

Automated Semantic Analysis of Changes in Image Sequences of Neurons in Culture

Omar Al-Kofahi^{†§}, Richard J. Radke[†], Badrinath Roysam^{*†}, and Gary Banker[‡]

[†] Rensselaer Polytechnic Institute

110 8th Street, Troy, New York 12180, USA

[‡] Oregon Health and Science University

3181 SW Sam Jackson Park Road, L606, Portland, Oregon 97239, USA

OAl-Kofahi@as-e.com, {rjradke, roysam}@ecse.rpi.edu, banker@ohsu.edu

Abstract—Quantitative studies of dynamic behaviors of live neurons are currently limited by the slowness, subjectivity, and tedium of manual analysis of changes in time-lapse image sequences. Challenges to automation include the complexity of the changes of interest, the presence of obfuscating and uninteresting changes due to illumination variations and other imaging artifacts, and the sheer volume of recorded data.

This paper describes a highly automated approach that not only detects the interesting changes selectively, but also generates quantitative analyses at multiple levels of detail. Detailed quantitative neuronal morphometry is generated for each frame. Frame-to-frame neuronal changes are measured and labeled as growth, shrinkage, merging or splitting, as would be done by a human expert. Finally, events unfolding over longer durations, such as apoptosis and axonal specification, are automatically inferred from the short-term changes.

The proposed method is based on a Bayesian model selection criterion that leverages a set of short-term neurite change models and takes into account additional evidence provided by an illumination-insensitive change mask. An automated neuron tracing algorithm is used to identify the objects of interest in each frame. A novel curve distance measure and weighted bipartite graph matching are used to compare and associate neurites in successive frames. A separate set of multi-image change models drives the identification of longer-term events.

The method achieved frame-to-frame change labeling accuracies ranging from 85–100% when tested on 8 representative recordings performed under varied imaging and culturing conditions, and successfully detected all higher-order events of interest. Two sequences were used for training the models and tuning their parameters; the learned parameter settings can be applied to hundreds of similar image sequences, provided imaging and culturing conditions are similar to the training set. The proposed approach is a substantial innovation over manual annotation and change analysis, accomplishing in minutes what it would take an expert hours to complete.

Index Terms—Change detection, change understanding, morphological dynamics, assay automation, statistical model selection, curve similarity, event analysis.

I. INTRODUCTION

Technologies for time-lapse imaging of live cells *in vitro* enable the observation of complex morphological and functional changes as a function of time. Of particular interest here is phase-contrast time-lapse imaging of live neurons grown in culture [1], [2]. The recorded image sequences are challenging to analyze due to the complexity of the neuronal structures, the presence of a mixture of interesting and nuisance changes, and their sheer size. Figure 1 illustrates several events observed in neuronal cultures (*in vitro*). The boxes in this figure highlight the merging of a pair of neurites, while the arrows

indicate neuronal growth and shrinkage. The pair of images in the top row (Sequence A) was drawn from a longer sequence composed of 147 frames, taken five minutes apart, where each frame is 1470×1180 pixels in size, and each pixel has a depth of 12 bits. The pair of images in the bottom row (Sequence B) was drawn from a similar sequence composed of 200 frames, implying a total data size close to 200 MB for each sequence.

Numerous such image sequences are currently captured in the neurobiology research community in the context of diverse studies (e.g. [3], [4]). There is a compelling need for automated methods to analyze these sequences in support of conventional hypothesis-driven studies, as well as emerging studies in cell and tissue engineering, automated assays, and high-throughput discovery research. Many studies require careful, detailed, and intelligent markup of the images, including tracing of neurites, quantitative morphometry, and analysis of complex spatio-temporal changes, to characterize neuronal structures and dynamic events of interest (e.g. [5]). While some of these changes occur from one image frame to the next, others only become apparent over much longer durations. Previous research on automating the analysis of these image sequences has mostly relied upon tedious interactions with computer-assisted markup tools to extract objects of biological interest, without any automatic higher-level reasoning about the changes they undergo.

Achieving highly automated, human-like analysis of neuronal imagery on a large scale requires multiple advances reported here. First, we adapt a fast, accurate, and fully automated algorithm for tracing neurites originally developed for fluorescence imagery [6] to phase-contrast microscopy. Second, we adopt robust registration, background modeling, and pixel-level change detection algorithms to generate a binary change mask for each image pair. Third, we design and validate mathematical models describing the morphological dynamics of cultured neuronal networks. To successfully fit these models to neurite traces, we propose a novel curve distance measure. Finally, we integrate the above building blocks into a principled Bayesian model selection framework that is robust to imaging artifacts and image analysis errors. The model selection uses a bipartite-graph-based algorithm to associate neurites between each image pair, even when such associations are not one-to-one. Finally, events unfolding over longer durations, such as apoptosis and axonal specification, are automatically inferred from the short-term changes.

The paper is organized as follows. The next section describes related work from the neurobiology and computer vision literature. Section III goes into detail on each of the building blocks in the change understanding framework. Section IV gives a performance analysis of the algorithm and presents automatically generated tables and graphical output resulting from applying the framework to several image sequences of cultured neurons *in vitro*. Finally, Section V concludes and discusses ideas for future work.

[§] Omar Al-Kofahi is now with American Science and Engineering, Inc., Billington, MA, USA.

* Please address correspondence to Badrinath Roysam. This research was supported in part by CenSSIS, the NSF Center for Subsurface Sensing and Imaging Systems under the Engineering Research Centers program of the National Science Foundation (Award Number EEC-9986821) and by Rensselaer Polytechnic Institute.

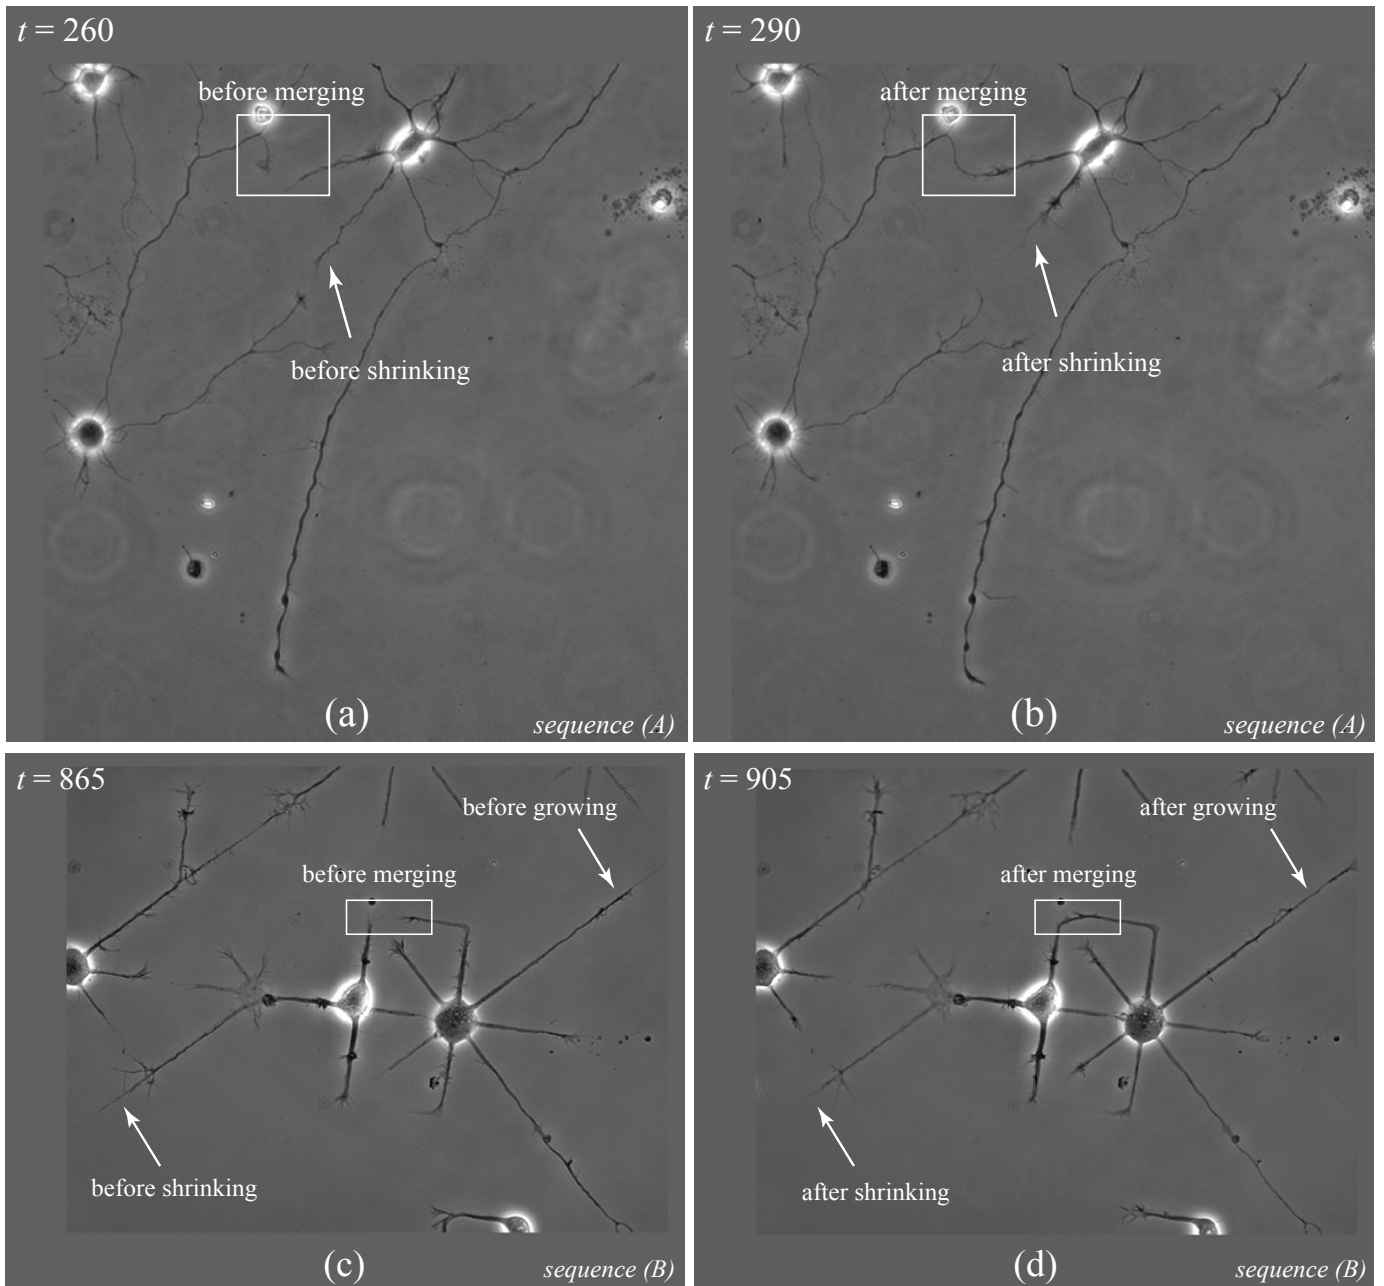


Fig. 1. Illustrating events in cultured neuron image sequences. Images were collected using a phase-contrast microscope. (a) and (b), sample frames from an image sequence of neurons cultured on a smooth surface. The sequence contains 147 frames taken five minutes apart, 1470×1180 pixels each, with 12 bits depth. (c) and (d), sample frames from an image sequence of neurons cultured on a patterned surface. The sequence contains 200 frames taken five minutes apart, 1300×1080 pixels each, with 12 bits depth. The boxes highlight the merging of two neurites, while the arrows indicate neurite growth/shrinkage. An automated analysis system should detect such phenomena automatically and describe them intuitively.

II. RELATED WORK

Analyzing changes in time-lapse images of live neurons in culture is of widespread interest [1], [2]. Neurobiologists are often interested in studying the effects of different regulators on neurite outgrowth and orientation. Understanding neurite outgrowth is important to brain and spinal cord research, because promoting neurite outgrowth potentially has curative effects in nerve-injury-related diseases such as Alzheimer's, Parkinson's disease, and spinal cord injury [7]. For the most part, time-lapse image sequences are currently analyzed manually using pointing tools (e.g. mice and tablets) and annotation software. Judgements about important changes in the image sequences are rarely made automatically. For example, Zhang et al. used

substrates with smooth and patterned surfaces to control neurite outgrowth and orientation [8]. Neurite lengths and orientations were measured using a computer-aided tool [9], and analysis of changes was performed manually. Rozovsky et al. studied the influence of aging astrocytes on neurons by associating aging with neurite atrophy [10]. After manually identifying image regions where neurites are present, neurite centerlines were extracted by thresholding followed by skeletonization. These centerlines were manually analyzed to determine atrophy. Lin et al. investigated the role of netrin-1 and slit-2 proteins in regulating and orienting axonal outgrowth [11]. While a computer assisted tool [12] was used to measure neurite lengths, axonal outgrowth was analyzed manually. Esch et al. investigated

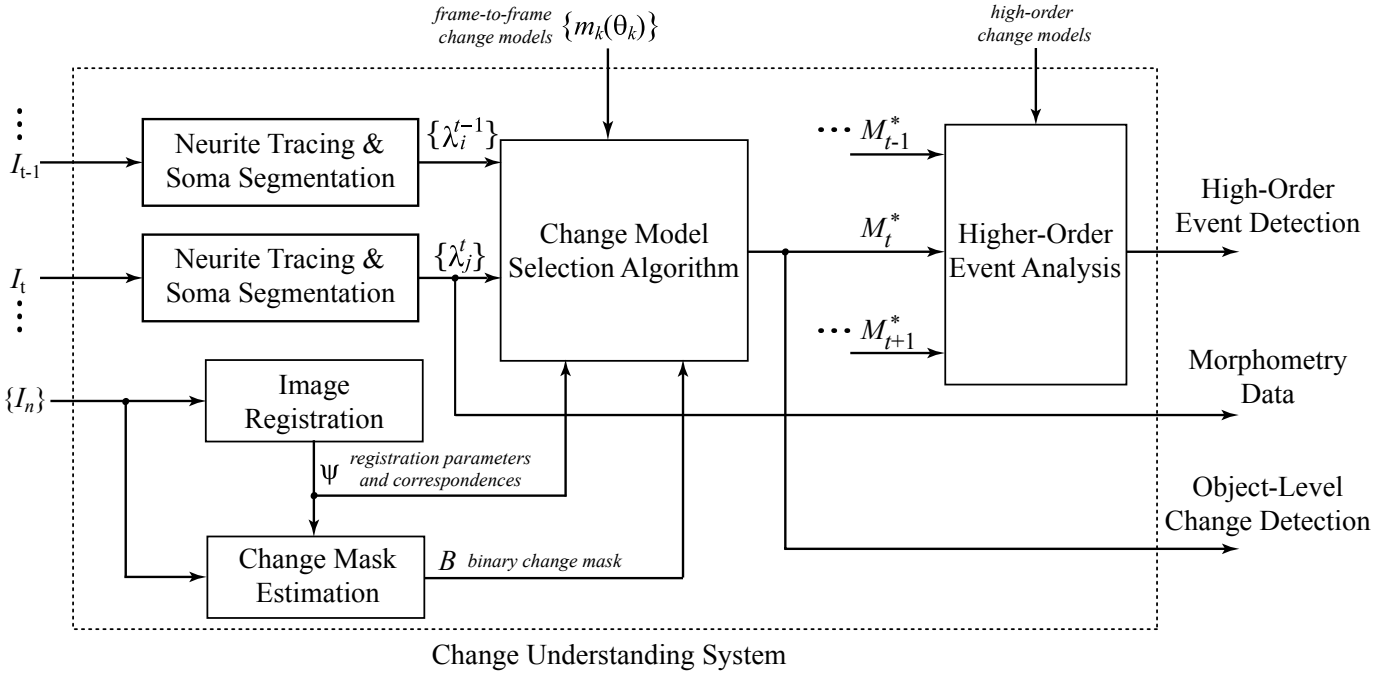


Fig. 2. Building blocks of the proposed change understanding system. A neurite tracing and soma segmentation algorithm operates on each of the sequence images, extracting the interesting objects in it along with their attributes. Images are registered, and a binary change mask is estimated for each pair of images. The neurite tracing output for a successive pair of images ($\{\lambda_i^{t-1}\}$ and $\{\lambda_j^t\}$), the registration parameters Ψ , and the change mask B are fed into a model selection algorithm, provided with a set of neurite-level change models $\{m_k(\theta_k)\}$. The model selection algorithm generates an optimal matching M^* between elements of $\{\lambda_i^{t-1}\}$ and $\{\lambda_j^t\}$, along with the corresponding change models and their parameters. The change results from the frame pairs can then be fed into a multi-image event analysis algorithm that detects and describes important sequence-level semantic changes.

the effects of cell adhesion molecules NgCAM and N-cadherin on neuronal outgrowth [5], [13]. While the neurite lengths were measured using a computer assisted tool, analysis of changes was performed manually. Nathan et al. studied the effect of apolipoproteins E3 and E4 on neurite outgrowth from cultured adult mouse cortical neurons [3]. Neurite length measurement as well as outgrowth analysis was performed manually. Withers et al. investigated the effect of bone morphogenetic protein-7 on dendritic development of cultured hippocampal neurons [4]. Neurite lengths were measured using a computer assisted tool [12], while outgrowth analysis was performed manually.

Several methods for detecting and describing changes in time-lapse image sequences have also been described in the broader biomedical image analysis literature. For example, Kauffmann et al. segmented cartilage volumes from three dimensional magnetic resonance imaging (3D MRI) images with an active contour model [14]. Cartilage volumes for the same patient at different sessions were then registered, and local volume changes were measured and reported. Meyer et al. segmented liver volumes from CT images using a liver atlas [15]. Segmented liver volumes for the same patient at different sessions were then registered, and local volume changes greater than a threshold were labeled as abnormal. Sbeh et al. segmented regions of retinal images corresponding to drusen (a retinal disease appearing as yellowish deposits) using active contours [16]. Retinal images for a patient at different sessions were registered, and drusen boundary changes were measured, which were manually used to describe the progression of the disease.

Our work is also inspired by advances in non-biomedical areas, especially surveillance. For example, Wren et al. used a human model for tracking people and interpreting their behaviors [17]. The model contained “blobs” for a person’s head, hands, torso, etc. The blob segmentation for the hands was also used with a hidden Markov

model to interpret the American Sign Language. Haritaoglu et al. [18] and Oliver et al. [19] used similar human models to track people and detect certain behaviors such as people carrying, exchanging and depositing objects. In another video surveillance application, Collins et al. segmented and tracked moving objects, and classified them into known types using neural networks [20]. Humans were additionally classified as either walking or running based on the changes in gait angle.

III. ALGORITHMS AND METHODS

The proposed method processes a sequence of images to generate a semantic, quantitative description of the changes, using the framework illustrated in Figure 2. We briefly summarize the algorithm here and provide the details in the following subsections. First, every adjacent pair of sequence images is registered using normalized cross correlation, to avoid introducing false changes due to misalignment (Section III-A). The registered image pair is fed to a pixel-level change detection algorithm, which produces a binary change mask B indicating the estimated presence or absence of change at each pixel in the second image (Section III-B). In parallel, a neurite tracing algorithm operates on each registered image, extracting neurite centerlines denoted $\{\lambda_i^t\}$ for the image at time t (Section III-C). Mathematically, these are described as curve segments, with attributes such as length, width, and branch points. While we concentrate here on changes that can be expressed in terms of the neurites, the somas are also segmented in each image.

The traced neurites and segmented somas, the registration parameters, and the binary change mask are processed by an integrative block termed the “Change Model Selection Algorithm” in Figure 2. This block is provided with a set of neurite change models $\{m_k\}$, each of which is associated with a semantic description of the change it represents (Section III-D). For each curve segment λ_i^{t-1}

corresponding to a neurite centerline trace in image I_{t-1} , we estimate the best change model to map it into another curve segment λ_j^t in I_t . Model parameters θ_k^{ij} for model m_k acting on λ_i^{t-1} and λ_j^t are estimated as:

$$\hat{\theta}_k^{ij} = \arg \min_{\theta_k} d(\lambda_j^t, m_k(\lambda_i^{t-1}, \theta_k)), \quad (1)$$

where $d(\cdot, \cdot)$ is a novel distance measure expressing the similarity between two curve segments (Section III-E). The model that best describes the change between λ_i^{t-1} and λ_j^t is selected as:

$$m_k^{ij} = \arg \max_k p(\lambda_i^{t-1}, \lambda_j^t | m_k(\hat{\theta}_k^{ij}), B), \quad (2)$$

where $p(\cdot)$ measures the *a posteriori* probability of observing the neurite pair given model m_k and the change mask B (Section III-F).

After identifying the model that best describes the change for every feasible match $(\lambda_i^{t-1}, \lambda_j^t)$, the model selection algorithm operates on $\{\lambda_i^{t-1}\}$ and $\{\lambda_j^t\}$ to determine the best overall matching between the two sets, using a bipartite-graph-based algorithm (Section III-G). Once the frame-to-frame changes have been estimated, they are further processed by a longer-term event analysis block that detects semantic events at the sequence level (Section III-I). In the following, each of the above building blocks is described in greater detail.

A. Registration of Successive Image Frames

Each neuron culture was mounted on a microscope x - y stage, and images were acquired every five minutes. The cultures were not removed or remounted throughout each experiment. However, sequence images were found to be misaligned by up to 10 pixels between consecutive images due to x - y stage vibrations. Hence, each image was registered with the previous one, using a translational model estimated with normalized cross correlation, i.e.

$$(\hat{u}, \hat{v}) = \arg \max_{(u,v)} \frac{\sum_{x,y} [(I_{t-1}(x,y) - \bar{I}_{t-1}) [I_t(x-u, y-v) - \bar{I}_t]]}{\left[\sum_{x,y} (I_{t-1}(x,y) - \bar{I}_{t-1})^2 \sum_{x,y} (I_t(x-u, y-v) - \bar{I}_t)^2 \right]^{1/2}}$$

where \bar{I}_{t-1}, \bar{I}_t are the average intensities of I_{t-1}, I_t respectively [21]. When tested over a sequence of 100 images, the average registration error was less than one pixel compared to manual registration ground truth. For different applications, one might need to consider more sophisticated transformations and estimation methods; see Zitova et al. [22] for a recent survey.

B. Illumination-Insensitive Change Mask Estimation

A simplified version of Toyama et al.'s Wallflower algorithm [23] is used to generate a binary change mask B_t for each pair of adjacent images (I_{t-1}, I_t) , where a pixel value of 1 indicates change and 0 indicates otherwise. The change mask is used to add robustness to errors in the neurite tracing and soma segmentation output by providing independent evidence as to the presence or absence of change.

First, the pixel intensity values in every image were normalized to have the same mean and variance as in the first image. Next, a Wiener filter [24] is used to make a linear prediction of the pixel values in I_t , based on the ten previous images. Any pixel whose intensity deviates significantly from the predicted value is declared as a change. The change mask was post-processed using a median filter to remove salt-and-pepper noise. As above, different applications may require more

sophisticated change detection algorithms; see Radke et al. [25] for a recent survey.

C. Neurite Tracing and Soma Segmentation

For cultured neuron sequences, the main objects of interest are the neurites, which can be described mathematically by curve segments. To extract and describe these curves, we use a tracing algorithm that automatically locates neurite centerlines [6], [26]. Based on the assumption that the neurites are bounded by nearly parallel edges, the algorithm starts by automatically finding seed points along neurites, then traces them in an exploratory manner by following the neuronal topology, guided by responses of multiple directional correlation kernels. We note that this algorithm is only one of many that have been proposed for automated neurite tracing (see, e.g. [27], [28], [29], [30], [31], [32]).

Somas were automatically segmented in the first frame of each sequence using the following simple algorithm. First, bright halos surrounding somas (due to the halo artifact in phase contrast microscopy [2]) were detected and segmented by adaptive thresholding [33]. Next, morphological closing was applied to the segmentation, followed by connected component analysis. Connected components smaller than a threshold (30 pixels) were rejected, and ellipses were fit to the remaining objects as illustrated in Figure 3. Somas segmented in the first frame of the sequence were then automatically tracked in subsequent frames using normalized cross correlation. The output of the neurite tracing and soma segmentation algorithm is illustrated in Figure 8-a, where the curve segments ($\{\lambda_i^{t-1}\}$) and the soma ellipses are overlaid on top of a gray scale image I_{t-1} .




Fig. 3. Illustrating automatic soma segmentation. (a) First frame of sequence C, (b) automatic thresholding followed by morphological closing, and (c) ellipses fit to the valid segmented objects as described in text.

(3) While in this paper, we only consider change models that can be described in terms of the neurites, one can imagine more sophisticated studies involving both neurons and somas, which can easily be accommodated into our framework.

D. Mathematical Models for Short-Term Changes in Neurons

Mathematical modeling of short-term (frame-to-frame) changes in neuronal structures enables us to formulate the problem of describing these changes in a statistical model selection framework. We constructed three basic models for neuronal changes: no change, growth/shrinkage, and merging/splitting. If the study warrants, this “library” of models can be extended to describe a richer set of behaviors using the principles described below.

In the following description, we fix a time t , and denote the set of traced neurites at times $t-1$ and t as $\{\lambda_i^{t-1}, i = 1, \dots, N_{t-1}\}$ and $\{\lambda_j^t, j = 1, \dots, N_t\}$ respectively. The numbers of elements in these sets are not assumed to be equal. Given a specific pair of neurites λ_i^{t-1} and λ_j^t , and a set of change models $\{m_k\}$, our first goal is to find the model that best describes the change between these two neurites, as well as optimal estimates of any necessary model parameters, denoted θ_k . A change model m_k acts on the curve segment λ_i^{t-1} and deforms it in a way that represents a valid neurite behavior, such that the result is as similar as possible to λ_j^t .

The similarity between λ_i^{t-1} after deformation and λ_j^t (described in the next subsection) reflects how well m_k describes the change. Realistically, we need not consider every possible match between curve segment pairs in $\{\lambda_i^{t-1}\}$ and $\{\lambda_j^t\}$, since many matchings are infeasible for a given scenario. For example, the neurites in Figure 1 do not move by a large amount between acquired images, but rather locally deform. Hence, when searching for matches, we need only consider candidates in a local spatial neighborhood. We note that this largely mitigates any problems with computational complexity as the number of neurites grows.

We model the likelihood of every change model as the weighted sum of two probability density functions: one that depends on the distance between curve segments, and one that depends on the binary change mask estimated previously. Namely,

$$p(\lambda_i^{t-1}, \lambda_j^t | m_k(\hat{\theta}_k^{ij}), B_t) = \delta \cdot p_{dist}(\lambda_i^{t-1}, \lambda_j^t | m_k(\hat{\theta}_k^{ij})) + (1 - \delta) \cdot p_{mask}(\lambda_i^{t-1}, \lambda_j^t | m_k(\hat{\theta}_k^{ij}), B_t), \quad (4)$$

where $\hat{\theta}_k^{ij}$ are the parameters of change model m_k estimated for curve segments λ_i^{t-1} and λ_j^t . The value of δ reflects our confidence in the output of the neurite tracing algorithm. In practice, the neurite tracing algorithm generates a few minor errors, so we use a value of δ slightly less than 1, to take into account the independent evidence provided by the change mask (see below). In the following, we describe the two component densities p_{dist} and p_{mask} for each of our basic change models. We note that at this stage, the various models include several tuning parameters that influence the overall performance of the algorithm. We discuss how each of these tuning parameters is estimated from processing a small set of manually labeled training data in Section IV-A.

1) *No Change*: The no-change model corresponds to the identity transformation— that is, we assume that any apparent position difference between the neurites λ_i^{t-1} and λ_j^t is due to sensor noise, tracing errors, or minor neurite deformation. The two components of (4) for the likelihood of the no-change model are:

$$p_{dist}(\lambda_i^{t-1}, \lambda_j^t | m_{no\ change}) = \tau e^{-\tau d(\lambda_i^{t-1}, \lambda_j^t)} \quad (5)$$

$$p_{mask}(\lambda_i^{t-1}, \lambda_j^t | m_{no\ change}, B_t) = \varphi e^{-\varphi r}. \quad (6)$$

We show in Section IV-A that exponential distributions are good approximations for these densities, and describe how the parameters τ, φ are estimated. Here, the function d is the integral area distance for curve segments, which we define in the next subsection. The variable r in (6) is the fraction of ones in the change mask B_t , evaluated over the union of the regions surrounding λ_i^{t-1} and λ_j^t . Under the no change model, we expect r to be small. Figure 4 illustrates the intuition behind this rule for an example containing a neuron tracing error. The neurite in Figure 4b is not fully traced. With no information other than the traces, we would mistakenly infer that the neurite shrank between the image frames. However, the binary change mask in the potential shrinkage region is almost entirely 0, indicating low evidence for shrinkage and biasing the model selection towards no-change.

2) *Growth/Shrinkage*: This model describes the growth or shrinkage of curve segment λ_i^{t-1} into the segment λ_j^t . To fit a growth model, we extrapolate λ_i^{t-1} beyond its growing end (the one not attached to the soma or another neurite) in the direction of the tangent at that end. This change model has a single parameter s , which indicates the positive length (in pixels) by which λ_i^{t-1} has grown to become λ_j^t . s is estimated as in (1) such that the distance between $m_{grow}(\lambda_i^{t-1}, \hat{s})$ and λ_j^t is minimized. This process is illustrated in Figure 5: curve segment λ_i^{t-1} is extended along the broken line (Panel




Fig. 4. Illustrating change mask contribution to the likelihood of different change models. (a,b) Corresponding regions and traces cropped from a pair of consecutive images from Sequence B. This situation should be labeled “no-change” but is in danger of being labeled “shrinkage” due to a tracing error. (c) The corresponding binary change mask region. (d,e) Regions surrounding λ_i^{t-1} and λ_j^t . (f) $d \oplus e$. (g) $f \cap c$. Since the change mask is almost entirely 0 in the region of interest, the likelihood p_{mask} is very low, which is enough to tip the model selection decision from shrinkage to no-change.

a), and s is estimated such that the integral area distance is minimized (Panel b). Hence,

$$p_{dist}(\lambda_i^{t-1}, \lambda_j^t | m_{grow}) = \tau e^{-\tau d(m_{grow}(\lambda_i^{t-1}, \hat{s}), \lambda_j^t)} \quad (7)$$

$$p_{mask}(\lambda_i^{t-1}, \lambda_j^t | m_{grow}, B_t) = \varphi e^{-\varphi(1-r)}, \quad (8)$$

where \hat{s} is the estimate for s as in (1). r in (8) is defined in the same way as the no-change case. However, p_{mask} has $1 - r$ in the exponent as opposed to r in the no-change model, since in this case we expect to see more change mask ones than zeros in the region of interest. The shrinkage model is the same as the growth model, with the roles of λ_i^{t-1} and λ_j^t reversed (in this case, the parameter s is negative).




Fig. 5. Illustrating the integral area distance measure between two curve segments. (a) We want to estimate the number of pixels to extend curve segment λ_i^{t-1} along the tangent at its end (broken line) to best match λ_j^t . (b) The integral area distance between λ_i^{t-1} and the extension of λ_i^{t-1} as a function of s . The minimum distance was found when λ_i^{t-1} was extended by \hat{s} pixels.

3) *Merge/Split*: The merge model describes the merging of two curve segments λ_i^{t-1} and λ_j^{t-1} in image I_{t-1} into one segment λ_k^t in image I_t . To evaluate this model, we create a new “neurite” λ_{ij}^{t-1} in I_{t-1} by connecting λ_i^{t-1} and λ_j^{t-1} by their two nearest endpoints, and measure the distance of this curve to λ_k^t . Hence,

$$p_{dist}(\lambda_i^{t-1}, \lambda_j^{t-1}, \lambda_k^t | m_{merge}) = \tau e^{-\tau d(\lambda_{ij}^{t-1}, \lambda_k^t)} \quad (9)$$

$$p_{mask}(\lambda_i^{t-1}, \lambda_j^{t-1}, \lambda_k^t | m_{merge}, B_t) = \varphi e^{-\varphi(1-r)}. \quad (10)$$

As in the growth model, $1 - r$ is used in the expression for p_{mask} . The split model is similar to the merge model, except the

new “neurite” is created from two neurites in the second image I_t and matched to a single neurite in I_{t-1} .

E. A Curve Distance Measure for Comparing Neurite Segments

As mentioned above, we require a measure for the distance between two curve segments that enables us to obtain reasonable estimates for how much one curve must be grown or shrunk to best match another. We propose a novel curve distance measure, the integral area distance, that is robust and enables us to obtain good estimates of the growth/shrink parameter. Let $\lambda_i : [0, 1] \rightarrow \mathbb{R}^2$, and $\lambda_j : [0, 1] \rightarrow \mathbb{R}^2$ be curve segments parameterized with constant speed. The integral area distance is defined as:

$$d(\lambda_i, \lambda_j) = \min_{\omega \in P} \int_0^1 \|\lambda_i(l) - \lambda_j(\omega(l))\| dl, \quad (11)$$

where P is the class of continuous, non-decreasing functions ω with $\omega(0) = 0$ and $\omega(1) = 1$. We use dynamic programming (i.e. Dijkstra’s algorithm [34]) to find the function ω that minimizes the integral in (11).

The integral area distance for a pair of curve segments as a function of the growth parameter is illustrated in Figure 5b. The integral area distance is minimized (suggesting the curves are most similar) when λ_i^{t-1} is extended by \hat{s} pixels, which reflects our notion of similarity between neurites. Recall that the images have already been registered to each other by a global transformation, so we assume any additional neurite deformations should be quantified by the distance measure. We considered several other curve distance measures, including the Fréchet distance [35], Hausdorff distance [36], and turning function distance [37]. However, we found these measures to be unsuitable for robust estimation of the growth parameter. For example, in Figure 5a, both the Fréchet and Hausdorff distances would suggest extending λ_i^{t-1} only to the point $s_{h,f}$, while the turning function distance would suggest no extension at all. As a secondary consideration, these distance measures are not robust, in that changing an arbitrarily small portion of one curve can produce large changes in the distance measure. Space precludes a detailed comparison of the different curve distance measures here, but we plan to present the details in a future publication.

F. Model Selection at the Neurite Level

The next step is to decide on the model that best describes the change between a given pair of neurites. One option is to use maximum likelihood estimation based on (4). However, in our case, we use the Bayesian maximum *a posteriori* (MAP) criterion, since we can obtain good estimates of the prior probabilities of the models and their parameters. Thus, we choose

$$m_k^{ij} = \arg \max_k \left\{ P(m_k(\theta_k)) p(\lambda_i^{t-1}, \lambda_j^t | m_k(\hat{\theta}_k^{ij}), B_t) \right\}. \quad (12)$$

The model priors $P(m_k(\theta_k))$ reflect initial domain-based beliefs about the likelihood of each model. We also need priors on any model parameters (e.g. the extension parameter in the growth model). For example, while the growth model itself may be likely, it would be unlikely that a segment grows ten times in size between two frames. Hence, we use a model prior of the form:

$$P(m_{\text{growth/shrinkage}}(s)) = P(m_{\text{growth/shrinkage}}) P(s | m_{\text{growth/shrinkage}}). \quad (13)$$

In Section IV-A, we describe how the priors are estimated based on manually labeled training data.

G. Model Selection at the Frame Level

Our goal is to find the set of matches and models that maximizes the total goodness of fit between the sets $\{\lambda_i^{t-1}, i = 1, \dots, N_{t-1}\}$ and $\{\lambda_j^t, j = 1, \dots, N_t\}$. We use a weighted bipartite graph [38] to solve this problem. As illustrated in Figure 6, a weighted bipartite graph G is represented by a set of vertices V and a set of undirected edges E connecting these vertices. V is the union of two disjoint sets, V_1 and V_2 , corresponding to the elements of $\{\lambda_i^{t-1}\}$ and $\{\lambda_j^t\}$, respectively. Each edge e_{ij} connects a vertex in V_1 with one in V_2 , and is associated with a positive weight w_{ij} . For an edge e_{ij} matching λ_i^{t-1} to λ_j^t , w_{ij} is set to the *a posteriori* probability of the change model chosen in (12). We again emphasize that in practice, the graph G is never complete, and is in fact quite sparse, containing only matches that are considered feasible based on distance considerations discussed below.

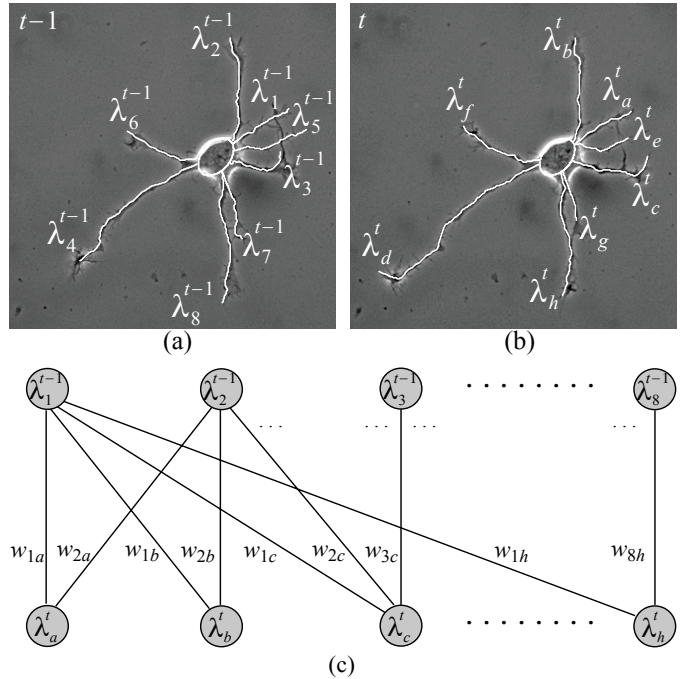


Fig. 6. Illustrating the bipartite graph method for establishing correspondences between neurites and change models. (a) and (b), sample sequence images taken at times $t-1$ and t , respectively, with neurite tracing ($\{\lambda_i^{t-1}\}$ and $\{\lambda_j^t\}$) and soma segmentation overlaid. (c) Bipartite graph used to match curve segments. The graph consists of two disjoint sets of vertices corresponding to $\{\lambda_i^{t-1}\}$ and $\{\lambda_j^t\}$, respectively. The weight associated with every edge, denoted w_{ij} , is set equal to the likelihood of the model that best describes the change between λ_i^{t-1} and λ_j^t .

In a bipartite graph, a *matching* is a subset of edges $M \subseteq E$ such that for all vertices $v \in V$, at most one edge of M is incident on v . A maximum weighted matching M^* is a matching such that the sum of edge weights is maximized:

$$M^* = \arg \max_M \left\{ \sum_{e \in M} w(e_{ij}) \right\}. \quad (14)$$

We use the standard Hungarian algorithm [39] to find the maximum weighted matching M^* .

We note that as stated above, the bipartite graph only allows for one-to-one matching between neurites. Whenever a merge/split change model $(\lambda_i^{t-1}, \lambda_j^{t-1}) \rightarrow \lambda_k^t$ is feasible, we create a new graph, G' , where the vertices corresponding to $(\lambda_i^{t-1}, \lambda_j^{t-1})$ in G are removed, as well as edges incident on them. A new vertex corresponding to λ_{ij}^{t-1} is added, and the edge connecting λ_{ij}^{t-1} with

λ_k^t is given a weight as described in Section III-D.3. The new graph G' is bipartite, and a maximum matching can again be found using the Hungarian algorithm. Given the two maximum matchings M^* and M'^* , we take the one with larger total weight as the solution. In practice, there are only a few merge/split hypotheses for the neurite application that pass the feasibility test. In particular, out of over 900 images used in the eight testing sequences, only 37 merge/split matchings were considered as feasible, out of which 28 were actually selected. Furthermore, there was never more than one merge/split hypothesis considered between a given image pair.

H. Derived Change Models

After the optimal matching of objects in each image and the determination of the basic change models describing each match (i.e. no change, growth/shrinkage, or split/merge), we can determine several additional semantic labels.

For example, if a segment λ_i^{t-1} was unmatched, we consider the intersection of the change mask B_t with a region surrounding λ_i^{t-1} . If the number of ones in this region is greater than the number of zeros (i.e. the mask suggests a change is present), then the segment is labeled “*deleted segment*”. Otherwise, it is labeled as “*incorrect neurite tracing*”. Similarly, an unmatched segment in I_t that has sufficient change mask evidence is labeled “*new segment*”, or otherwise “*incorrect neurite tracing*”. If one of a new or deleted segment’s end points is closer to an existing segment than twice the average neurite width, it is labeled as “*new/deleted branch*” instead.

I. Multi-Image Change Models Describing Longer-Term Events

Finally, we can build change models that combine frame-to-frame change results along the sequence to detect phenomena that unfold gradually. For example, a key phenomenon of interest is the development of cell polarity, or axonal specification (see Figure 9). When neurites first form, they cannot be distinguished as either axons or dendrites. Until the axon reaches a length at least 10-15 μm more than that of any other neurite stemming from the same soma, it is morphologically and molecularly indistinguishable from the other neurites [5], [13]. We can detect axonal specification by aggregating frame-level changes for all neurites stemming from each soma; the neurite that grows to become 15 μm longer than the rest is labeled as an axon, and the corresponding time index is reported.

A second phenomenon of interest is cell apoptosis (Figure 9). This process is characterized by the collapse of all the neurites stemming from a particular soma, which can be automatically detected when all the neurites associated with a soma are labeled as deleted.

IV. EXPERIMENTAL RESULTS

We developed and tested our algorithms on ten diverse image sequences of cultured neurons. Two image sequences were reserved for model training, and the remaining eight were used for testing. In the next section we describe how the different parameters used in the frame-level models were estimated from the training data. In Section IV-B, we present example analysis and graphical output automatically generated from the experiments to validate our method. We show that the semantic frame-to-frame changes of interest can be robustly and correctly detected, that these changes enable the accurate detection of axonal specification and apoptosis, and that the intelligent data reduction resulting from the proposed method is quite useful.

A. Model Training and Parameter Selection

As mentioned in Section III-D, the values of several parameters defining the model distributions must be estimated in order to

optimize performance of the overall framework. The key observation is that the same parameter values can be applied to any image sequence collected with similar conditions. Hence, the effort required to estimate the tuning parameters once can be tolerated when a large number of image sequences is collected under similar conditions, which is often true in hypothesis testing and assays.

We selected two training image sequences to estimate the various parameters used in our models. The first sequence consisted of 100 images of neurites growing on a smooth surface, while the second sequence consisted of 200 images of neurites growing on a patterned surface. The two sequences were selected based on visual inspection to ensure that many examples of the various changes of interest were present. For each adjacent training image pair, we manually traced the neurites, determined the neurite matches, and recorded the change model relating each match (i.e. no change, grow/shrink, or merge/split). We note that this process was extremely tedious, requiring more than nine hours to manually process both sequences, underscoring the need for a fast, automatic algorithm to do the same job.

We began by estimating the most appropriate sampling period for processing a neuronal image sequence. Ideally, we could analyze every frame in each sequence for changes. However, the tracing algorithm we use has a certain inherent precision; it can terminate prematurely before the neurite ends, or it can continue beyond the neurite tip by few pixels (usually due to the complex structure of growth cones). Hence, we must subsample the image sequences with an appropriate period Δt in order to avoid detecting “changes” below the tracing precision. We compared the automated tracing results with the manual ground truth, and estimated the tracing precision as the average number of pixels the automated tracing adds or misses, counting only from the neurite end. For the training data, the neurite tracing algorithm was found to have a precision of 8 pixels. Next, for a range of sampling periods Δt , we measured the average pixel difference in centerline traces (counting only pixels from the neurite end) corresponding to the same neurites between consecutive images. The first value of Δt for which the average difference was greater than the tracing precision was selected as the sampling period. For neurons cultured on patterned surfaces, the sampling period was estimated as $\Delta t = 8$, while for neurons cultured on smooth surfaces, the sampling period was estimated as $\Delta t = 6$.

The parameter δ used to weigh the two distributions in (4) was estimated to maximize the probability of making the correct frame-level model decision P_D , which is plotted as a function of δ for the training data in Figure 7-a. The value of $\delta = 0.7$ was found to maximize performance ($P_D = 0.91$) for the training sequences. This figure also reveals that if the complementary information provided by the change mask is ignored ($\delta = 1$), the performance drops to $P_D = 0.79$. Conversely, if the binary change mask is used as the sole evidence ($\delta = 0$), the performance is very poor ($P_D = 0.18$).

As shown in Figure 7-b, the distribution of the distances between manually matched no-change segments in (5) was best fit by an exponential distribution with parameter $\tau = 0.12$ (estimated with maximum-likelihood). Similarly, the distribution for the ratio r in (6) was best fit by an exponential distribution with parameter $\varphi = 1.8$ (Figure 7-c). To estimate the prior on the growth model parameter s in (13), a histogram of the number of growth/shrinkage pixels was fit to a Gaussian distribution with mean $\mu = 4.0$ and standard deviation $\sigma = 48$ (Figure 7-d). Finally, based on the manual labeling of the frame-level changes, we computed the model priors as $P(\text{no change}) = 0.46$, $P(\text{growth/shrinkage}) = 0.45$, and $P(\text{merge/split}) = 0.09$. Except for the sampling period, estimates for all model parameters were found to be similar under both smooth and patterned culturing conditions.




Fig. 7. Illustrating model training and parameter estimation. (a) δ in (4) was estimated to maximize the probability of correct frame-level model decision in the training data. (b-d) Normalized histograms of the observations were used to estimate the probability density functions that best describe them. (b) $d(\lambda_i^{t-1}, \lambda_j^t)$ for the no-change model was estimated to follow an exponential distribution with parameter $\tau = 0.12$; (c) the ratio r was also estimated to follow an exponential distribution with parameter $\varphi = 1.8$; (d) the prior on the growth model parameter s in (13) was estimated to follow a Gaussian with mean $\mu = 4$ and standard deviation $\sigma = 48$.

B. Experiments on Testing Data

All the image sequences used for validation of the algorithms were collected using a phase-contrast microscope. The pixel resolution for the collected images was $0.1067 \mu\text{m}$ along the x and y dimensions. Images were collected every 5 minutes. All of our experiments were implemented using unoptimized Matlab 7.0 code, except for the tracing algorithm, which was implemented in C++. All experiments were run on a Pentium III 1.8 GHz computer with 1 GB of memory. The algorithms are entirely automatic, and required no manual intervention or editing.

Table I reports the overall performance of the algorithm on each of the eight testing sequences. The median probability of correct frame-level model decision P_D was 96.5%. The detection performance results were computed with respect to manual ground truth generated by a single human expert aided by the automated tracing results. We note that quantitative ground-truthing of the detailed neurite morphometry would be incredibly tedious and time-consuming. For example, we estimate that manually tracing a neurite 150 pixels in length requires 25 seconds. Hence, Sequence C of 200 frames, with an average of 37 neurites per image, would require over 6 hours to manually trace every eighth frame, and a comparable amount of time to determine and quantify changes, whereas our algorithm does the same task in 43 minutes.

The reported execution times include the time required to run the entire algorithm on each sequence, and were found to increase roughly linearly with the number of neurites per image. This is the case because we impose spatial constraints on our feasible matches; only those segments within the same local neighborhood were considered for matching (i.e. the two nearest endpoints between λ_i^{t-1} and λ_j^t must be closer than fifty pixels). Without the spatial constraint, running times would increase quadratically with the number of segments.

Here, we provide a few examples of the graphical output and automatically extracted semantic descriptions from the testing dataset to indicate the utility of the technique. First, we consider Sequence A, consisting of 147 images, in which every sixth image was processed. This sequence was processed in 29 minutes. Two images from this sequence were illustrated in Panels a & b of Figure 1. The neurite tracing and soma segmentation output for Figure 1-a is illustrated in Figure 8-a, and the extracted semantic change description for the frame pair is illustrated in Figure 8-b. Table II-a is an accompanying high-level text summary generated by the proposed method, expressed in terms of the neurites and their behaviors. The first two columns indicate the corresponding neurite indices

as they appear in Panel 8-a. The third column shows the type of change automatically detected between the corresponding neurites. The value between parentheses is the *a posteriori* probability of change, as detailed in Section III-F. Table II-b illustrates the wealth of quantitative analysis results automatically generated by our algorithm that is immediately available to the user, using the second row of Table II-a as an example. Morphometric measurements are shown for each neurite, such as length, width, start, end, the originating soma, and curvature. The curvature reflects the neurite tortuosity; a line has zero curvature, whereas a circle has a curvature of one. Quantitative measurements of change are also presented, including the direction and amount by which the neurite shrank. These types of quantitative measurements, automatically generated by our method, would require significant effort to generate manually. In this sequence, an average of nineteen neurites were extracted by the neurite tracing algorithm in every frame. 85% of the detected neurite changes in the sequence were labeled correctly; a detailed analysis of the results is shown in Table II-c. A significant fraction of the errors reported in Table II-c were due to tracing errors; for example, short neurites stemming from the soma are sometimes not traced due to the halo artifact around the soma body.

Next, we consider Sequence B, consisting of 200 images of neurons cultured on a patterned surface, in which every eighth image was processed. This sequence was processed in 22 minutes. Two images from this sequence were illustrated in Panels c & d of Figure 1. The neurite tracing and soma segmentation output for Figure 1-c is illustrated in Figure 8-c, and the extracted semantic change description for the frame pair is illustrated in Figure 8-d. 89% of the detected neurite changes in this sequence were labeled correctly.

Figure 9 shows several frames from Sequence C, consisting of 200 images. An invisible biochemical grid pattern was imprinted on the surface, and growth patterns as well as interactions between the different neurons on this imprinted surface were studied. The goal was to detect sequence-level events that span more than a pair of consecutive images. In particular, two such events of special interest were automatically detected in this sequence: axonal specification and apoptosis. A total of four neurons were captured in the microscope's field of view. Somas and neurites were automatically extracted as illustrated in Figure 10-a, and neuronal changes were automatically detected and described between successive images. Neurite lengths were automatically measured as illustrated in Figure 10-b. Given the sequence of frame-to-frame changes and neurite lengths, axonal specification was automatically detected as described in Section III-I for two of the neurites (indexed A-1 and C-6) at times $t = 325$

TABLE I
PERFORMANCE SUMMARY OF THE EXPERIMENTS ON THE EIGHT TEST IMAGE SEQUENCES.

Sequence	# Frames	Frame Size	Avg. # Segments	P_D	Run Time (min)	Smooth/Patterned
A	147	1470 × 1180	19	85%	29	smooth
B	200	1300 × 1080	17	89%	22	patterned
C	200	1301 × 1080	37	87%	43	patterned
D	109	314 × 295	6	100%	6	smooth
E	60	474 × 484	7	99%	8	smooth
F	52	638 × 812	11	98%	10	smooth
G	70	348 × 438	12	95%	10	smooth
H	74	650 × 515	9	99%	9	smooth




Fig. 8. Illustrating sample semantic change analysis results for the image pairs in Figure 1. (a) Output of the neurite tracing and soma segmentation overlaid on the grayscale image of Figure 1-a. (b) Automatically generated graphical description of the changes between the image pair in Figure 1-a,b. 85% of the detected changes in this 147-frame sequence were labeled correctly. (c) Output of the neurite tracing and soma segmentation overlaid on the grayscale image of Figure 1-c. (d) Automatically generated graphical description of the changes between the image pair in Figure 1-c,d. 89% of the detected changes in this 200-frame sequence were labeled correctly.

TABLE II

SEMANTIC DESCRIPTION OF THE CHANGES BETWEEN THE PAIR OF IMAGE FRAMES SHOWN IN FIGURE 1-A,B. (A) TEXTUAL SUMMARY ACCOMPANYING THE GRAPHICAL OUTPUT IN FIGURE 8-B, INDICATING ONLY THE CHANGES. A “U” INDICATES A NEURITE UNASSOCIATED WITH ANY SOMA IN THE FIELD OF VIEW. (B) AUTOMATICALLY GENERATED EXCERPT DETAILING THE CHANGES FOR THE SECOND ROW IN (A). (C) CONFUSION MATRIX, THE ROW LABELS INDICATE THE TRUE CHANGE MODEL (DETERMINED MANUALLY), WHILE THE COLUMN LABELS INDICATE THE AUTOMATICALLY DETECTED CHANGE MODELS. THE ENTRIES IN THE MATRIX ARE THE NUMBER OF TIMES THAT A COMBINATION OF TRUE AND DETECTED MODELS OCCURRED IN THE EXPERIMENT.

Frame 1 segments	Frame 2 segments	Detected Change (<i>a posteriori</i> probability)
–	A-6	<i>tracing error</i> (0.023)
B-3	B-3	<i>shrinkage</i> (0.027)
B-9	B-9	<i>shrinkage</i> (0.033)
–	B-11	<i>tracing error</i> (0.019)
C-2, B-4	C-13	<i>merge</i> (0.045)
C-7	C-7	<i>shrinkage</i> (0.031)
–	C-12	<i>tracing error</i> (0.022)
–	C-13	<i>tracing error</i> (0.019)
U-6	U-6	<i>shrinkage</i> (0.036)
U-10	U-10	<i>growth</i> (0.039)
–	U-16	<i>tracing error</i> (0.025)

(a) Object-level Summary

Measurement	λ_i^{t-1}	λ_j^t
length (pixels)	149	97
avg. width (pixels)	7.6	7.3
start	(207, 514)	(207, 515)
end	(352, 410)	(298, 442)
soma index	B	B
total curvature	0.082	0.079
shrinkage length		51
shrinkage direction		126°

(b) Morphometric Detail

True	Detected				
	<i>no change</i>	<i>growth/shrinkage</i>	<i>merge/split</i>	<i>new/deleted</i>	<i>tracing error</i>
<i>no change</i>	194 (92%)	7	0	5	5
<i>growth/shrinkage</i>	18	218 (87%)	0	7	7
<i>merge/split</i>	0	0	2 (100%)	0	0
<i>new/deleted</i>	3	2	0	20 (62%)	7
<i>tracing error</i>	1	5	1	16	41 (64%)

(c) Confusion Matrix

and $t = 445$ min, respectively. In addition, the neuron indexed B was automatically detected to undergo apoptosis at time $t = 565$ min. Overall, an average of 37 segments were traced per frame in this sequence, and every eighth frame was processed for changes as described earlier. 87% of the detected neuronal changes in the sequence were labeled correctly, and the sequence was processed in 43 minutes.

V. CONCLUSIONS AND DISCUSSION

We demonstrated the practical feasibility of automatically generating concise, high-level descriptions of changes occurring in time-lapse imagery of cultured neurons. This enables far more rapid quantitative and qualitative understanding of the image changes compared to manual markup methods and simple computational strategies such as pixel-level change detection. Rapid and clear understanding is aided by data reduction, a high level of abstraction in the computer vision hierarchy (neurites instead of pixels), and analysis of both short and long time scales. Behind the concise summary is a wealth of detailed quantitative image analysis results that are available to the user. For example, when a curve segment is labeled as growing, the amount of growth is also measured, along with a list of these pixels, and the growth direction.

The high rate of success of the system in scoring the changes for the studied data sets indicates that it presents a significant improvement over the manual approaches in current use. Such robust automation can enable the neuroscientist to design far more ambitious studies. For instance, it is desirable to process a large number of neurons when a subtle neurobiological effect is being studied. The proposed approach can eliminate the subjectivity inherent to manual analysis and potentially improve both speed and accuracy. Though the system described here was not specifically optimized for speed,

it is still significantly faster than comparable manual scoring. Indeed, it runs more quickly than the rate of time-lapse data collection.

As demonstrated in [6], automatically generated traces tend to be more accurate compared to manual traces. On the other hand, automated tracing systems also tend to make more gross errors such as missing neurites/fragments, compared to manual tracing, especially in low-contrast, ambiguous image regions. For the studies described here, almost half of the (however few) change model errors ultimately arose from neurite tracing errors that in turn resulted from imaging artifacts such as the halo effect in phase-contrast microscopy. With other modalities, the sources of errors are likely to be different. For instance, use of fluorescence microscopy will likely yield tracing errors from non-uniform staining and/or noise. However, depending on the study, the longer-term events in the image sequence may be of more interest, and the detection of these should be robust to occasional model selection errors, as demonstrated here.

On a broader scale, the integrative framework, as depicted in Figure 2 and described in the text, is simple, extensible, modular, and adaptable to time-lapse image sequences from many other domains. Each of the major modules in Figure 2 can potentially be improved separately, or substituted entirely for a given application. If higher throughput is desired (e.g. for culture arrays), there are many opportunities for pipelined and/or parallel processing (e.g. evaluation of the change models).

In summary, this work not only presents the live-cell biologist with a practically useful tool that addresses a real need, but also enables future advances in the field. Future studies can be far more ambitious in terms of the numbers of cells analyzed, as well as the types of structural and functional measurements that can be made. The algorithms could facilitate high-throughput automated inspection systems to be used in a variety of areas including assay automation,




Fig. 9. Illustrating the multi-image events axonal specification and apoptosis. A sample showing nine frames from a sequence (Sequence C) with 200 frames collected for neurons cultured on a patterned surface. Two of the four neurons captured undergo axonal specification (arrows highlighted with the letter *s*), while a third neuron undergoes apoptosis (arrow highlighted with the letter *a*).

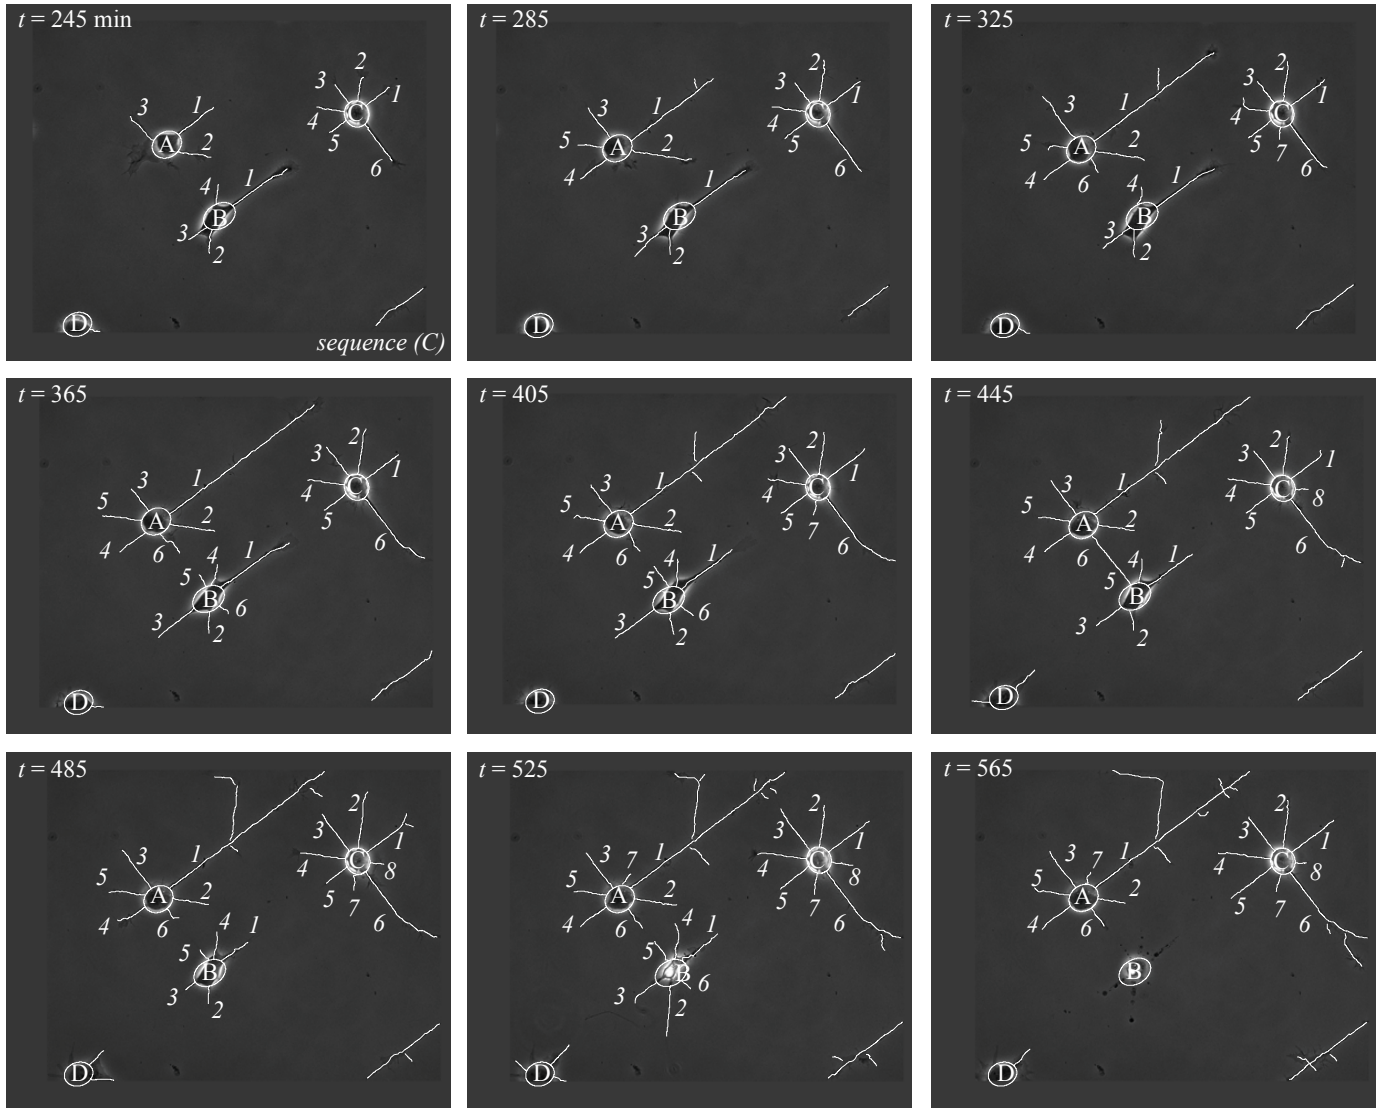
drug discovery, and optimization of culturing parameters for neuronal engineering. They also set the stage for quantifying the relationship between neuronal structure revealed by phase-contrast microscopy and a potentially unlimited number of functional indicators revealed by fluorescent labels. We have applied the same framework discussed here to multi-spectral (phase-contrast and fluorescence) image sequences to investigate the role of protein transport in axonal specification, and will report the results in a separate publication.

VI. ACKNOWLEDGMENTS


This work was supported in part by CenSSIS, the Center for Subsurface Sensing and Imaging Systems, under the Engineering Research Centers Program of the National Science Foundation (Award Number EEC-9986821), and by Rensselaer Polytechnic Institute. The authors would like to thank Harihar N. Iyer for his help in reviewing the paper.

REFERENCES

- [1] G. Banker and K. Goslin, Eds., *Culturing Nerve Cells*, 2nd ed. Cambridge, Massachusetts: Bradford Books, September 1998.
- [2] A. Bennett, H. Osterberg, H. Jupnik, and O. Richards, *Phase Microscopy: Principles and Applications*. New York: John Wiley and Sons, 1951.
- [3] B. Nathan, Y. Jiang, G. K. Wong, F. Shen, G. J. Brewer, and R. G. Struble, "Apolipoprotein E4 inhibits, and apolipoprotein E3 promotes neurite outgrowth in cultured adult mouse cortical neurons through the low-density lipoprotein receptor-related protein," *Brain Research*, vol. 928, no. 1-2, pp. 96–105, February 2002.
- [4] G. S. Withers, D. Higgins, M. Charette, and G. Banker, "Bone morphogenetic protein-7 enhances dendritic growth and receptivity to innervation in cultured hippocampal neurons," *European Journal of Neuroscience*, vol. 12, no. 1, pp. 106–116, January 2000.
- [5] T. Esch, V. Lemmon, and G. Banker, "Local presentation of substrate molecules directs axon specification by cultured hippocampal neurons," *The Journal of Neuroscience*, vol. 19, no. 15, pp. 6417–6426, August 1999.
- [6] K. A. Al-Kofahi, A. Can, S. Lasek, D. Szarowski, D. M. Natalie, W. Shain, J. N. Turner, and B. Roysam, "Median-based robust algorithms for tracing neurons from noisy confocal microscope images," *IEEE Trans. Inform. Technol. Biomed.*, vol. 7, no. 4, pp. 302–317, Dec 2003.
- [7] J. V. Pelt, I. Vajda, P. Wolters, M. Corner, and G. Ramakers, "Dynamics and plasticity in developing neuronal networks in vitro," in *Development, Dynamics and Pathology of Neuronal Networks From Molecules to Functional Circuits, Progress in Brain Research*, J. V. Pelt, M. Kamermans, C. N. Levelt, A. V. Ooyen, G. J. A. Ramakers, and P. R.



(a)



(b)

Fig. 10. Illustrating detection of axonal specification and apoptosis for Sequence C. (a) Neurite tracing and soma segmentation output for the images in Figure 9. (b) Neurite lengths as a function of time for the images in (a). Axonal specification and apoptosis are automatically detected from the frame-to-frame changes and neurite lengths.

TABLE III
GLOSSARY OF SYMBOLS USED.

Symbol	Description
I_t	Image captured at time t .
B_t	Binary change mask between I_{t-1} and I_t .
λ_i^t	i^{th} neurite centerline trace extracted by the neurite tracing algorithm for image I_t .
m_k	Change model k .
θ_k	Parameters for model k .
$\hat{\theta}_k^{ij}$	Parameters for model m_k estimated to map λ_i^{t-1} into λ_j^t .
m_k^{ij}	The change model selected to describe the change between λ_i^{t-1} and λ_j^t .
δ	Weight used for the mixture density for model selection in (4), estimated as $\delta = 0.70$ in our experiments.
$P(m_k(\theta_k))$	Prior for model m_k with parameters θ_k .
τ	Exponential parameter for the likelihood term that depends on the distance between curve segments. Estimated as $\tau = 0.12$ in our experiments.
φ	Exponential parameter for the likelihood term that depends on the change mask. Estimated as $\varphi = 1.8$ in our experiments.
s	Growth/shrinkage parameter. Indicates the number of pixels a curve segment grew or shrunk.
μ	Mean of the growth parameter s , estimated as $\mu = 4$.
σ	Standard deviation of the growth parameter s , estimated as $\sigma = 48$.
G, E, V	A graph, its edges, and its vertices, respectively.
e_{ij}, w_{ij}	A graph edge connecting vertex v_i to v_j and its non-negative weight.
M^*	Maximum weighted matching on a graph.
ω	Continuous and non-decreasing parametrization of a curve segment.

Roelfsema, Eds. Amsterdam: Elsevier, 2005, vol. 147.

- [8] Z. Zhang, R. Yoo, M. Wells, T. Beebe Jr., R. Biran, and P. Tresco, "Neurite outgrowth on well-characterized surfaces: preparation and characterization of chemically and spatially controlled fibronectin and RGD substrates with good bioactivity," *Biomaterials*, vol. 26, no. 1, pp. 47–61, January 2005.
- [9] Image Pro Plus, Media Cybernetics. San Diego, California, USA. [Online]. Available: <http://www.mediacy.com>
- [10] I. Rozovsky, M. Wei, T. E. Morgan, and C. E. Finch, "Reversible age impairments in neurite outgrowth by manipulations of astrocytic GFAP," *Neurobiology of Aging*, vol. 26, no. 5, pp. 705–715, May 2005.
- [11] L. Lin, Y. Rao, and O. Isacson, "Netrin-1 and slit-2 regulate and direct neurite growth of ventral midbrain dopaminergic neurons," *Molecular and Cellular Neuroscience*, vol. 28, no. 3, pp. 547–555, March 2005.
- [12] Image-1 software, Universal Imaging Corporation, Downingtown, PA, USA. [Online]. Available: www.image1.com
- [13] T. Esch, V. Lemmon, and G. Banker, "Differential effects of NgCAM and N-cadherin on the development of axons and dendrites by cultured hippocampal neurons," *Journal of Neurocytology*, vol. 29, no. 3, pp. 215–223, March 2000.
- [14] C. Kauffmann, P. Gravel, B. Godbout, A. Gravel, G. Beaudoin, J.-P. Raynauld, J. Martel-Pelletier, J.-P. Pelletier, and J. A. de Guise, "Computer-aided method for quantification of cartilage thickness and volume changes using MRI: Validation study using a synthetic model," *IEEE Trans. Biomed. Eng.*, vol. 50, no. 8, pp. 978–988, August 2003.
- [15] C. Meyer, H. Park, J. M. Balter, and P. H. Bland, "Method for quantifying volumetric lesion change in interval liver CT examinations," *IEEE Trans. Med. Imag.*, vol. 22, no. 6, pp. 776–781, June 2003.
- [16] Z. B. Sbeh, L. D. Cohen, G. Mimoun, and G. Coscas, "A new approach of geodesic reconstruction for drusen segmentation in eye fundus images," *IEEE Trans. Med. Imag.*, vol. 20, no. 12, pp. 1321–1333, December 2001.
- [17] C. R. Wren, A. Azarbayejani, T. Darrell, and A. Pentland, "Pfinder: Real-time tracking of the human body," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 19, no. 7, pp. 780–785, 1997.
- [18] I. Haritaoglu, D. Harwood, and L. S. Davis, "W⁴: Real-time surveillance of people and their activities," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 22, no. 8, pp. 809–830, August 2000.
- [19] N. M. Oliver, B. Rosario, and A. P. Pentland, "A Bayesian computer vision system for modeling human interactions," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 22, no. 8, pp. 831–843, August 2000.
- [20] R. Collins, A. Lipton, T. Kanade, H. Fujiyoshi, D. Duggins, Y. Tsin, D. Tolliver, N. Enomoto, and O. Hasegawa, "A system for video surveillance and monitoring," Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep., May 2000.
- [21] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, 2nd ed. John Wiley and Sons, Inc., 2001.
- [22] B. Zitova and J. Flusser, "Image registration methods: A survey," *Image and Vision Computing*, vol. 21, no. 11, pp. 977–1000, October 2003.
- [23] K. Toyama, J. Krumm, B. Brumitt, and B. Meyers, "Wallflower: Principles and practice of background maintenance," in *ICCV99*, 1999, pp. 255–261.
- [24] S. Haykin, *Adaptive Filter Theory*, 4th ed. Prentice Hall, 2001.
- [25] R. J. Radke, S. Andra, O. Al-Kofahi, and B. Roysam, "Image change detection algorithms: A systematic survey," *IEEE Trans. Image Processing*, vol. 14, no. 3, pp. 294–307, March 2005.
- [26] K. A. Al-Kofahi, S. Lasek, D. Szarowski, C. Pace, G. Nagy, J. N. Turner, and B. Roysam, "Rapid automated three-dimensional tracing of neurons from confocal image stacks," *IEEE Trans. Inform. Technol. Biomed.*, vol. 6, no. 2, pp. 171–187, June 2002.
- [27] I. Y. Koh, W. B. Lindquist, K. Zito, E. A. Nimchinsky, and K. Svoboda, "An image analysis algorithm for dendritic spines," *Neural Computation*, vol. 14, no. 6, pp. 1283–1310, June 2002.
- [28] E. Meijering, M. Jacob, J. Sarria, P. Steiner, H. Hirling, and M. Unser, "Design and validation of a tool for neurite tracing and analysis in fluorescence microscopy images," *Cytometry*, vol. 58A, no. 2, pp. 167–176, April 2004.
- [29] W. He, T. Hamilton, A. Cohen, T. Holmes, C. Pace, D. Szarowski, J. Turner, and B. Roysam, "Automated three-dimensional tracing of neurons in confocal and brightfield images," *Microsc Microanal.*, vol. 9, no. 4, pp. 296–310, August 2003.
- [30] G. Streekstra and J. van Pelt, "Analysis of tubular structures in three-dimensional confocal images," *Network*, vol. 13, no. 3, pp. 381–395, August 2002.
- [31] C. M. Weaver, P. R. Hof, S. L. Wearne, and W. B. Lindquist, "Automated algorithms for multiscale morphometry of neuronal dendrites," *Neural Computation*, vol. 16, no. 7, pp. 1353–1383, July 2004.
- [32] D. Jaeger, "Accurate reconstruction of neuronal morphology," in *Computational neuroscience: Realistic modeling for experimentalists*, E. De Schutter, Ed. Boca Raton, Fla. : CRC Press, 2001, ch. 6, pp. 159–178.
- [33] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Trans. Syst., Man, Cybern.*, vol. 9, no. 1, pp. 62–66, 1979.
- [34] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 2nd ed. New York: McGraw-Hill, 2001.
- [35] H. Alt and M. Godau, "Computing the the Fréchet distance between two polygonal curves," *International Journal of Computational Geometry and Applications*, vol. 5, no. 1, pp. 75–91, 1995.
- [36] G. Rote, "Computing the minimum Hausdorff distance between two point sets on a line under translation," *Information Processing Letters*, vol. 38, no. 3, pp. 123–127, 1991.
- [37] E. M. Arkin, L. P. Chew, D. P. Huttenlocher, K. Kedem, and J. S. B. Mitchell, "An efficiently computable metric for comparing polygonal shapes," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 13, no. 3, pp. 209–216, March 1991.
- [38] B. Korte and J. Vygen, *Combinatorial Optimization - Theory and Algorithms*, ser. Algorithms and Combinatorics. Springer, 2000, vol. 21.
- [39] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, pp. 83–87, 1955.



Omar Al-Kofahi Omar A. Al-Kofahi (S '03 - M '05) received the B.A. degree in Electrical Engineering from Jordan University of Science and Technology, Jordan, in 1998, and the M.A. and Ph.D. degrees from the Electrical, Computer and Systems Engineering Department, Rensselaer Polytechnic Institute, Troy, NY, in 2001 and 2005, respectively. His Ph.D. research was focused on developing an integrated change understanding framework that generates high-level descriptions of changes in time-lapse biomedical image sequences with applications for

changes in *in-vitro* neuronal and stem cell image sequences. His M.A. research was focused on two-view attenuation correction laser scanning confocal microscope image stacks, for which he received the Microscopy Society of America's student presidential award. During his Ph.D., he was an active member of the National Science Foundation Engineering Research Center for Subsurface Sensing and Imaging Systems (CenSSIS), where he served as the president of the student leadership council.

Dr. Al-Kofahi joined American Science and Engineering, Inc., Billerica, MA, in September, 2005, as a research scientist in image processing and understanding. His current research includes noise reduction and image deblurring in short- and long-range X-ray transmission and backscatter images.



Gary Banker Gary Banker received the B.S. degree in mathematics and psychology from the University of Washington in 1968 and the Ph.D. in neuroscience from the University of California at Irvine in 1973. He has been at Oregon Health and Science University since 1998, where he currently is Senior Scientist in the Center for Research on Occupational and Environmental Toxicology and Professor of Cell and Developmental Biology. He directs OHSU's Multidisciplinary Neuroscience Training Program, serves on the Executive Committee for the Nanobiotechnology Center at Cornell University, and is a member of the editorial board of Brain Cell Biology. Dr. Banker's research focuses on the development and maintenance of neuronal polarity, using novel methods for visualizing neuronal development and protein trafficking in living neurons in culture. Current projects aim to elucidate the molecular machinery that underlies the fidelity of protein targeting and its role in the specification of polarity during neuronal development. The recipient of a Javitz award from the National Institutes of Health for his research in neuroscience, Dr. Banker is a member of the Society for Neuroscience and the American Society for Cell Biology.

Dr. Banker's research focuses on the development and maintenance of neuronal polarity, using novel methods for visualizing neuronal development and protein trafficking in living neurons in culture. Current projects aim to elucidate the molecular machinery that underlies the fidelity of protein targeting and its role in the specification of polarity during neuronal development. The recipient of a Javitz award from the National Institutes of Health for his research in neuroscience, Dr. Banker is a member of the Society for Neuroscience and the American Society for Cell Biology.



Richard J. Radke Richard J. Radke (S '98 - M '01) received the B.A. degree in mathematics and the B.A. and M.A. degrees in computational and applied mathematics, all from Rice University, Houston, TX, in 1996, and the Ph.D. degree from the Electrical Engineering Department, Princeton University, Princeton, NJ, in 2001. For his Ph.D. research, he investigated several estimation problems in digital video, including the synthesis of photorealistic "virtual video", in collaboration with IBM's Tokyo Research Laboratory. He has also worked at the

Mathworks, Inc., Natick, MA, developing numerical linear algebra and signal processing routines.

He joined the faculty of the Department of Electrical, Computer, and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY, in August, 2001, where he is also associated with the National Science Foundation Engineering Research Center for Subsurface Sensing and Imaging Systems (CenSSIS). His current research interests include deformable registration and segmentation of three- and four-dimensional biomedical volumes, machine learning for radiotherapy applications, distributed computer vision problems on large camera networks, and modeling 3D environments with visual and range imagery. Dr. Radke received a National Science Foundation CAREER Award in 2003.



Badrinath Roysam Badrinath Roysam (M '89) received the B. Tech degree in electronics engineering from the Indian Institute of Technology, Madras, India, in 1984, and the M.S. and D.Sc. degrees from Washington University, St. Louis, in 1987, and 1989, respectively. He has been at Rensselaer Polytechnic Institute, Troy, New York since 1989, where he is currently a Professor in the Electrical, Computer and Systems Engineering Department. He is an associate director of the Center for Subsurface Sensing and Imaging Systems (CenSSIS) - a multi-university

NSF-sponsored engineering research center. He also holds an appointment in the Biomedical Engineering Department. He is an associate editor for the IEEE Transactions on IT for Biomedicine. His ongoing projects are in the areas of 2-D, 3-D, and 4-D biomedical image analysis, biotechnology automation, optical instrumentation, high-speed and real-time computing architectures, and parallel algorithms. Dr. Roysam is a member of IEEE, ASEE, the Microscopy Society of America, Society for Neuroscience, the Association for Research in Vision and Ophthalmology, and the Society for Molecular Imaging.