogeneity. Despite the fluorophore being expressed from a constitutive promoter, some areas within the root image appear brighter than others, due to a combination of slight differences in expression levels and shading effects [26]. The shading effect can usually be attributed to several known factors, including fluorescent attenuation along the confocal (depth) axis, image acquisition factors, variations of illumination exposure rate, spatially uneven distribution of the fluorophore and the spatial characteristics of illumination beams, and fluorochrome micro-environment, e.g. pH, temperature, and embedding medium. Comparing the image in the second row with the ones in the first and third rows, the motion estimates in the division zone near the root tip are very sparse. This sparsity was caused by the sequence having been acquired at a lower gain (to avoid clipping). Furthermore, since cell movements are in



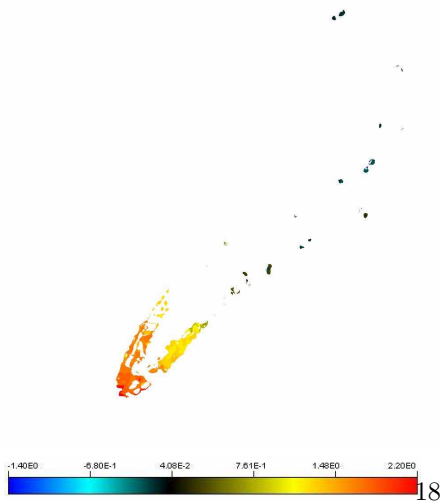






Figure 7: Estimated motion in time-lapse sequences of CLSM images of three different growing *Arabidopsis thaliana* roots. Column (a) shows the first confocal image in each sequence. The corresponding estimation results for horizontal and vertical motions and their standard deviations and covariances are shown in columns (b) - (f), respectively. (This figure should be viewed in colour).



(c)

(d)

Fig. 7 (continued)




Fig. 7 (continued)

3D, the failure in some regions with different motion types is unavoidable due to the use of 2D analysis in this work, despite acquisition being arranged so that out of plane motion is kept to a minimum.

The velocity profile, or rate of displacement per unit of time along the root axis, is a commonly used tool to visualize and quantify the difference in expansion rate along the root axis. Fig. 8 shows an example of the projection of multiframe motion estimation data to a spline fitted to the central axis of the root, resulting in the velocity profile. The reference for distance is located in the root tip (typically at the quiescent centre). The root tip has the highest velocity, while at a certain distance from the quiescent centre, velocity decreases abruptly and significantly to approach zero at the end of the growing zone. The results agree with previous investigations using high resolution images captured by non-fluorescence microscopy [38].

Particular cell regions (tissues, individual cells and cell membranes) can also be selected and tracked allowing the study of small changes in cell expansion rate with a temporal resolution of less than minutes. This represents an exciting tool for plant physiology studies.

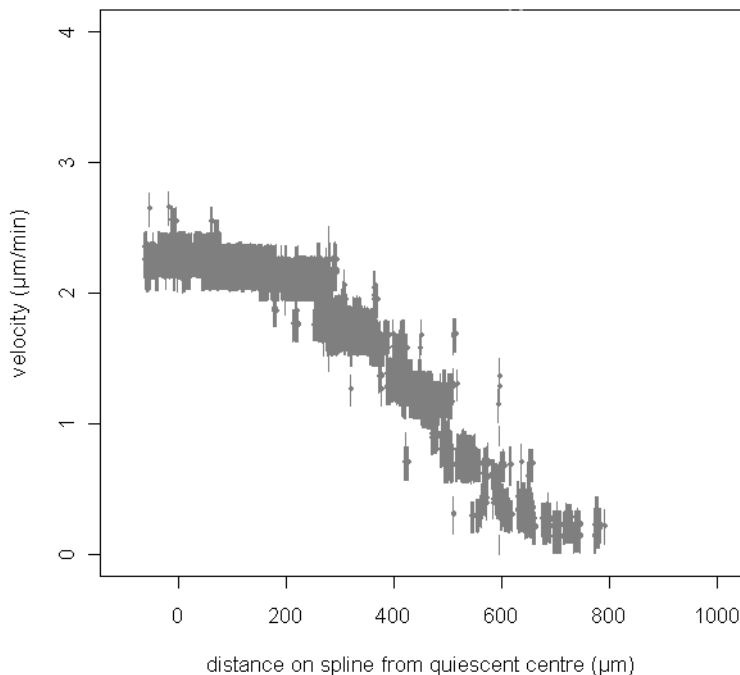


Figure 8: Velocity as a function of distance along the root from a reference in the root tip (the quiescent centre). Motion was estimated in a time-lapse sequence of CLSM images of a growing *Arabidopsis thaliana* root (8 days after seed germination) (see Fig. 4).

Estimating motion from an image sequence can require an hour or longer of processing time. However, in this study, speed of inference was not a primary concern. The focus was on an algorithm that could reliably infer local motion in large datasets and give useful measures of uncertainty.

4.4 Comparison with Other Methods

A dataset (see Fig. 9) that contains a mixture of root growth and global translation was selected to compare the performance of the different methods. Results are presented after subjectively tuning the parameters to their reported optimal values (automatic parameter learning being non-trivial). Due to the lack of ground truth we compared the methods directly.






Figure 9: The first frame in the sequence used for comparative study.

Results from the KLT method are shown in Fig. 10. The default energy functions and parameters were used (linearized quadratic) with 200 feature tracks requested. Its performance deteriorated dramatically when the region window size was below 11×11 pixels (for which results are shown). Fig. 10(a) is a quiver plot of the sparse motion estimated. The method fails to tell that the largest velocity should be around the root tip area. The motion speeds and directions are estimated wrongly for the cells at the end of elongation zone and the differentiation zone. The top row of Fig. 10(b) shows that the motion estimated is not accurate in the root tip region where confusion between the small, similar



(a)



(b)


Figure 10: Results from the Kanade-Lucas-Tomasi method: (a) a scaled quiver plot of the sparse motion; (b) two zoomed views illustrating some typical errors. The top row shows gross errors in the root tip area where cells are small and similar. Ideally, the white points showing motion occurrence should be located on the cell membranes. The bottom row illustrates the detection of features on less stable structures in the elongation zone.

cells is considerable. Ideally, the white points showing motion occurrence should be located on the cell membranes. However, around 10% of the feature points were assigned to intra-cellular regions where the appearance is more variable (as compared to cell junctions) and the associated motions are more often incorrect. A more global method will be necessary to remove such ambiguities. As shown in the bottom row of Fig. 10(b), worse results were obtained when applying the KLT method to the elongation zone with less stable structures.


Fig. 11 shows the results obtained by using Black and Anandan’s robust smoothing method. The third smoothing level is shown, with parameters close to their default and recommended values ($l_1=10, l_2=1, S_1=10.0, S_2=1.0, s_1=5.0, s_2=0.05$). As one might expect, the performance of this method with these parameters agrees with the previous spatially localised method around highly informative data (e.g. the region 1/3 up from the root tip). However, in regions of ambiguity, such as the root tip, the low-level spatial constraints fail to provide any significant advantages.

Reliably fitting MRFs to our data using those algorithms described in Section 3 proved difficult. The algorithms were quite sensitive to the parameters used, often providing clearly inaccurate results that were often inconsistent between algorithms. Partially successful results (using the regularised quadratic image noise and linear smoothing function) for expansion and swap GC algorithms are shown in Fig. 12. These best results were achieved after extensive manual parameter tuning on this data set. Lower velocity estimates in the root tip region are incorrect but the results are better than Black and Anandan’s method in this region, which suggests that the linear smoothing function is more appropriate for these data. However, in comparison with that method, smoothing elsewhere provides inaccurate estimates. Loopy BP in the form of max-product (Middlebury and in-house implementation) and min-sum (in-house implementation) also produced poor results and had prohibitively high memory requirements (a recognised problem with BP that has drawn some interest recently, e.g. [16, 9]). We note that all methods were able to produce accurate results when applied to synthetic transformations (i.e. root images with simple translation) with a wide range of energies/parameters. The final algorithm in the Middlebury suite was iterated conditional modes (ICM) and, as would be expected, this gave very poor results due to getting trapped in local minima early on (with energy values much higher than graph cuts).

The results obtained by the developed method using 2 or 3 frames with an 11 pixel circular window are shown in Fig. 13. It can be clearly seen that using a constrained motion model over 3 frames reduces the gross errors present in the above results. Using even a single extra frame with a constrained motion model can increase motion estimation certainties as can be seen by comparing the images in Fig. 13(c). Gross motion estimation errors are reduced as can be seen by inspecting the inset magnified region in Fig. 13(a) and comparing it to the corresponding region in Fig. 13(b).




(a)



(b)

Figure 11: Results from Black and Anandan's robust smoothing method: (a) and (b) show horizontal and vertical components of flow scaled to be comparable with the other results. The method is accurate around informative structures but low-level spatial smoothing provides little advantage, sometimes even reducing accuracy due to localised over-smoothing, compared with local methods. (c) shows violations of the spatial coherence model.



(c)


Fig. 11 (continued)

4.5 Discussion


The algorithm presented in Section 2 was motivated by the study of the dynamic behaviors of plant cell growth and expansion. The results obtained by the software are now discussed with some biological background and interpretation.

Root elongation or growth is the result of a combination of cell division and cell expansion. The division zone of the Arabidopsis comprises of the quiescent centre surrounded by root initial cells. These cells are generally isodiametric in shape. Within the division zone, the root initials undergo several rounds of very tightly controlled division only expanding slightly in size. Hence the horizontal motion of cells shown in Fig. 4(b) exhibits no clear dominant motion. As the distance between the cells and the quiescent centre increases, the cells begin to rapidly expand longitudinally resulting in cells which are longer than their cylindrical diameter. The vertical motion estimated in Fig. 4(c) is mainly attributed to longitudinal expansion, while only minor contribution is from cell division. Therefore, low uncertainty was estimated as shown in Fig. 4(e). Cells undergo further development in the differentiation zone where some cells may expand in alternative directions to become root hair cells and highly oriented in the longitudinal axis. At this stage, both longitudinal and radial components of motion are approximately zero, as correctly estimated by the algorithm as illustrated in Figs. 4(b) and 4(c).

Another interesting result is that the pixels detected with high certainty motion are mostly located at the outer cell layers as seen in Figs. 4 and 7. These cell layers are particularly important for root growth [33]. From the



(a)



(b)

Figure 12: Results from the MRF model with various inference schemes: (a) and (b) show the horizontal and vertical components using the expansion graph cuts algorithm, respectively; (c) and (d) show the horizontal and vertical components using the swap graph cuts algorithm, respectively. These algorithms proved sensitive to the choice of energy functions and parameters and often, as in this case, provided inconsistent results. Note that the sharp changes are due to the discrete encoding of velocities and are not an inherent problem with the method.



(c)



(d)

Fig. 12 (continued)




Figure 13: Results from the proposed method. The left column shows results using two frames and the right column shows results using three frames. (a) and (b) show estimated horizontal and vertical motion, respectively. Mid-grey values correspond to zero motion, white to right/down and black to left/up with a range of 8 in all directions. Points where velocity is not estimated are shown in grey. (c) Uncertainty (with black denoting certainty).



(c)

Fig. 13 (continued)

image shown, one may argue that the inner root tissues are not as clear as those of the outer layer cells, which might cause the algorithm to fail to report motion estimation. The possible role of outer vs. inner cell layers in restraining plant root growth is still a subject of much debate and is well worth further investigation [2, 4].

In comparing motion estimation methods, it is useful to contrast them in terms of how much spatial and temporal information is used. As a spatially and temporally localized method, KLT was capable of tracking accurately around resolvable structure. However, such a method is susceptible to gross errors due to local ambiguity and large motions. Using a method other than eigenvalue inspection, that is more application specific, could reduce such errors.

Black and Anandan’s robust smoothing method does not provide estimates of spatial uncertainty. The global smoothing reduces accuracy in some regions due to localised over-smoothing. This can be clearly seen by the incorrect underestimated motion at the root tip. Notice that there is no visual evidence in the motion estimate for a discontinuity at the root agar boundary. The method, without further prior knowledge, could not be expected to identify this motion boundary.

One disadvantage of the MAP methods used in MRFs is that, although they use probabilistic models, they do not produce spatial estimates of uncertainty (although marginals could be estimated using min-sum BP). It is important when discussing MRFs to mention the problem of automatic parameter learning, for example ML estimation of θ in $p(x; \theta) = Z^{-1}(\theta)e^{-E(x; \theta)}$ given a collection of *i.i.d.* datasets. This is complicated due to the global nature of the partition function, Z , and the large amount of data. Some recent progress, e.g. Hinton’s contrastive divergence method [7], may make this more feasible, but sensitivity to the smoothing parameters and poor or inconsistent inference sug-

gest this will be difficult on these data. There have been criticisms elsewhere of simple MRF models, e.g. those involving pairwise smoothing functions, for other applications, leading to proposals such as the product-of-experts methods for modelling, and learning, higher order clique potentials, e.g. [20].

For the proposed method, a significant degree of local ambiguity was removed by employing temporal constraints. Notice the errors due to aperture problems on cell membranes. The certainty is correctly estimated to be low where such errors occur. The large motions, lack of evidence in many regions, secondary motions and significant non-Gaussian noise justify the need for more sophisticated spatial constraints in order to produce more accurate motion estimates.

5 Conclusions and Future Work

A robust method was developed for estimating motion in CLSM images of living plant roots. Experimental results demonstrated that by employing robust likelihoods, temporal constraints, and estimates of uncertainty, the algorithm performed well for motion estimation in this domain. Comparative results showed that a spatially and temporally localized method was not able to track the smaller, repetitive cell structure in the root tip and was somewhat prone to identifying spurious features. Generic low-level smoothing functions do not perform well in this application, due to large regions with little stable structure. In contrast to recent positive investigations, GC and loopy BP for MAP estimation of the MRF labels provided inconsistent results on these data. For the large motions, lack of evidence in many regions, secondary motions and significant non-Gaussian noise, a key conclusion that can be drawn is the need for more sophisticated spatial constraints in order to produce more accurate motion estimates.

The image velocity was assumed to be approximately constant within a local neighbourhood. The size of the neighbourhood was chosen manually and depended on scale. Future work could explore setting this automatically, possibly differently across the image. The likelihood function also embodied the assumption that the appearance of a local neighbourhood is constant over a short period of time (given the motion). The error in the assumption of appearance constancy was modelled as statistically independent pixel variation. Various forms of pixel-wise likelihood should be further investigated, e.g. mixtures. Careful examination of Fig. 2(a) suggests that much of the variation is highly localised and that the independence assumption is a reasonable modelling strategy. However, the pixel variations given the motion are to some extent dependent. Changing the model of local appearance to account for these dependencies between pixel variations in a local neighbourhood would be more computationally expensive and non-trivial. It would be interesting to investigate whether such models would yield improved motion estimation given the available data from which to learn such dependencies. However, such an investigation lies beyond the scope of this paper.

Since the prior is very uninformative, all the pixels were used to test for multiple strong likelihood peaks. For faster motion estimation, one possible approach is to use thresholding to identify cells firstly, but this should be employed with care. Due to the unusual properties of plant root images acquired by CLSM, it is inevitable that some thresholded pixels that ideally should form parts of cells will be omitted or vice versa.

It is likely that we will be placing more emphasis on developing task specific, higher level, spatial constraints, such as those discussed in [27], that are able to incorporate stronger prior knowledge on cell structure and kinematics. One novel challenge of building and inferring such high level models is varying topology, both between datasets and over time.

In addition, further investigation should be conducted to take into account the effects of cell division. Recently, Chavarria-Krauser and Schurr [8] attempted to describe the process of root elongation using a model of cell division and expansion. Cell division was modeled at a constant rate in the **division zone**, and during this time the cell extension rate is increasing linearly with cell length.

Here we only considered recovering locally rigid translations. A more realistic model would also account for deformation, and would be important for longer-term tracking. Since cell movements are essentially in 3D, we plan to extend the current 2D motion estimation to 3D.

There are many possible directions for future work to improve the performance of MRF models for plant cell tracking. One direction is to develop an MRF model for joint segmentation and motion estimation, allowing for the incorporation of non-isotropic and more accurate smoothing and observation functions. Another direction would be to learn higher order MRF potentials. However, the issue of reliable inference and high resource requirements would be compounded. Our plan when investigating MRFs was to develop the global smoothing approach by doing joint motion estimation and segmentation, an approach that has had some success in other applications, e.g. [40].

Acknowledgements: We thank the Biotechnology and Biological Sciences Research Council (BBSRC Grant Ref. BBSB16720) and the Scottish Government Rural and Environmental Research and Analysis Directorate (RERAD) for funding this work. We are also grateful to Michael Black, Stan Birchfield and those who contributed to the Middlebury MRF suite, for their source code. Thanks to Dr Jim Haseloff for providing LTI-eGFP lines and Fraser Bransby for helpful discussions.

References

- [1] S. Baker and I. Matthews. Lucas-Kanade 20 years on: A unifying framework. *International Journal of Computer Vision*, 56(3):221–255, 2004.
- [2] P.W. Barlow. Anatomical controls of root growth. *Aspects of Applied Biology*, 22:57–66, 1989.

- [3] J.L. Barron and L. Liptay. Measuring 3D Plant Growth Using Optical Flow. *BioImaging*, 5:82–86, 1997.
- [4] 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*, 57(2):437–447, 2006.
- [5] 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, 1996.
- [6] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. In *IEEE International Conference on Computer Vision*, pages 377–384, 1999.
- [7] M.A. Carreira-Perpiñán and G.E. Hinton. On contrastive divergence learning. In *Workshop on Artificial Intelligence and Statistics*, pages 33–40, 2005.
- [8] A. Chavarria-Krauser and U. Schurr. A cellular growth model for root tips. *Journal of Theoretical Biology*, 230:21–32, 2004.
- [9] P.F. Felzenszwalb and D.P. Huttenlocher. Efficient belief propagation for early vision. *International Journal of Computer Vision*, 70(1):261–268, 2006.
- [10] A. Genovesio, T. Liedl, V. Emiliani, W.J. Parak, M. Coppey-Moisan, and J.C. Olivo-Marin. Multiple particle tracking in 3-d+t microscopy: method and application to the tracking of endocytosed quantum dots. *IEEE Transactions on Image Processing*, 15(5):1062–1070, 2006.
- [11] S. Gilroy. Fluorescence microscopy of living plant cells. *Annual Review of Plant Physiology and Plant Molecular Biology*, 48:165–190, 1997.
- [12] M. R. Hanson and R. H. Köhler. GFP imaging: methodology and application to investigate cellular compartmentation in plants. *Journal of Experimental Botany*, 52(356):529–539, 2001.
- [13] G. Harauz and F. P. Ottensmeyer. Interpolation in computing forward projections in direct three-dimensional reconstruction. *Physics in Medicine and Biology*, 28(12):1419–1427, 1983.
- [14] B.K.P. Horn and B.G. Schunk. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
- [15] Improvisation. Online. Available: <http://www.improvisation.com/>.
- [16] P. Kohli and P.H.S. Torr. Efficiently solving dynamic Markov random fields using graph cuts. In *IEEE International Conference on Computer Vision*, volume II, pages 922–929, 2005.

- [17] V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? In *European Conference on Computer Vision*, pages 65–81, 2002.
- [18] S. Kurup, J. Runions, U. Köhler, L. Laplaze, S. Hodge, and J. Haseloff. Marking cell lineages in living tissues. *Plant*, 42(3):444–453, 2005.
- [19] C.-Y. Lam. *Applied Numerical Methods for Partial Differential Equations*. Prentice Hall, Inc., New Jersey, 1994.
- [20] X. Lan, S. Roth, D. Huttenlocher, and M.J. Black. Efficient belief propagation with learned higher-order markov random fields. In *European Conference on Computer Vision*, 2006.
- [21] S. Li. *Markov Random Field Modeling in Computer Vision*. Springer Verlag, 1995.
- [22] B. Likar, J. B. Maintz, M. A. Viergever, and F. Pernus. Retrospective shading correction based on entropy minimization. *Journal of Microscopy*, 197(3):285–295, 2000.
- [23] B.D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *International Joint Conference on Artificial Intelligence*, pages 674–679, 1981.
- [24] D.J.C. Mackay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, The Edinburgh Building, Cambridge, CB2 2RU, UK, 2002.
- [25] S. de F. Maraschin, M. Vennik, G. E. M. Lamers, H. P. Spaink, and M. Wang. Time-lapse tracking of barley androgenesis reveals position-determined cell death within pro-embryos. *Planta*, 220:531–540, 2005.
- [26] M.A. Model and J.K. Burkhardt. A standard for calibration and shading correction of a fluorescence microscope. *Cytometry*, 44(4):309–316, 2001.
- [27] T. J. Roberts, S.J. McKenna, J. Hans, T. A. Valentine, and A. G. Bengough. Part-based multi-frame registration for estimation of the growth of cellular networks in plant roots. In *International Conference on Pattern Recognition*, August 2006.
- [28] B. Scheres, P. Benfey, and L. Dolan. Root development. In C.R. Somerville and E.M. Meyerowitz, editors, *The Arabidopsis Book*. American Society of Plant Biologists, Rockville, MD, doi:10.1199/tab.0101, 2002.
- [29] R.E. Sharp, W.K. Silk, and T.C. Hsiao. Growth of the primary maize root at low water potentials. i. spatial distribution of expansive growth. *Plant Physiology*, 87:50–57, 1988.

- [30] M. Shimizu and M. Okutomi. Precise sub-pixel estimation on area-based matching. In *IEEE International Conference on Computer Vision*, pages 90–97, 2001.
- [31] A. Singh. *Optic Flow Computation: A Unified Perspective*. IEEE Computer Society Press, California, 1991.
- [32] M.N. Somleva, E.D.L. Schmidt, and S.C. de Vries. Embryogenic cells in *Dactylis glomerata* L. (Poaceae) explants identified by cell tracking and by serk expression. *Plant Cell Report*, 19:718–726, 2000.
- [33] R. Swarup, E. Kramer, P. Perry, K. Knox, H. M. Ottoline Leyser, J. Haseloff, G. T. S. Beemster, R. Bhalerao, and M. J. Bennett. Root gravitropism requires lateral root cap and epidermal cells for transport and response to a mobile auxin signal. *Nature Cell Biology*, 7:1057–1065, 2005.
- [34] R. Szeliski, R. Zabih, D. Scharstein, O. Veksler, V. Kolmogorov, A. Agarwala, M. Tappen, and C. Rother. A comparative study of energy minimization methods for markov random fields. In *European Conference on Computer Vision*, volume 2, pages 19–26, Graz, Austria, May 2006.
- [35] M. Tappen and W. Freeman. Comparison of graph cuts with belief propagation for stereo using identical MRF parameters. In *IEEE International Conference on Computer Vision*, volume 2, pages 900–908, October 2003.
- [36] M.A.J. Toonen, T. Hendriks, E.D.L. Schmidt, H.A. Verhoeven, A. van Kammen, and S.C. de Vries. Description of somatic-embryoforming single cells in carrot suspension cultures employing video cell tracking. *Planta*, 194:565–572, 1994.
- [37] MetaMorph Imaging System. Universal Imaging Corporation. Online. Available: <http://www.image1.com/>.
- [38] C.M. van der Weele, H.S. Jiang, K.K. Palaniappan, V.B Ivanov, K. Palaniappan, and T.I Baskin. A new algorithm for computational image analysis of deformable motion at high spatial and temporal resolution applied to root growth. *Plant Physiol*, 132(3):1138–1148, 2003.
- [39] A. Walter, H. Spies, S. Terjung, R. Küsters, N. Kirchgessner, and U. Schurr. Spatio-temporal dynamics of expansion growth in roots: automatic quantification of diurnal course and temperature response by digital image sequence processing. *Journal of Experimental Botany*, 53(369):689–698, 2002.
- [40] P.P. Wyatt and J.A. Noble. MAP MRF joint segmentation and registration of medical images. *Medical Image Analysis*, 7(4):539–552, 2003.
- [41] G. Zeng, S. Birchfield, and C. E. Wells. Detecting and measuring fine roots in minirhizotron images using matched filtering and local entropy thresholding. *Machine Vision & Applications*, 17(4):265–278, 2006.