Detection of Mitosis within a Stem Cell Population of High Cell Confluence in Phase-Contrast Microscopy Images

Seungil Huh Robotics Institute Carnegie Mellon University 5000 Forbes Ave. Pittsburgh, PA15213 seungilh@cs.cmu.edu

Abstract

Computer vision analysis of cells in phase-contrast microscopy images enables long-term continuous monitoring of live cells, which has not been feasible using the existing cellular staining methods due to the use of fluorescence reagents or fixatives. In cell culture analysis, accurate detection of mitosis, or cell division, is critical for quantitative study of cell proliferation. In this work, we present an approach that can detect mitosis within a cell population of high cell confluence, or high cell density, which has proven challenging because of the difficulty in separating individual cells. We first detect the candidates for birth events that are defined as the time and location at which mitosis is complete and two daughter cells first appear. Each candidate is then examined whether it is real or not after incorporating spatio-temporal information by tracking the candidate in the neighboring frames. For the examination, we design a probabilistic model named Two-Labeled Hidden Conditional Random Field (TL-HCRF) that can use the information on the timing of the candidate birth event in addition to the visual change of cells over time. Applied to two cell populations of high cell confluence, our method considerably outperforms previous methods. Comparisons with related statistical models also show the superiority of TL-HCRF on the proposed task.

1. Introduction

Computer vision for automated analysis of cell populations has gained increasing attention due to its enormous potential for discoveries in cell biology and pharmacology as well as stem cell engineering [5, 11, 6]. Among such visionbased systems, the ones adopting phase-contrast time-lapse microscopy allow for long-term continuous monitoring of live and intact cells because phase-contrast microscopy is a non-destructive imaging modality. On the other hand, curMei Chen Intel Science & Technology Center Carnegie Mellon University 4720 Forbes Ave. Pittsburgh, PA15213 mei.chen@intel.com



Frame 130 Frame 132 Frame 134 Figure 1. Examples of mitotic cells in a bovine aortic endothelial cell (BAEC) population in phase-contrast microscopy images. The typical process of mitosis can be characterized by increased brightness and circularity, and decreased size. In frame 132, we can observe two birth events; birth event is defined as the time and

location at which two daughter cells first appear right after mitosis.

rent automated or semi-automated systems that use cellular staining methods do not permit long-term monitoring of live cells due to the use of fluorescence reagents or fixatives, which can cause phototoxicity or kill cells, respectively.

Mitosis is the process of cell division which is typically shown as in Figure 1. Mitosis detection is important for assessing proliferative activity of stem cells [16]. Automated mitosis detection in phase-contrast time-lapse microscopy can provide quantitative information regarding cell proliferation on a continuous basis, which is critical for studying biological phenomena of stem cells.

To date, several computer vision based mitosis detection methods have been proposed. Some use results from cell tracking [19, 2, 1, 14]. Mitotic cells are detected during tracking by examining change in visual properties of cells or by associating cell regions or their trajectories. These approaches are intuitive but limited in that mitosis detection is dependent on cell tracking, which is more challenging. Other approaches do not involve cell tracking [3, 11, 12, 7]. Most of these methods first reduce search space by detecting candidates that are likely to contain mitosis [3, 12, 7]. Probabilistic or statistical models are often employed to identify true mitosis among candidates by learning visual change of mitotic cells based on human-annotated samples [11, 12, 7].





Figure 2. Patches containing (a) positive and (b) negative birth event samples. A birth event is typically shown as a figure eight shape; but, such a shape is often not clearly observed. Negative samples are randomly extracted where the mean and standard deviation of pixel intensities in the patch is within the ranges of positive birth event samples.

Recently, one mitosis detection approach demonstrated success in precisely detecting birth events that are defined as the time and location at which mitosis is complete and two daughter cells first appear [7]. In this approach, candidate patch sequences are first constructed by linking the patches with increased brightness, which is a major visual characteristic of mitotic cells. A probabilistic model named Event Detection Conditional Random Field (EDCRF) is then devised to identify whether each candidate contains mitosis or not. For the candidates determined to contain mitosis, the birth event timing is simultaneously localized. Owing to the effectiveness of candidate patch construction and the discriminating power of EDCRF, this approach achieved good performance both in birth event detection accuracy and computational efficiency. However, as cell confluence, or cell density, increases, this method has difficulty in separately identifying individual mitotic cells, resulting in decreased performance. In fact, no existing approach has been successful at handling a cell population of high cell confluence.

To address this challenging problem, we present a mitosis detection approach that can separately identify individual mitosis in a cell population of high cell confluence where a mitotic cell is often in contact with other mitotic cells or non-mitotic cells with bright halos. We first detect candidates of birth events by initial screening and learning visual characteristics of birth events from a single image. Each candidate is then examined whether it is real after incorporating spatio-temporal information in neighboring frames. For this decision task, we develop a probabilistic model named Two-Labeled Hidden Conditional Random Field (TL-HCRF) that can model the temporal location of the candidate birth event in addition to the visual change of cells over time. Candidate birth event detection enables more accurate detection of individual mitosis at high cell confluence, resulting in significant improvement in birth event detection compared to previous work. Furthermore, the new model suited for the given task additionally boosts the performance. Empirical evidence on two cell populations of high cell confluence clearly demonstrates the superiority of our approach on the proposed task.

The remainder of this paper is organized as follows. We overview our approach and describe the details of the candi-

date extraction step in Section 2. We then formulate a probabilistic model to validate mitosis candidates in Section 3. The experiments and results are presented in Sections 4 and 5, respectively, followed by conclusion in Section 6.

2. Approach

Given a sequence of phase-contrast microscopy images, our goal is to detect birth events in the sequence. Specifically, we would like to find when (in which frame) and where (at which x and y positions) birth events occur. To achieve this goal, we form a process consisting of three steps: detection of candidate birth events, construction of candidate patch sequences, and identification of birth events, which are described in this and following sections. Before these steps, we apply a preprocessing method [7] to correct intrinsic illumination variation in phase-contrast microscopy images.

2.1. Candidate Birth Event Detection

In this step, we detect candidate birth events by initial screening and learning visual characteristics of birth events from a single frame. The goal of this step is more to reduce search space. Ideally, we expect to detect all actual birth events while detecting as few false positives as possible. In practice, sacrificing a small number of true mitosis that are outliers in morphology or intensity often eliminates a large number of false positives, resulting in higher overall performance on birth event detection.

Suppose that we have a training image sequence where all birth events are annotated and test image sequences where we would like to detect birth events. From the training sequence, we collect square patches extracted at each birth event as the center position as shown in Figure 2(a). The size of a patch is set to roughly enclose one mitotic cell. To increase the number of the positive samples as well as to achieve rotation invariance, each patch is duplicated by rotation at orientations 90°, 180° , and 270° . Through mirror reflection, additional four patches are generated; thus, a total 8 patches are produced from one birth event.

For efficient detection of birth event candidates, we first compute the statistics of the positive patches on intensity distributions and changes over time. Specifically, for each patch P located at (x, y) in frame t, we compute the mean and standard deviation of pixel intensities of the following three patches:

- P
- P subtracted from the patch at (x, y) in frame t 1
- P subtracted from the patch at (x, y) in frame t + 1

We use these statistics to initially remove the patches that are unlikely to contain birth events from both training and test sequences. Specifically, if any of these statistics of a square patch is not in the 99% confidence interval of the corresponding statistic of the positive patches, the patch is excluded for further consideration. Because pixel intensities increase as well as their change over consecutive frames is considerable during mitosis, these statistics are effective in excluding the patches not containing a birth event. These statistics can also be efficiently calculated for every patch in a given image through convolutions with an average filter.

To extract negative birth event patches from the training sequence, among the remaining patches after the initial screening, we randomly select patches whose center positions are not too close to a birth event as shown in Figure 2(b). The patches close to a birth event both spatially and temporally are excluded because they may cause confusion in detecting birth events due to their similar appearance to the positive samples. In our experiments, 30000 negative samples are randomly extracted. The rotation scheme is not applied since the number of negative samples is large enough. We exclude the patches whose center positions are located within spatially 5 pixels and temporally 5 frames from a birth event.

A Support Vector Machine (SVM) is then trained using the positive and negative birth event samples and applied to the patches in test sequences that have survived the initial screening. Unique scale histograms [7] are used as visual features. We set the decision boundary to correctly identify 99% of positive samples in the training set because the goal of this step is to reduce search space while minimizing false negatives. To efficiently handle a large number of training patches, SVM^{light} [8] was used for implementation. After the SVM is applied, multiple patches neighboring to a birth event are typically determined as positive patches that contain the birth event. Among the positive patches neighboring to one another, only one patch that is located at the centroid of the cluster is selected as the patch containing a candidate birth event and the other patches are discarded.

This candidate birth event detection enables precise identification of birth events when they occur in the vicinity of other cells at high cell confluence, or high cell density. Note that previous patch sequence construction methods [3, 12, 7] do not have such capability; as a result, their



Figure 3. An example of a candidate patch sequence. Each candidate patch sequence is constructed by tracking a candidate birth event before and after the birth event in the sequence. In this example, a candidate birth event is detected at time t.

performance may significantly decrease at high cell confluence.

2.2. Candidate Patch Sequence Construction

To accurately determine whether a candidate birth event is a true birth event, temporal information should be investigated because a birth event is defined as a specific event in mitosis, which is a process through several frames. The information from both previous and following frames is important because it provides a clue of birth event timing when a figure eight shape, which indicates separate two daughter cells, lasts for several frames or is vaguely observed.

To incorporate temporal information, we construct a candidate patch sequence by tracking each candidate birth event forward and backward in time as shown in Figure 3. We track candidate birth events rather than simply linking patches at the same location in the neighboring frames because mitotic cells often move by themselves or by external forces. For tracking, we link the nearest patches in terms of the Euclidean distance between the pixel values within the neighboring region in consecutive frames. In our experiments, we track five frames before and two frames after a birth event, obtaining a candidate patch sequence consisting of eight frames for each candidate birth event. The number of patches tracked is empirically determined based on birth event detection performance on the training samples.

2.3. Identification of Birth Events

Now the problem reduces to determining whether each candidate patch sequence contains an actual birth event given the most likely temporal location of the birth event. To solve this problem, we propose a probabilistic model named Two Labeled Hidden Conditional Random Field (TL-HCRF). The details of TL-HCRF will be described in the following section. Through modeling with TL-HCRF,



Figure 4. Graphical representations of two previous models (HCRF and EDCRF) and the TL-HCRF model. x_i and h_i represent the *i*th observation (the *i*-th patch in a given candidate sequence in this work) and the hidden state assigned on x_i , respectively. In HCRF, y represents mitosis occurrence. In EDCRF, y represents mitosis occurrence as well as birth event timing and sub-labels s_1, \dots, s_m represent birth event timing by their transition. In TL-HCRF, y represents mitosis occurrence and determines two sublabels s_1 and s_2 . Another label z represents birth event timing and determines the model connectivity. Gray circles denote observed variables for training. For testing, y and s are not observed. Note that z is observed as the timing of the candidate birth event for testing.

we can obtain the probabilities that each candidate birth event is a real one.

After applying TL-HCRF, we need to perform postprocessing to prevent duplicate detection of one birth event as well as to localize each birth event timing. As a figure eight shape may last for several frames, candidate birth events typically capture not only a birth event, but also its previous or following stages. To find the most probable birth event timing among them, we first build a graph in which each node represents a candidate birth event. An edge is added between two nodes if the corresponding candidate birth events are located at a neighboring location in consecutive frames. We then select the candidate birth event that has the greatest probability in each connected graph and examine whether the candidate is real or not. The other candidates are considered to be extracted from the previous or following stages of the same birth event and thus ignored.

3. Two Labeled Hidden CRF

In this section, we formulate the Two Labeled Hidden CRF (TL-HCRF) model and describe its learning and inference processes.

3.1. Related models

We review two existing probabilistic models: Hidden Conditional Random Fields (HCRF) [17] and Event Detection Conditional Random Fields (EDCRF) [7], both of which have been used for mitosis detection. The graphical representations of these two models are shown in Figure 4.

HCRF, which was originally applied to speech and gesture recognition [4, 18], has been applied to mitosis detection and demonstrated its superiority to Support Vector Machines (SVM) and Conditional Random Fields (CRF) [12]. Since HCRF has only one label variable which represents the occurrence of mitosis, it is intrinsically not capable of modeling birth event timing when mitosis occurs. EDCRF was devised to additionally model birth event timing based on the idea of HCRF. In EDCRF, the timing of a birth event is modeled as a sub-class transition. By adding this timing information, EDCRF is able to not only determine the timing of birth events, but also more accurately detect mitosis occurrence than HCRF. EDCRF is, however, not suitable for our task because it assumes that a birth event can occur at any patch in a candidate patch sequence. In other words, EDCRF is not designed to use the information on candidate birth event timing, which is helpful to determine the real birth event timing.

3.2. TL-HCRF formulation

In our task, for each candidate patch sequence, we are given the most likely timing for a birth event, which is the timing of the candidate birth event in the sequence. To utilize this information, we propose Two Labeled Hidden Conditional Random Field (TL-HCRF), the graphical representations of which is shown in Figure 4.

TL-HCRF has two label variables y and z; given a sequence x, y indicates whether x contains a birth event and z represents the event timing. More formally, in the training phase,

$$y = \begin{cases} 1 & \text{if } \mathbf{x} \text{ contains a birth event} \\ 0 & \text{otherwise} \end{cases}$$
(1)

$$z = \begin{cases} p & \text{if the } p \text{-th patch of } \mathbf{x} \text{ contains a birth event} \\ q & \text{if there exists no birth event in } \mathbf{x} \end{cases}$$
(2)

In the test phase, y is not observed and

$$z = q \tag{3}$$

where the q-th patch contains a candidate birth event. In other words, if a birth event exists in the given sequence, z represents its timing; if a birth event does not occur or the information is not known, z represents the timing that a birth event most likely occurs. Note that p is not necessarily equal to q since p is determined by ground truth. q is determined by the number of frames tracked before the candidate birth event in the candidate patch sequence construction step. If more than one mitosis are contained in a candidate patch sequence, the mitosis whose timing is closer to the candidate birth event in the sequence is used to set p.

Suppose that **x** consists of *m* patches; i.e., **x** = (x_1, x_2, \dots, x_m) where x_j denotes the *j*-th patch (*m* can be varied for different sequences.). We assume hidden variables $\mathbf{h} = (h_1, h_2, \dots, h_m)$ where $h_j \in \mathcal{H}$ corresponds to x_j and \mathcal{H} is a set of hidden states in the model. We also assume two sub-labels s_1 and s_2 . Given z, s_1 is connected with h_1, h_2, \dots, h_{z-1} while s_2 is connected with h_z, h_{z+1}, \dots, h_m as shown in Figure 4. According to y, the sub-labels s_1 and s_2 are set as

$$\begin{cases} s_1 = N, \ s_2 = N & \text{if } y = 0 \\ s_1 = B, \ s_2 = A & \text{if } y = 1 \end{cases}$$
(4)

where label N, B, and A represent No event, Before the event, and After the event (including the event), respectively. In other words, if there exists a birth event in a given candidate sequence, the patches and hidden variables before the event are associated with a sub-label B while those after the event with a sub-label A. Otherwise, all the patches and hidden variables are associated with a sub-label N.

Under these definitions, we define a latent conditional model:

$$P(y|\mathbf{x}, z; \theta) = P(s_1, s_2|\mathbf{x}, z; \theta) = \sum_{\mathbf{h}} P(\mathbf{h}, s_1, s_2|\mathbf{x}, z; \theta)$$
(5)

where θ is a set of parameters of the model.

We define $P(\mathbf{h}, s_1, s_2 | \mathbf{x}, z; \theta)$ using a log linear model as follows:

$$P(\mathbf{h}, s_1, s_2 | \mathbf{x}, z; \theta) = \frac{1}{Z} exp(\Psi(\mathbf{h}, s_1, s_2, \mathbf{x}, z; \theta)) \quad (6)$$

where Z is a partition function which is defined as

$$Z = \sum_{\substack{(s_1, s_2) \in \\ \{(N,N), (B,A)\}}} \sum_{\mathbf{h} \in \mathcal{H}^m} exp\big(\Psi(\mathbf{h}, s_1, s_2, \mathbf{x}, z; \theta)\big) \quad (7)$$

We define $\Psi(\mathbf{h}, s_1, s_2, \mathbf{x}, z; \theta)$ by extending the formulations of the conditional random field (CRF) and the hidden conditional random field (HCRF).

$$\Psi(\mathbf{h}, s_1, s_2, \mathbf{x}, z; \theta) = \sum_{j=1}^m f^{(s)}(h_j, \mathbf{x}, j) \cdot \theta^{(s)}(h_j)$$
$$+ \sum_{j=2}^{z-1} f^{(t)}(h_{j-1}, h_j, \mathbf{x}, j) \cdot \theta^{(t)}(h_{j-1}, h_j, s_1, s_1)$$

$$+ f^{(t)}(h_{z-1}, h_z, \mathbf{x}, k) \cdot \theta^{(t)}(h_{z-1}, h_z, s_1, s_2) + \sum_{j=z+1}^{m} f^{(t)}(h_{j-1}, h_j, \mathbf{x}, j) \cdot \theta^{(t)}(h_{j-1}, h_j, s_2, s_2) + \sum_{j=1}^{z-1} \theta^{(l)}(h_j, s_1) + \sum_{j=z}^{m} \theta^{(l)}(h_j, s_2)$$
(8)

where $f^{(s)}(h_j, \mathbf{x}, j)$ and $f^{(t)}(h_{j-1}, h_j, \mathbf{x}, j)$ are state and transition functions, respectively. $\theta^{(s)}$ and $\theta^{(t)}$ are the parameters of state and transition functions, respectively, and $\theta^{(l)}$ is the parameter associated with sub-labels; thus, $\theta = \{\theta^{(s)}, \theta^{(t)}, \theta^{(l)}\}.$

Following previous works [17, 7], state and transition functions are defined as

$$f^{(s)}(h_j, \mathbf{x}, j) = \phi(x_j), \quad f^{(t)}(h_{j-1}, h_j, \mathbf{x}, j) = 1$$
 (9)

where $\phi(x_j)$ is a visual feature vector of x_j . We use unique scale gradient histograms [7] as visual features in our experiments.

3.3. Learning model parameters

Suppose that *n* pairs of candidate patch sequences and corresponding labels $\{(\mathbf{x_1}, (y_1, z_1)), (\mathbf{x_2}, (y_2, z_2)), \cdots, (\mathbf{x_n}, (y_n, z_n))\}$ are given for training. For learning parameters, we maximize the following regularized log-likelihood function [10, 9].

$$L(\theta) = \sum_{i=1}^{n} \log P(y_i | \mathbf{x_i}, z; \theta) - \frac{1}{2\sigma^2} ||\theta||^2$$
(10)

where σ is the variance of a Gaussian prior.

This optimization problem can be solved by gradient ascent methods. Let $\theta^{(s)}[h]$ be the parameter vector in $\theta^{(s)}$ that corresponds to a certain hidden state h and $\theta_k^{(s)}[h]$ be the k-th element of $\theta^{(s)}[h]$. Then, the derivative of $\log P(y|\mathbf{x}, z; \theta)$ with respect to $\theta_k^{(s)}[h]$ is computed as

$$\frac{\partial \log P(y|\mathbf{x}, z; \theta)}{\partial \theta_k^{(s)}[h]} = \sum_{j=1}^m P(h_j = h|s_1, s_2, \mathbf{x}, z; \theta)\phi_k(x_j)$$
$$- \sum_{\substack{(s_1, s_2) \in \\ \{(N,N), (B,A)\}}} \sum_{j=1}^m P(h_j = h, s_1, s_2|\mathbf{x}, z; \theta)\phi_k(x_j)$$
(11)

where $\phi_k(x_j)$ is the k-th element of $\phi(x_j)$. $P(h_j = h|s_1, s_2, \mathbf{x}, z; \theta)$ can be computed by belief propagation [15] in O(m) [13].

Let $\theta^{(t)}[h', h'', s', s'']$ be the parameter in $\theta^{(t)}$ that corresponds to two sub-labels s' and s'' and two hidden states h' and h''. Then the derivative of $\log P(y|\mathbf{x}, z; \theta)$ with respect

to $\theta^{(t)}[h^\prime,h^{\prime\prime},s^\prime,s^{\prime\prime}]$ is computed as

$$\frac{\partial \log P(y|\mathbf{x}, z; \theta)}{\partial \theta^{(t)}[h', h'', s', s'']}$$
(12)

$$= \sum_{j=2}^{m} P(h_{j-1} = h', h_j = h''|s_1, s_2, \mathbf{x}, z; \theta) g_1(s', s'', j, y, z_j)$$

$$- \sum_{\substack{(s_1, s_2) \in \\ \{(N, N), (B, A)\}}} \sum_{j=2}^{m} P(h_{j-1} = h', h_j = h'', s_1, s_2 | \mathbf{x}, z; \theta)$$

where $g_1(s', s'', j, y, z)$ is defined as

$$g_{1}(s', s'', j, y, z) =$$

$$\begin{cases} 1 & \text{if } y = 0 \text{ and } (s', s'') = (N, N) \\ 0 & \text{if } y = 0 \text{ and } (s', s'') \neq (N, N) \\ 0 & \text{if } y \neq 0 \text{ and } (s', s'') = (N, N) \\ I(j < z) & \text{if } y \neq 0 \text{ and } (s', s'') = (B, B) \\ I(j = z) & \text{if } y \neq 0 \text{ and } (s', s'') = (B, A) \\ I(j > z) & \text{if } y \neq 0 \text{ and } (s', s'') = (A, A) \end{cases}$$

$$(13)$$

where I(x) is an indicator function; i.e., I(x) is 1 if x is true, 0 otherwise. Note that when optimizing the parameter for the transition function, we consider four different sub-label transitions, namely, $(s', s'') \in \{(N, N), (B, B), (B, A), (A, A)\}$, which represent no event, before the event, during the event, and after the event, respectively. Similar to the optimization of $\theta^{(s)}, P(h_{j-1} = h', h_j = h''|s_1, s_2, \mathbf{x}, z; \theta)$ can also be efficiently computed by belief propagation.

Lastly, let $\theta^{(l)}[h, s]$ be the parameter in $\theta^{(l)}$ that corresponds to a sub-labels s and a hidden states h. Then the derivative of $\log P(y|\mathbf{x}, z; \theta)$ with respect to $\theta^{(l)}[h, s]$ is computed as

$$\frac{\partial \log P(y|\mathbf{x}, z; \theta)}{\partial \theta^{(l)}[h, s]}$$

$$= \sum_{j=1}^{m} P(h_j = h|s_1, s_2, \mathbf{x}, z; \theta) g_2(s, j, y, z)$$

$$- \sum_{\substack{(s_1, s_2) \in \\ \{(N, N), (B, A)\}}} \sum_{j=1}^{m} P(h_j = h, s_1, s_2 | \mathbf{x}, z; \theta)$$
(14)

where $g_2(s, j, y, z)$ is defined as

$$g_{2}(s, j, y, z) =$$
(15)
$$\begin{cases}
1 & \text{if } y = 0 \text{ and } s = N \\
0 & \text{if } y = 0 \text{ and } s \neq N \\
0 & \text{if } y \neq 0 \text{ and } s = N \\
I(j < z) & \text{if } y \neq 0 \text{ and } s = B \\
I(j \ge z) & \text{if } y \neq 0 \text{ and } s = A
\end{cases}$$

Using the derivatives in Eqs. (11), (13), and (15), we can find the optimal parameter θ^* that maximizes the objective in Eq. (10).

3.4. Inferences

For the *i*-th patch sequence in test set, in addition to \mathbf{x}_i , we set z_i to be the timing of the candidate birth event. Therefore, we need to infer y_i given \mathbf{x}_i and z_i , which is a binary classification problem. The label is determined to be

$$y_i^* = \begin{cases} 1 & \text{if } P(y_i = 1 | \mathbf{x}_i, z_i; \theta^*) > 0.5 \\ 0 & \text{otherwise} \end{cases}$$
(16)

This conditional distribution, which can be reformulated by Eq. (5), can be computed by belief propagation. After y_i^* is obtained for each *i*-th patch, we perform the postprocessing described in the previous section to finalize detected birth events.

4. Experiments

We tested mitosis detection algorithm on two stem cell populations of high cell confluence: C2C12 myoblastic stem cells (C2C12) and bovine aortic endothelial cells (BAEC).

4.1. Data and Ground truth

For C2C12 and BAEC sequences, images were acquired every 5 minutes for approximately 84 and 19 hours, obtaining 1013 and 228 images, respectively.¹ Manual annotation of birth events was performed on both image sequences. For each birth event, the center of the boundary between two daughter cells was marked when the boundary is first clearly observed. As a result, 673 and 422 mitotic events were annotated in the C2C12 and BAEC sequences, respectively. In addition, we manually tracked each mitotic cell for three frames both before and after the birth event to include the case that a birth event is detected with a small timing error less or equal to three frames.

For both C2C12 and BAEC sequences, we tested our approach on the last 100 frames, and the rest frames were used for model training. Assessed by two biologists, the last 100 frames of C2C12 start with a confluence of $65 \sim 75\%$ and end with a confluence of $80 \sim 90\%$ while the BAEC sequence is at 100% confluence from the beginning to the end.

¹In detail, for C2C12, during the growth of stem cells, microscopy cell images were acquired using a Zeiss Axiovert T135V microscope (Carl Zeiss Microimaging, Thornwood, NY) equipped with a 5X, 0.15 N.A. phase-contrast objective, a custom-stage incubator, and the InVitro software (Media Cybernetics Inc., Bethesda, MD). For BAEC, the open area, or wound area, in each dish is observed with a Leica DMI 6000B inverted microscope equipped with a 10X objective with phase optics until neighboring cells completely fill.

	Threshold	th=1		th=3		
C2C12	Approach	[7]	Our approach	[7]	Our approach	
	F-measure	0.746	0.800	0.794	0.845	
	(Precision/Recall)	(0.755/0.738)	(0.883/0.731)	(0.803/0.786)	(0.933/0.772)	
	AUC	0.672	0.706	0.758	0.761	
	Threshold	th	=1	th	=3	
	Threshold Approach	[7]	=1 Our approach	[7]	=3 Our approach	
BAEC	Threshold Approach F-measure	[7] 0.351	=1 Our approach 0.716	[7] 0.372	=3 Our approach 0.761	
BAEC	Threshold Approach F-measure (Precision/Recall)	th [7] 0.351 (0.454/0.286)	=1 Our approach 0.716 (0.796/0.650)	th [7] 0.372 (0.480/0.303)	=3 Our approach 0.761 (0.847/0.692)	

Table 1. Mitosis detection performance comparison between the previous approach [7] and the proposed approach on C2C12 and BAEC in terms of F-measure and AUC of the PR-curve. Detection results are considered true positive when the timing error in terms of the number of frames is not greater than a given threshold (th=1 or 3). Our approach outperforms the previous approach, particularly on BAEC, which shows much higher cell confluence than C2C12.

4.2. Evaluation and Comparison

To measure the performance of our approach, we compute precision and recall on the conditions of birth event location and timing. Specifically, if a birth event is detected spatially within 10 pixels and temporally within th frames (th is one of 1 and 3 in our experiments) from the manual annotation, it is considered correctly detected (true positive). A small timing error is allowed because there is a few frames' variation even among human annotations. If a detection result does not satisfy these conditions, the detection is regarded incorrect (false positive). If one birth event is detected more than once, the detection results except the one closest to the birth event are considered incorrect (false positive).

We compare our approach with the previous work [7], which constructs candidate patch sequences by thresholding brightness without the candidate birth event detection step. To the best of our knowledge, [7] is the only work that explicitly detects birth events during mitosis. We also compare our TL-HCRF model with HCRF [18] and ED-CRF [7] on the task that classifies candidate patch sequences whether each of them contains a birth event or not. To compare the detection results, we compute F-measure, which is the harmonic mean of precision and recall, and the area under the curve (AUC) of the precision-recall curve (PR-curve). PR-curves are obtained by varying the decision probability in Eq. (16).

5. Results

Our approach significantly outperforms the previous approach [7] in terms of F-measure and AUC of the PR graph as shown in Table 1. In particular, performance improvement is significant on BAEC which is at 100% confluence from the beginning to the end. The low performance of the previous approach on BAEC clearly indicates that it is challenged by high cell confluence. In cell populations of high cell confluence, where cells are closer to each other, mitotic

cells are often in contact with other mitotic cells or nonmitotic cells with bright halos. In such cases, the candidate patch construction scheme used in the previous work tends to identify such clustered cells as one entity rather than discerning each of them as an individual entity. As a result, the accuracy of birth event detection in terms of both occurrence and timing is hampered. On the other hand, the proposed candidate patch sequence construction scheme is more capable of detecting each individual cell because each candidate birth event, which typically shows a figure eight shape, is first detected by a learning-based method. Candidate birth event detection preceding candidate patch sequence construction significantly reduces false negative or undetected mitosis compared to the previous work.

Table 2 shows that TL-HCRF outperforms the previous two models, HCRF and EDCRF, on the proposed task. To produce these results, we replaced TL-HCRF with HCRF or EDCRF without changing any other settings. HCRF is inferior to TL-HCRF on the given task because HCRF is not designed to use the information on birth event timing. The higher performance of HCRF than EDCRF on C2C12 can be ascribed to the postprocessing step, which finds the most probable birth event timing based on the confidence levels of candidate birth events neighboring to one another. ED-CRF can model the timing of a birth event, but it assumes that a birth event can occur at any patch rather than referring to candidate birth event timing. Therefore, in our framework which involves candidate birth event detection, ED-CRF has excessive modeling power than necessary, which is realized by sub-labels assigned on each observation. To make the inference of all sub-labels tractable, EDCRF restricts that each sub-class label is associated only with hidden states in a disjoint set [7], which may degrade the overall performance. Unlike these two previous models, TL-HCRF has appropriate modeling power for the given task by assuming that a birth event occurs at the candidate birth event patch if mitosis occurs; as a result, TL-HCRF does

C2C12 myoblastic stem cells (C2C12)										
Threshold	Threshold th=1			th=3						
Model	HCRF	EDCRF	TL-HCRF	HCRF	EDCRF	TL-HCRF				
F-measure	0.679	0.643	0.800	0.793	0.765	0.845				
(Precision/Recall)	(0.704/0.655)	(0.674/0.614)	(0.883/0.731)	(0.822/0.766)	(0.803/0.731)	(0.933/0.772)				
AUC	0.578	0.529	0.706	0.724	0.689	0.761				
Bovine aortic endothelial cells (BAEC)										
		Bovine aortic	endothelial cells	(BAEC)						
Threshold		Bovine aortic th=1	endothelial cells	(BAEC)	th=3					
Threshold Model	HCRF	Bovine aortic th=1 EDCRF	endothelial cells TL-HCRF	(BAEC) HCRF	th=3 EDCRF	TL-HCRF				
Threshold Model F-measure	HCRF 0.613	Bovine aortic th=1 EDCRF 0.647	endothelial cells TL-HCRF 0.716	(BAEC) HCRF 0.685	<i>th</i> =3 EDCRF 0.726	TL-HCRF 0.761				
Threshold Model F-measure (Precision/Recall)	HCRF 0.613 (0.622/0.604)	Bovine aortic th=1 EDCRF 0.647 (0.641/0.654)	endothelial cells TL-HCRF 0.716 (0.796/0.650)	(BAEC) HCRF 0.685 (0.695/0.675)	<i>th</i> =3 EDCRF 0.726 (0.718/0.733)	TL-HCRF 0.761 (0.847/0.692)				

Table 2. Mitosis detection performance comparison among the three probabilistic models: HCRF, EDCRF, and TL-HCRF on C2C12 and BAEC in terms of F-measure and AUC of the PR-curve. For this comparison, each model is applied for the binary classification of candidate patch sequences generated by our proposed approach. Detection results are considered true positive when the timing error of birth event is not greater than a given threshold (th=1 or 3). TL-HCRF outperforms the other two previous models in terms of both F-measure and AUC.

not require additional restrictions, achieving the best performance.

6. Conclusion

In this paper, we present an effective method that can detect mitosis in a cell population of high cell confluence, which has proven challenging due to the difficulty in identifying individual cells. We propose a method that detects the candidates of the time and location at which two daughter cells first appear, which enables separate detection of mitotic cells. We also design a new probabilistic model named Two-Labeled Hidden Conditional Random Field (TL-HCRF) that can examine each candidate whether it is real or not after incorporating spatio-temporal information. Experimental results on two stem cell populations clearly demonstrate that our method considerably outperforms previous methods for mitosis detection in a stem cell population of high cell confluence.

References

- O. Al-Kofahi, R. Radke, S. Goderie, Q. Shen, S. Temple, and B. Roysam. Automated cell lineage construction: A rapid method to analyze clonal development established with murine neural progenitor cells. *Cell Cycle*, 5(3):327–335, 2006.
- [2] O. Debeir, P. V. Ham, R. Kiss, and C. Decaestecker. Tracking of migrating cells under phase-contrast video microscopy with combined mean-shift processes. *IEEE Trans. Med. Imaging*, 24(6):697–711, 2005.
- [3] O. Debeir, V. Mégalizzi, N. Warzée, R. Kiss, and C. Decaestecker. Videomicroscopic extraction of specific information on cell proliferation and migration *in virto*. *Exp. Cell Res.*, 314(16):2985–2998, 2008.
- [4] A. Gunawardana, M. Mahajan, A. Acero, and J. Platt. Hidden conditional random fields for phone classification. In *Proc. Interspeech*, pages 1117–1120, 2005.
- [5] A. Hand, T. Sun, D. Barber, D. Hose, and S. Macneil. Automated tracking of migrating cells in phase-contrast video microscopy sequences using image registration. J. Microsc., 234(1):62–79, 2008.

- [6] D. House, M. Walker, Z. Wu, J. Wong, and M. Betke. Tracking of cell populations to understand thier spatio-temporal behavior in response to physical stimuli. In Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshop on Mathematical Methods in Biomedical Image Analysis, pages 186–193, 2009.
- [7] S. Huh, D. Ker, R. Bise, M. Chen, and T. Kanade. Automated mitosis detection of stem cell populations in phase-contrast microscopy images. *IEEE Trans. Med. Imaging*, 30(3):586–596, 2011.
- [8] T. Joachims. Making large-Scale SVM Learning Practical. Advances in Kernel Methods - Support Vector Learning. MIT-Press, 1999.
- [9] S. Kumar and M. Herbert. Discriminative random fields: A framework for contextual interaction in classification. In *Proc. International Conference on Computer Vision*, pages 1150–1157, 2003.
- [10] J. Lafferty, A. McCallum, and F. Pereira. Conditional random fields: probabilistic models for segmenting and labelling sequence data. In *Proc. International Conference on Machine Learning*, pages 282– 289, 2001.
- [11] K. Li, E. Miller, M. Chen, T. Kanade, L. Weiss, and P. Campbell. Computer vision tracking of stemness. In *Proc. IEEE International Symposium on Biomedical Imaging*, pages 847–850, 2008.
- [12] A.-A. Liu, K. Li, and T. Kanade. Mitosis sequence detection using hidden conditional random fields. In *Proc. IEEE International Symposium on Biomedical Imaging*, 2010.
- [13] L. Morency, A. Quattoni, and T. Darrell. Latent-dynamic discriminative models for continuous gesture recognition. In *Proc. Computer Vision and Pattern Recognition*, pages 1–8, 2007.
- [14] D. Padfield, J. Rittscher, N. Thomas, and B. Roysam. Spatiotemporal cell cycle phase analysis using level sets and fast marching methods. *Med. Image Anal.*, 13(1):143–155, 2009.
- [15] J. Pearl. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann, 1988.
- [16] H. Quastler and F. Sherman. Cell population kinetics in the intestinal epithelium of the mouse. *Exp. Cell Res.*, 17(3):420–438, 1959.
- [17] A. Quattoni, S. Wang, L. Morency, M. Collins, and T. Darrell. Hidden conditional random fields. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(10):1848–1853, 2007.
- [18] S. Wang, A. Quattoni, L.-P. Morency, D. Demirdjian, and T. Darrell. Hidden conditional random fields for gesture recognition. In *Proc. Computer Vision and Pattern Recognition*, pages 1521–1527, 2006.
- [19] F. Yang, M. Mackey, F. Ianzini, G. Gallardo, and M. Sonka. Cell segmentation, tracking, and mitosis detection using temporal context. In Proc. International Conference on Medical Image Computing and Computer Assisted Intervention, pages 302–309, 2005.