

Cell morphology classification and clutter mitigation in phase-contrast microscopy images using machine learning

Diane H. Theriault · Matthew L. Walker ·
Joyce Y. Wong · Margrit Betke

Received: 15 August 2010 / Revised: 25 March 2011 / Accepted: 9 May 2011 / Published online: 11 June 2011
© Springer-Verlag 2011

Abstract We propose using machine learning techniques to analyze the shape of living cells in phase-contrast microscopy images. Large scale studies of cell shape are needed to understand the response of cells to their environment. Manual analysis of thousands of microscopy images, however, is time-consuming and error-prone and necessitates automated tools. We show how a combination of shape-based and appearance-based features of fibroblast cells can be used to classify their morphological state, using the Adaboost algorithm. The classification accuracy of our method approaches the agreement between two expert observers. We also address the important issue of clutter mitigation by developing a machine learning approach to distinguish between clutter and cells in time-lapse microscopy image sequences.

Keywords Microscopy imaging · Cell morphology · Adaboost · Machine learning

D. H. Theriault · M. Betke (✉)
Department of Computer Science, Boston University,
111 Cummington St., Boston, MA 02215, USA
e-mail: betke@cs.bu.edu

D. H. Theriault
e-mail: deht@cs.bu.edu

M. L. Walker
Department of Biology, Boston University,
111 Cummington St., Boston, MA 02215, USA
e-mail: mwalk79@bu.edu

J. Y. Wong
Department of Biomedical Engineering, Boston University,
111 Cummington St., Boston, MA 02215, USA
e-mail: jywong@bu.edu

1 Introduction

Phase-contrast microscopy is a well-established tool to image living cells in research concerning the engineering of biomaterials. Studies of in vitro cell behavior typically involve thousands of cells. Interpreting these vast amounts of data via manual analysis is time-consuming, costly, and prone to human error.

In response to the growing need for automated image analysis tools, machine vision methods have recently been developed for detection and tracking of cells in time-lapse microscopy video (e.g. [1, 10, 14, 15]). Segmentation of cells can be extremely challenging and has been addressed by many authors (e.g. [5, 16, 19, 21, 26]).

In early work on classification of cell shape, Olson et al. [20] showed how machine learning techniques known at that time, i.e., hierarchical clustering and nearest neighbor analysis, could be used to classify cells into three classes by interpreting various shape descriptors. More recent classification techniques, in particular, support vector machines, have been applied by Han et al. [8] to the problem of detecting cell nuclei and by Ruusuvaori et al. [23] for classifying rod, spherical and spiral bacteria. Bradhurst et al. [3] evaluated four machine learning algorithms, naïve Bayes, multi-layer perceptrons, radial basis function networks, and support vector machines, to evaluate bone marrow stromal cell culture quality. There is also an extensive literature regarding the use of automated classification techniques to identify sub-cellular structures in fluorescence microscopy images [2, 12, 17, 18].

Our contributions are as follows:

- We present a method based on the machine learning technique Adaboost [6] to classify cell shape in the presence of significant intra-class variation.




Fig. 1 Two cells undergoing changes in shape. The time stamps of the images are 0:00, 1:00, and 2:00 h (*top*) and 0:00, 0:50, and 1:15 h (*bottom*)

- We show that we can distinguish living cells from substrate clutter using an Adaboost classifier.
- We illustrate the relationship between segmentation and classification accuracy.
- We provide a solution to the particular problem of analyzing living Balb/c mouse fibroblast cells on tissue-culture plastic substrates imaged using phase-contrast microscopy.

In a dataset containing over 800 cells in 125 images, we demonstrate classification accuracy comparable to the agreement achieved between expert human observers when applied on manually segmented images, and a 3–9 percentage point reduction in performance when automatic segmentation is applied. We have also applied our technique to 30 time-lapse sequences to identify clutter and classify cell shape. We show a significant reduction in the number of erroneous detections, with only a small reduction in true detections. We demonstrate that even though our segmentation technique has known limitations, our downstream analysis of cell morphology classification produces reliable results.

2 Methods

2.1 Challenges in analyzing cell morphology

Living cells respond to environmental signals that regulate behaviors, such as cell differentiation, division, and migration. Biomaterial scientists and engineers study cell

migration and changes in cell morphology to understand the influence of different environmental signals. The behavior of vascular smooth muscle cells and fibroblast cells is of particular interest to engineers who investigate cell-substrate interactions in vitro [25]. The morphological characteristics of a fibroblast cell depend on its current state and can change during its lifetime (Fig. 1).

Initially, cells are defined to be “non-spread.” Their image is characterized by a small area, rounded cell perimeter, and high contrast due to reflection of light off of the spherical membrane (Fig. 2, top left). This appearance is also typical for cells that are undergoing cellular division or apoptosis (programmed cell death). The transition from the “non-spread” to the “spread” state of a cell involves the polymerization of cytoskeletal elements that push out against the cell membrane, increasing the area of the cell in the image. The transition process yields a cell that has an either “polarized” or “non-polarized” morphology. Non-polarized cells are characterized by a smooth (rounded) cell membrane (Fig. 2, top right), whereas polarized cells have extended protrusions supported by elements of the internal cytoskeleton (Fig. 2, bottom left). Cells in a polarized state can further be subdivided into cells with and without orientation. When the protrusions are evenly distributed around the cell, the cell has no orientation. If the cell is strongly polarized in one direction and thereby shows a significant level of bilateral symmetry, it may be defined as “oriented” (Fig. 2, bottom right).

The appearance of cells belonging to a single morphological class varies significantly (Fig. 3). This makes the problem




Fig. 2 Examples of different cell shapes: non-spread (*top left*), spread and non-polarized (*top right*), spread and polarized (*bottom left*) and polarized with orientation (*bottom right*)

of automatically recognizing the state of a cell in an image very challenging.

Another difficulty in designing automated tools for cell image analysis is the fact that non-cell objects may appear in phase-contrast microscopy images. These clutter objects may be contaminations of the substrate and may or may not move and change appearance (Fig. 4).




Fig. 4 Several examples of non-cell clutter objects (*boxes*)

2.2 Creation of cell morphology library

Studying the behavior of slow-moving cells requires observations of many cells over long periods of time (18–24 h) at a granularity that permits observation of the nuances of how the cells are moving and changing shape. Fibroblast cells of the Balb/c 3T3 mouse strain (American Type Culture Collection, Manassas, VA, USA) were cultured at 37°C in 5% CO₂ in Dulbecco's modified Eagles medium (Invitrogen, Grand Island, NY, USA) supplemented with 50 µg mL⁻¹ penicillin, 50 U mL⁻¹ streptomycin, 200 mM L-glutamine and 10% bovine calf serum (Hyclone, Logan, UT, USA). In each experiment, we seeded the fibroblast cells onto a tissue culture plastic substrate at a density of 10,000 cells per cm² and allowed them to adhere for 24 h.




Fig. 3 Examples of cells in the spread and non-polarized (*top*) and spread and polarized (*bottom*) states

Imaging was carried out using an Zeiss Axiovert S100 phase-contrast microscope (Carl Zeiss, Thornwood, NY, USA) equipped with a motorized stage (Ludly, Hawthorne, NY, USA), a custom incubator box designed to maintain 5% CO₂ and a humidified environment, a cooled CCD camera (Princeton Instruments, Trenton, NJ, USA) and Metamorph Imagine software. Image dimensions were 1,030 by 1,300 pixels, and pixel size was 0.6614 μm².

Substrates with cells were imaged at 30-s intervals for 16–24 h. We used the first 600 images (5 h) from each of 30 resulting image sequences to create a “time-lapse migration library”. A cell biologist manually tracked the 258 cells contained in these 30 sequences, using every tenth frame (5 min). In every annotated frame, each cell was labeled with a set of attribute labels indicating whether the cell was spread, polarized, and/or oriented.

We also assembled a “cell morphology library” with the goal of representing the variation of cell shape within each morphological class. Two cell biologists localized and classified cells in 125 images. We used labels of cells identified by the biologists to establish benchmarks for our machine learning experiments. One of the two biologists also manually traced the boundary of 883 of the cells, providing “ground-truth” manual segmentation of cells in the images.

2.3 Image segmentation

Automated segmentation of cells in microscopy images is a challenging task, and various approaches have been proposed (e.g., [5, 16, 19, 21]). One of the challenges of segmenting cells in phase-contrast microscopy images is that the intensity values of the pixels representing the inside of the cell are very similar to the values representing the outside of the cell, and so applying intensity thresholds is not effective. Fortunately, cells are often at least partially surrounded by a “halo” (e.g., the bright vertical lines along the cell in Fig. 3, top right), and so edge-based methods, such as active contours, are useful. Since segmentation is not the focus of this paper, we use a simple segmentation approach, which we briefly discuss here for completeness.

Our segmentation method uses a multi-scale approach. The input image $I(x, y)$ is downsampled by a factor of 2^{-k} , where $k = \{1, 2, 3\}$. The magnitudes of the intensity gradient $\nabla I = \left(\frac{\partial I(x,y)}{\partial x}, \frac{\partial I(x,y)}{\partial y} \right)$ are computed at each scale and upsampled to create gradient magnitude images G_k at the original image size. The magnitude of the intensity gradient is then approximated by

$$G^*(x, y) = \max_k \{ G_k(x, y) \mid k \in \{1, 2, 3\} \}, \quad (1)$$

by the value of the intensity change that is largest among the original image and the two downsampled versions. The

motivation for this approach is that it automatically selects the resolution scale in which an edge is most pronounced. An adaptive threshold is then applied [7] to $G^*(x, y)$ which produces a binary mask $M_G(x, y)$ that defines image locations with significant intensity changes.

To supplement the edge-based technique, we adopted the “standard-deviation transform” used by Bradhurst et al. [3] for segmenting bone marrow cells. This method uses the observation that, while the background of the image is fairly uniform, regions inside of cells have a varied appearance. The standard-deviation transform $S(x, y)$ is computed by calculating the standard deviation of intensity values in the region $I(x \pm m, y \pm m)$ of L1-distance m surrounding pixel (x, y) . Using the “integral image” [24], we compute the transform $S(x, y)$ efficiently. For our experiments, we chose $m = 5$.

A threshold is automatically computed using the p-tile thresholding method [13] with $p = 92\%$ and applied to $S(x, y)$ to yield a binary mask $M_S(x, y)$ that indicates the locations in the image that are surrounded by pixels of high intensity variations.

A final segmentation mask $M(x, y)$ is formed by combining binary masks $M_G(x, y)$ and $M_S(x, y)$ with an OR operation. If $M(x, y) = 1$, the corresponding pixel $I(x, y)$ belongs to an object of interest. We then apply connected component labeling [9], for which $M(x, y) = 1$ and fill holes. These regions correspond to segmented objects of interest, i.e., cells or clutter. We denote a segmented region by r and the set of segmented regions by \mathbf{R} .

Our segmentation method is vulnerable to the halo effect commonly found in phase-contrast microscopy images and there is also a tendency to cut off thin, extended protrusions (Fig. 5). However, we show in the remainder of this paper that we are still able to obtain good classification results.

2.4 Feature extraction

To apply our classification method, we compute 18 shape-based and 21 appearance-based features for each segmented image region $r \in \mathbf{R}$. The shape-based features are computed using only the binary mask $M(x, y)$, whereas the appearance-based features use the intensity image $I(x, y)$. The first 13 shape-based features are functions of image moments [9]. We use the first two moments, i.e., object area A and centroid (\bar{x}, \bar{y}) , and a circularity measure C that is based on the central moments

$$v_{pq} = \sum_{(x,y) \in r} (x - \bar{x})^p (y - \bar{y})^q, \quad (2)$$

where $p, q \in \{0, 1, 2\}$. We define circularity by $C = m - n / (m + n)$, where $m = \frac{1}{2}(v_{20} + v_{02})$ and $n = \frac{1}{2}(\sqrt{4v_{11}^2 + (v_{20} - v_{02})^2})$. We also use the first seven Hu moments [11], which are functions of the normalized




Fig. 5 Manual (left) and automatic (right) segmentation of three cells. Left to right the automatic segmentation method identified 49.0, 94.1, and 99.4% of the pixels included in the manual segmentation, and 65.4,

63.6 and 51.8% of the pixels identified by the automatic segmentation were not included in the manual segmentation

moments $\eta_{pq} = v_{pq}/A^{(1+\frac{p+q}{2})}$, $p, q \in \{0, 1, 2, 3\}$. In addition, we incorporate measures of extension, dispersion, and elongation [4], which are functions of the first two Hu moments.

The remaining five shape-based features are derived from characteristics of the boundary \mathbf{B} , the set of pixels in region r that have neighboring pixels not belonging to r . Boundary \mathbf{B} is defined using an 8-connected neighborhood definition:

$$\mathbf{B} = \left\{ (x, y) \in r \mid \sum_{d_x=-1}^1 \sum_{d_y=-1}^1 M(x+d_x, y+d_y) < 9 \right\}. \tag{3}$$

The Euclidean distance between the j th point of the boundary \mathbf{B} and the centroid of the region is given by $d_j = \sqrt{(x_j - \bar{x})^2 + (y_j - \bar{y})^2}$. We include the mean, standard deviation, minimum, and maximum of the distances, computed for all points on the boundary, in the shape-based feature vector.

The length of a cell boundary can be represented by $|\mathbf{B}|$, the sum of L1-distances between consecutive pixels on the boundary, or by the perimeter P , the sum of the L2-distances. Because some consecutive pixels on the boundary may be diagonal to each other, $P \neq |\mathbf{B}|$, in general. We use P to compute the “shape factor,” or isoperimetric quotient, $\frac{4\pi A}{P^2}$, which is the ratio of the area of the region and that of a circle with the same perimeter, i.e., $\frac{A}{P^2}/\frac{1}{4\pi}$.

The appearance-based features are statistics of the image $I(x, y)$, approximations of the intensity gradient $G_k(x, y)$, as computed by our multi-scale approach (Sect. 2.3), and approximations of the Laplacian

$$L(x, y) \approx \frac{\partial^2 I(x, y)}{\partial x^2} + \frac{\partial^2 I(x, y)}{\partial y^2} \tag{4}$$

of the image, which are also computed by a multi-scale approach. In particular, the appearance-based feature vector for each region r includes the mean, standard deviation, and

skew of $I(x, y)$, $G_k(x, y)$, and $L_k(x, y)$ with $k = \{1, 2, 3\}$, such that $(x, y) \in r$.

2.5 Design of machine learning methods and analysis methodology

Adaboost [6] is a popular machine learning algorithm that combines weak classifiers to create a strong classifier. The training phase of the Adaboost algorithm requires as input a set of labeled feature vectors, a set of weak classifiers, and a parameter that determines the number of weak classifiers to be used. During training, Adaboost produces as output a set of chosen classifiers and associated weights. In the prediction phase, the Adaboost algorithm takes as input an unlabeled feature vector, computes a score, and applies a threshold to produce a binary label for the vector. A threshold of zero is used when it can be assumed that class labels are equally likely.

We have chosen to use a bank of binary naïve Bayes classifiers, one for each element of the feature vector, as weak classifiers. Each classifier in the bank operates on a single element v_i of the feature vector. Given the data, the probability that a sample belongs to a particular class c is given by Bayes rule: $p(c|v_i) = \frac{p(v_i|c)p(c)}{p(v_i)}$. A class label is applied using the likelihood ratio test. We chose to model $p(v_i|c)$ as a univariate Gaussian probability density, and we assigned equal prior probabilities to the class labels. To train the i th naïve Bayes classifier, we computed the parameters of the model for $p(v_i|c)$ for each class, i.e., mean and standard deviation.

As described in Sect. 2.1, the classification of the shape of a cell may be characterized by a set of three labels, indicating whether the cell is “spread” or “non-spread”, “polarized” or “non-polarized”, and “oriented” or “non-oriented”. We trained separate, independent classifiers for each attribute label, and then combined the three predicted labels to produce the final classification. When testing our automated classification technique, we used the labels applied by one of the

expert observers. Accuracy metrics and receiver operating curves (ROC) were computed using tenfold cross validation.

We compute 2×2 confusion matrices for each attribute, in addition to an 8×8 confusion matrix for the final, combined classification. We report “accuracy” for each attribute as the sum of the diagonal of each 2×2 confusion matrix. We also report the “overall accuracy” as the sum of the diagonal of the 8×8 confusion matrix. For the final classification label, we compute a “weighted accuracy” score which gives credit for partially correct answers. Given N labeled examples, L different attributes for each example ($L = 3$), and T and C , the sets of true and predicted labels, respectively, we define a weighted accuracy function

$$W_8 = \frac{1}{N \cdot L} \sum_{i=1}^N \sum_{l=1}^L \delta(C_i(l), T_i(l)), \quad (5)$$

where $\delta(\alpha, \beta) = 1$ if $\alpha = \beta$ and zero otherwise. For example, if the true class of a single cell is (spread/ non-polarized/ non-oriented), but it is classified as (spread/ polarized/ non-oriented), $W_8 = \frac{2}{3}$.

2.6 Tracking applied to time-lapse migration library

After segmentation, a simple tracker that uses a greedy data association method is used to link detected objects across frames to create tracked objects. The tracker does not use state estimation but represents each object track by the sequence of centroids of the detected objects. At each frame, the matrix that represents the costs of associations between the existing tracks and new detections is computed. No gating is used because there are no more than a few dozen tracks or detections at any one time. We use a sigmoidal cost function whose argument is the distance between the centroid of the detection and the previous position of the track.

For each track, the detection with the lowest association cost is identified. Similarly, for each detection, the track with the lowest association cost is identified. If the results match, the track is appended with the new detection. If not, no association is made. If a detection does not match any track, it is used to start a new track. If a track is not associated with a detection, it may be allowed to “coast” forward for no more than 10 min. Tracks that are less than 2:30 min long are not allowed to coast. Tracks that are not allowed to coast forward are terminated. Tracks that last for less than 2:30 min are discarded as spurious.

3 Results

In this section, we first present the morphology classification results obtained for cells in the morphology library. Next, we report the classification results of our clutter identification

strategy on a training dataset. Finally, we present the impact of our clutter mitigation strategy and the effectiveness of our morphology classification methodology in our analysis of the time-lapse migration library.

3.1 Morphology classification results

We first describe our benchmark: the agreement between two expert human observers. We next present our classification results when given the benefit of ground-truth manual segmentation. Then, we report our classification results when using a fully automated recognition system combining automated segmentation with automated classification. In Sect. 6, we present results obtained when pursuing two alternate classification methodologies.

According to the labels assigned by one of the expert observers, the morphology library contains the following distribution of class labels: 82.8% of the examples were spread, 60.5% of the examples were polarized, and 31.9% of the examples were oriented. According to the other expert observer, 84.0% of the examples were spread, 63.9% of the examples were polarized, and 17.0% of the examples were oriented.

3.1.1 Benchmark establishment: inter-observer classification agreement

Our first set of results concerns the establishment of a performance benchmark for our machine learning experiments. It is noteworthy that the two expert human observers, who labeled the same 838 cells in the morphology library, were in complete agreement for less than 60% of the cells.

Both observers applied three labels to each cell, one for each attribute (spread, polarized, and oriented). We combined these labels to obtain a single classification for each cell and created an 8×8 confusion matrix (Fig. 6, left). Examination of the confusion matrix reveals significant structure. For example, the observers might disagree on the spread attribute, and while cells classified as non-spread by one observer may be classified as spread by the other, they are not typically classified as both spread and polarized by the second observer. The accuracy measured for each attribute is summarized in Table 1, column 1.

3.1.2 Classification results using manual segmentation

Using the ground-truth segmentation established by one of the cell biologists, we computed a feature vector for each of 883 cells, using the 39 features described in Sect. 2.4. Using manual segmentation and all available features, our accuracy, when compared to one of the two experts, is as good as the agreement observed between the human observers (Table 1, column 2).

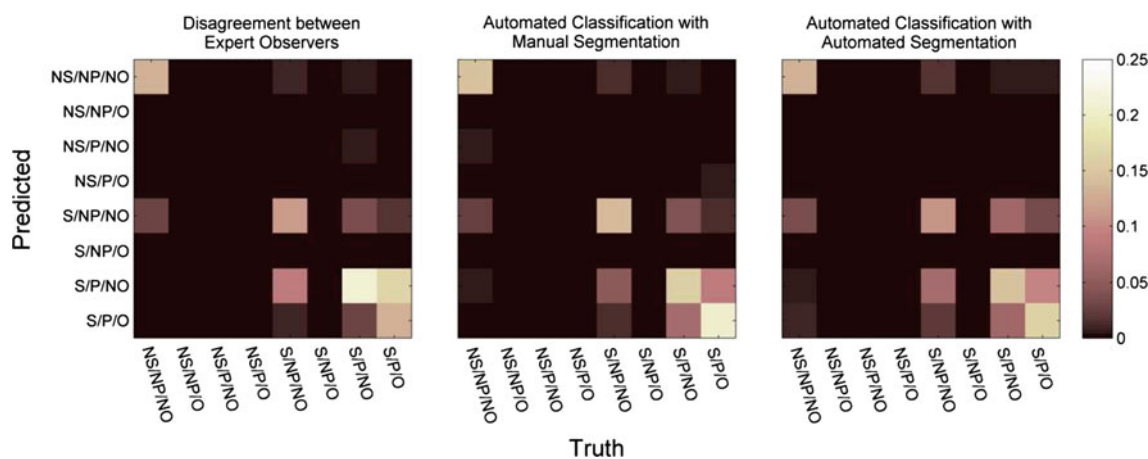


Fig. 6 We show the disagreement between human observers on 838 instances (*left*), classification accuracy on 883 instances using manual segmentation, compared to one of the expert observers (*middle*), and classification accuracy on 841 instances using automated segmentation, compared to one of the expert observers (*right*). Classification accuracy

scores were obtained using tenfold cross validation. Each example is classified by combining three binary attribute labels: spread (*S*) or non-spread (*NS*), polarized (*P*) or non-polarized (*NP*), and oriented (*O*) or non-oriented (*NO*)

Table 1 Agreement between the cell labels determined by expert observers (column 1), accuracy of automated classification when compared to labels of one of the expert observers (columns 2 and 3), and

accuracy of automated classification when compared to the cells where the labels applied by both expert observers were the same (columns 4 and 5)

	Agreement between expert observers	Classification accuracy (one observer)		Classification accuracy (both observers)	
		Manual segmentation	Automated segmentation	Manual segmentation	Automated segmentation
No. of cell instances	838	883	841	490	476
Overall accuracy (%)	58.5	65.3	54.5	72.0	62.4
Weighted accuracy (%)	84.6	86.7	81.1	90.3	86.2
Spread attribute (%)	93.3	93.9	90.1	98.2	95.2
Polarized attribute (%)	83.3	85.6	76.9	91.2	84.9
Oriented attribute (%)	77.2	80.6	76.1	81.4	78.6

Accuracy of automated classification determined via tenfold cross validation, using shape-based and appearance-based features computed from manual or automated segmentations

As can be seen by inspection (Fig. 6, middle), the structure of the confusion matrix seen when using automated classification is similar to that found between the humans. There are some differences; e.g., the algorithm classified some non-spread cells as both spread and polarized, whereas cells classified as non-spread by one expert observer were not typically classified as spread *and* polarized by the other.

We also applied the classification procedure to the set of 490 cells where the labels applied by both expert observers were in complete agreement (Table 1, columns 4 and 5). When we used only these samples, our classification results improved for all metrics. For example, the overall classification accuracy based on automated segmentation improved by 7.9 percentage points. For the remaining experiments, we tested our classification procedure using only the labels applied by the first expert observer.

We tested the performance of our methodology using different combinations of features. We compared using shape-based features, appearance-based features, and a combination of both (Fig. 7, top). Inspection shows that the performance obtained using only shape-based features is comparable to the results obtained using all features, whereas the performance obtained when using only appearance-based features suffers significantly, particularly for the polarized and oriented attributes.

3.1.3 Classification results using automated segmentation

In practical application, our classification strategy will not have the benefit of manual segmentation. Shapes resulting from automated segmentation are less precise than manual segmentation. An important question is how well our classification strategy works, given these known shortcomings.

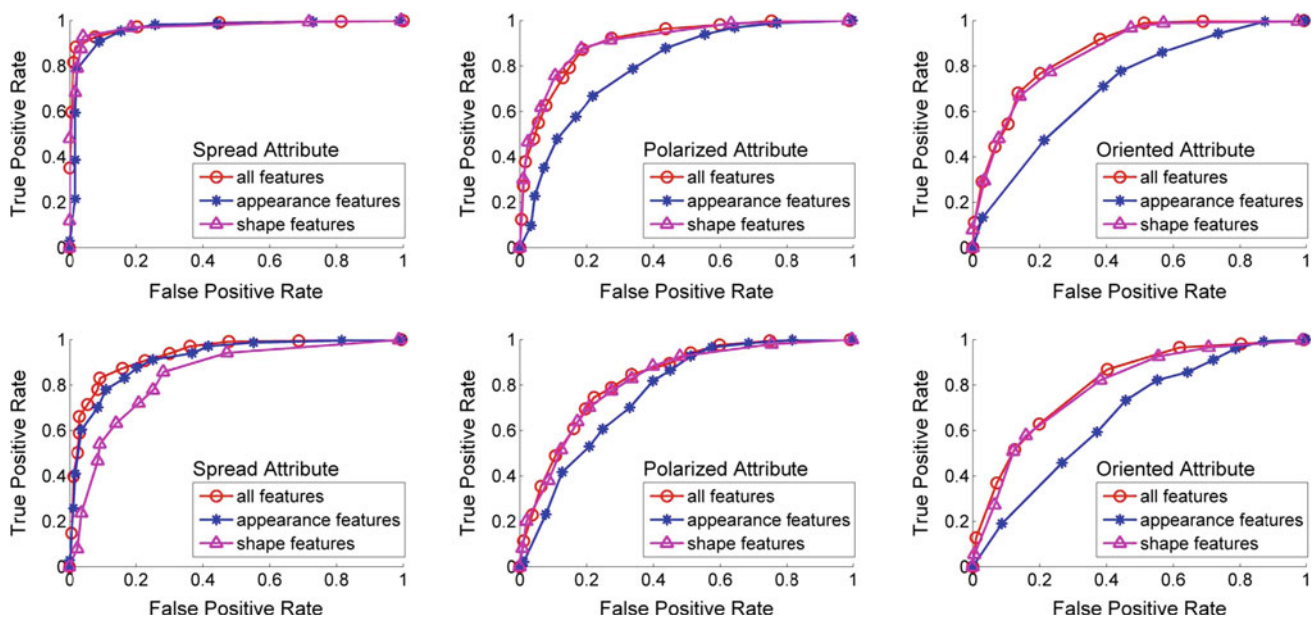


Fig. 7 ROC curves showing the performance of the Adaboost classifier on the spread (*left*), polarized (*middle*), and oriented (*right*) attributes when used with manual segmentation (*top*) or automated segmentation (*bottom*). We compare performance achieved when using shape-based

features, appearance-based features, or a combination of both. Results were obtained by comparing classification results with the labels produced by one of the expert observers, using tenfold cross validation

We applied our segmentation technique to the 125 images in the morphology library. The ground-truth labels used in the Sect. 3.1.2 were aligned with the automatically segmented regions by determining which region contained each labeled object. Due to the imperfect segmentation, some of the truth labels had no associated detected region, some detected regions had no associated truth label, and some detected regions had multiple associated truth labels. These objects were excluded from our classification experiment, leaving 841 examples.

Overall, the accuracy we obtained using automatic segmentation was 10.8 percentage points less than the accuracy obtained using manual segmentation, when using both shape-based and appearance-based features (Table 1, column 3). The decrease in performance for the weighted accuracy, and the accuracy for each individual attribute was between 3 and 9 percentage points. The polarized attribute suffered the worst loss at 8.7 percentage points. The structure of the confusion matrix obtained using automatic segmentation is similar to the structure found using manual segmentation (Fig. 6, right).

As before, we tested the performance of our classification methodology using different combinations of features (Fig. 7, bottom). An interesting difference between the manual and automatic segmentation is that the ROC curve for the spread attribute shows that classifiers using only appearance-based features performed better than classifiers using only shape-based features when automatic segmentation was applied.

3.2 Clutter identification

We used the first ten annotated images (50 min) from each of the 30 sequences in the time-lapse migration library to assemble a dataset containing 1,870 cell and 2,252 clutter objects. To extract objects from the image data, automated segmentation was applied, and the ground truth labels were aligned with the detected regions. Regions that did not correspond to any truth object were identified as clutter. Regions that were in one-to-one correspondence with truth objects were identified as cells. Shape-based and appearance-based features were computed for these objects. Using this dataset, we then trained an Adaboost classifier to distinguish between clutter and cells. Tenfold cross validation was used to evaluate the classifier performance using different feature sets (Fig. 8).

By examining the true positive rate, we chose a prediction threshold that favored correctly identifying cells at the expense of misidentifying clutter. Using this adjusted threshold, we were able to identify 90.5% of clutter objects in the training data set (2,038 out of 2,252), while preserving 97.9% of the cell detections (1,831 out of 1,870 cells), measured using tenfold cross validation (Table 2).

3.3 Classification in time-lapse migration library

We now present the impact of our clutter mitigation strategy and morphology classification technique on our analysis of sequences in the time-lapse migration library. These image

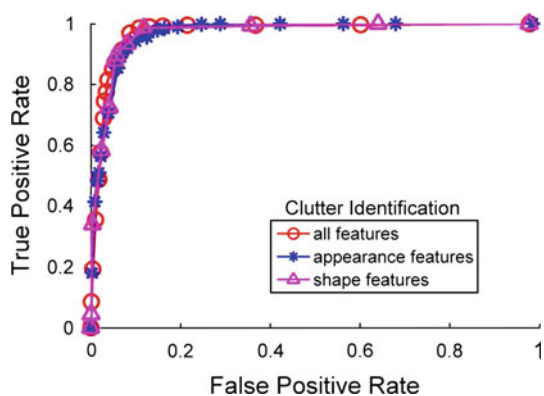


Fig. 8 ROC curves showing the performance of the Adaboost classifier distinguishing between cells and clutter when using shape-based features, appearance-based features, or both

Table 2 Confusion matrix showing the classification results for 4,122 clutter and cell objects, with a threshold adjusted to favor correctly identifying cells

Predicted	True	
	Cell	Clutter
Cell	1, 831	214
Clutter	39	2, 038

sequences are representative of the data used in biomaterial studies conducted to address important issues, e.g., cell–substrate interactions.

There are two sets of results to report. One result is the change in the number of correctly and incorrectly identified cells found with and without the clutter mitigation strategy. The other result is the morphology classification performance on the correctly identified cells.

In each sequence, cells were imaged at 30-s intervals, and cells in the images were manually identified and classified by a cell biologist every ten frames (5 min). No manual segmentation was performed. The first ten marked images (50 min) from each sequence were used to assemble two training datasets to train the clutter and cell morphology classifiers. The subsequent 500 frames (50 annotated images, approximately 4:10 h) from each sequence were used as testing data.

3.3.1 Clutter mitigation success

The tracker described in Sect. 2.6 was applied to 500 frames from each sequence in the time-lapse migration library. Although the expert-supplied annotations were only available for every tenth frame, the tracker had the benefit of access to every frame in the sequence. When the clutter mitigation strategy was used, clutter objects were removed from consideration before the tracker had access to the detected objects.

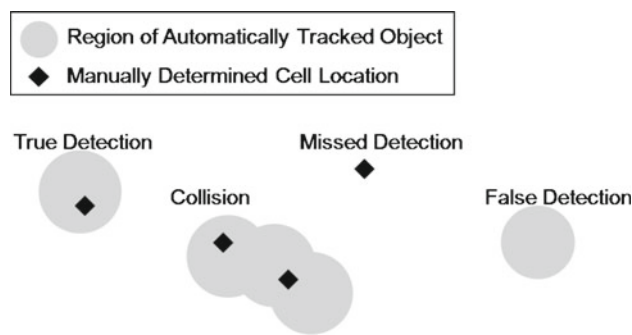


Fig. 9 An illustration of the definition of a true detection, missed detection, collision, and false detection

To evaluate the success of our clutter mitigation strategy, we computed several performance metrics to compare the objects found by our segmentation and tracking algorithms with the annotations made by the cell biologist. We computed the scores for our metrics using the tracked objects produced by the tracker, instead of the raw detections. (The tracker discarded detections contained in tracks shorter than 2.5 min.) A true detection is defined as an annotated cell in one-to-one correspondence with a tracked object. A tracked object may be a detection from the image, or a coast inserted by the tracker. A missed detection is defined as an annotated cell that does not correspond to any tracked object. A collision is defined as an annotated cell that corresponds to a tracked object that corresponds to more than one annotated cell (the set of annotated cells in a many-to-one relationship with tracked objects). A false detection is defined as a tracked object that does not correspond to an annotated cell (Fig. 9).

The first metric we computed was the probability of true detection—the number of true detections, over the total number of annotated cells. The second metric was the probability of a missed detection—the number of missed detections over the total number of annotated cells. The third metric was the probability of false detection—the number of false detections over the total number of tracked objects. The probability of collision—the number of truth objects in many-to-one correspondence with tracked objects over the total number of annotated cells—is not affected by our clutter mitigation strategy, but it does reduce our apparent probability of detection. The correspondence between detected objects and annotated cells was determined as described in the Sect. 3.1.3. We also counted the total number of true detections, missed detections, and false detections. Tabulation of performance metrics was performed every ten frames (5 min).

We observed our desired result, namely, significantly decreasing false detections and only slightly affecting our detection rate (Table 3). The true detection rate was reduced from 82.8 to 80.0%, but the false detection rate was reduced from 51.2 to 15.8%. The total number of false detections dropped by 82.8%, while the total number of true detections only dropped by 3.4%.

Table 3 Using a machine learning approach to reducing clutter causes a significant reduction in false detections, with only a small reduction in true detections

	No clutter mitigation	With clutter mitigation
Probability of true detection	82.8%	80.0%
Probability of missed detection	0.2%	3.1%
Probability of collision	16.9%	16.9%
Probability of false detection	51.2%	15.8%
Total annotated cells	11,469	11,469
Total tracked objects	20,385	11,791
Total true detections	9,502	9,180
Total missed detections	25	351
Total false detections	9,983	1,718

Table 4 Morphology classification accuracy achieved in the time-lapse migration library

Overall accuracy	55.7%
Weighted accuracy	82.1%
Spread accuracy	88.8%
Polarized accuracy	74.0%
Oriented accuracy	83.4%




Fig. 10 Overall confusion matrix obtained for the classification performance on 9,180 true detections found in the time-lapse migration library, using automatic segmentation, a simple tracker, and both shape and appearance-based features

3.3.2 Morphology classification performance

Our final task was to analyze the performance of the morphology classifier on the true detections. Using the clutter mitigation strategy, there were 9,180 true detections. The true class distribution of the cells was 82.1% spread, 44.9% polarized, and 14.4% oriented.

Measured by the overall accuracy, we found that our classification performance was as good as with the performance achieved on cells in the morphology library (Table 4; Fig. 10).

We observed in the structure of the confusion matrix that the confusion rate between polarized/non-oriented and polarized/oriented cells is worse than the confusion rate in the morphology library, although when measured by classification accuracy for the oriented attribute, our performance is quite good. We also note that the accuracy of our method for the polarized attribute is lower than the accuracy observed in the morphology library. Both of these issues may be due to the class prevalence being different in the time-lapse migration library than in the morphology library.

4 Discussion

4.1 Relationship to previous work

Machine learning methods have been applied to analysis of microscopy images for many years. We do not compare our methods to previously published work because of the difficulty of direct comparison with methods that are designed for different situations, including different imaging modalities (e.g., fluorescence microscopy [2, 12, 17, 18]), microorganisms [23], or cell types. To the best of our knowledge, public benchmark datasets of phase-contrast microscopy images for evaluating image analysis techniques are not available and need to be created.

We are not aware of any work where the issue of clutter mitigation has been addressed using machine learning, or where classification accuracy using automated segmentation is compared with classification accuracy using manual segmentation. Experimenting with different machine learning techniques and feature sets proposed in the literature remains an important part of ongoing work.

4.2 Feature selection

The set of features to use as input to a machine learning algorithm is one of the most important decisions when designing a classification methodology. In our methodology, we used

Adaboost to choose among weak classifiers that each operated on a single element of the feature vector, and so we can use the classifiers built by Adaboost to gain some insight into the discriminative power of the features, based on the way each weak classifier was chosen and weighted. Care must be taken when interpreting this information. Adaboost selects weak classifiers which are in some sense orthogonal to each other. So, in our case, if a set of features is strongly correlated, only one feature may appear in the finished classifier, but this does not mean that the other features are not useful. Another consideration is that the number of weak classifiers to use is determined *a priori*, and so classifiers chosen towards the end of training may not contribute strongly to the overall result, and can actually degrade performance in the worst case. The strongest evidence for the usefulness of an attribute is that it is consistently chosen by the algorithm under differing conditions.

We examined the classifiers built for classifying each attribute under the three conditions tested; manual and automatic segmentation in the morphology library and automatic segmentation in the migration library. It is interesting that for a given attribute, the classifiers built under different conditions generally used different features.

In our experiments, we found that shape-based features made the strongest positive contribution to the classification. The shape-based features that appear frequently include measures of cell size, including area, perimeter, and the mean distance between the centroid and pixels belonging to the boundary. We also see features related to the cell's roundness, including circularity and the isoperimetric quotient. The lower Hu moments and the extension, dispersion, and elongation image moments, which are functions of the first and second Hu moments, appear especially in the classifiers for the oriented attribute.

Rodenacker and Bengtsson [22] discussed a wide variety of features and reported that statistics of the gradient magnitude and Laplacian were very useful. The appearance-based features that appear in our finished classifiers are related to texture, including the mean of the absolute value of the Laplacian and the standard deviation and skew of the gradient magnitude.

4.3 Additional experimental methodologies

Guided by domain knowledge about cell morphology, we experimented with two additional approaches for morphology classification. Our methods and quantitative results are described in Sect. 6.

Our first alternative approach was to treat the attributes as an ordered cascade by enforcing a classification hierarchy using rules (Sect. 6.1). To implement this approach, we modified the training and testing phases of Adaboost. We found that the hierarchical rules imposed a dependency structure that, as the experiments revealed, was not needed. Inspection of the structure of the confusion matrix computed without hierarchical rules (Fig. 6) showed that the prediction structure that the hierarchical rules were designed to impose already existed. For example, one rule in the system stated that non-spread cells cannot be polarized. In the experiments without hierarchical rules, no cells were classified as both non-spread and oriented, so this rule was shown to be redundant.

Our second alternative approach was to use ternary labels, instead of binary labels. Our decision to experiment with ternary labels was informed by the biological reality that, since the shape of a cell changes continuously, a cell may pass through a partially spread (or partially polarized or oriented) state in between the non-spread and fully-spread states. Using ternary labels, we introduced an intermediate or partial class between the positive and negative classes for each attribute. This gives us slightly finer granularity for classifying cell shape, which may be important in downstream studies of population dynamics.

To evaluate our classification results using ternary labels, we compared the automatic classification results with ternary labels provided by one of the cell biologists. We computed a modified accuracy metric that gave partial credit for partially correct results. Using this evaluation technique, we did not find that our quantitative performance was substantially different from the experiments where we used binary labels.

On the other hand, qualitatively, we found some differences. We define a positive or negative example classified as belonging to the intermediate class to be partially correct. Similarly, an intermediate example classified as positive or negative is also defined to be to be partially correct. Positive examples classified as negative and negative examples classified as positive are defined to be fully incorrect. We found that while there were fewer fully correct results, there were far fewer fully incorrect results because many of the cells that were not identified correctly were assigned to the intermediate class. Depending on the penalty for incorrect classification, and the acceptability of intermediate labels, ternary labels may be beneficial. Applications may also benefit from the finer classification granularity available when using ternary labels.

5 Conclusions

In this paper, we addressed two problems: classifying cell shape and identifying clutter. Our goal was to identify and classify cells in a completely automated way. To benchmark our performance, we used the disagreement between human experts. To understand how well our classification technique would work under the best circumstances, we applied our methodology to manually segmented images. We found that our classification accuracy was comparable to the agreement between human observers. Then, we applied our classification technique to automatically segmented images, and we found a small reduction in classification performance.

We emphasize that even when given the imprecise cell boundaries produced by a simple automated segmentation method, our classification method was successful. This is an important result because it reveals that image analysis tasks can be performed effectively without perfecting automated segmentation.

Misinterpreting non-cell clutter objects can waste significant resources, both computational resources and, more importantly, the time and efforts of the cell biologist who may have to carefully inspect the results for correctness. Without this inspection, false detections jeopardize the quality of the results of subsequent analysis. This lack of trust may stifle research in the development of downstream algorithms.

Previous work on cell segmentation and tracking has not sufficiently addressed the important issue of mitigating clutter. Our contribution, a machine learning approach that identified clutter objects successfully, is a significant step towards addressing this issue. In one of our experiments, our method correctly identified over 90% of clutter objects while preserving nearly 98% of the cell objects (Sect. 3.3). Future work is needed to improve clutter identification while simultaneously correctly identifying cells.

Acknowledgments We would like to thank Matthew Antone and Evimaria Terzi for helpful discussions and Patricia A. Solski for providing us with ground-truth cell labels. Funding for this work was provided by the National Science Foundation, HCC grant IIS-0910908.

6 Appendix: Additional experimental methodologies

Guided by domain knowledge about cell morphology, we experimented with two additional approaches for morphology classification. For these two experiments, we only report our results using manual segmentation.

6.1 Classification results using hierarchical rules

Our first alternative approach was to treat the attributes as an ordered cascade by enforcing a classification hierarchy. To

Table 5 Performance of two alternate classification strategies are compared with our primary classification method

Classification method	Binary	Hierarchical	Ternary
Cell instances	883	883	883
Overall accuracy (%)	65.3	63.1	45.8
Weighted accuracy (%)	86.7	85.9	90.8
Spread attribute (%)	93.9	93.5	93.0
Polarized attribute (%)	85.6	85.7	85.9
Oriented attribute (%)	80.6	78.4	75.3

Results were obtained using tenfold cross validation on instances found using automated segmentation

implement this approach, we modified the training and testing of the Adaboost classifier. During the training phase, we removed all non-spread instances from the training set for the polarized classifier. Similarly, we removed all non-spread and non-polarized instances from the training set for the oriented classifier. During the prediction phase, any cell classified as non-spread was automatically classified as non-polarized and non-oriented.

Using manual segmentation, we found that the performance obtained when using hierarchical rules was similar to the performance obtained without them (Table 5, column 2) and that the confusion matrix (not shown) had the same structure.

6.2 Classification results using ternary labels

Our second alternative methodology uses ternary labels, instead of binary labels. We discuss the justification for this in Sect. 4.3. To evaluate our approach, we compare our classification results to ternary labels applied to all of the cells in the morphology library by one of the two cell biologists.

For each attribute, we trained two one-versus-all classifiers: one classifier for the positive class and one classifier for the negative class. In the prediction phase, cells were classified as positive or negative only if the score computed by the appropriate classifier was positive and the score computed by the other classifier had the opposite sign. If the scores from the two classifiers had the same sign, then the example was assigned the intermediate label. After a ternary value was assigned for each attribute of an instance of a cell, our method combined the labels to form the final classification of the shape of the cell.

To evaluate the ternary classifier, we compute 3×3 confusion matrices for each attribute. To measure the accuracy, we compute a cost function between the true class t and the predicted class c that gives credit for a partially correct answer. The labels t and c have a value of 1 if they are the positive class, -1 if they are the negative class, and 0 if they are the

Table 6 Binary classification results for each attribute, using manual segmentation, using both shape-based and appearance-based features

Predicted	True	
	Non-spread (%)	Spread (%)
Non-spread	15.5	2.7
Spread	3.4	78.4
	Non-polarized (%)	Polarized (%)
	Non-polarized	32.4
Polarized	8.0	53.2
	Non-oriented (%)	Oriented (%)
	Non-oriented	59.5
Oriented	8.9	21.2

Results were obtained using tenfold cross validation on 883 instances

intermediate class. We define

$$f(t, c) = \begin{cases} 1 & \text{if } t = c, \\ 0.375 & \text{if } t = 0 \text{ and } c = \{1, -1\}, \\ 0.375 & \text{if } t = \{1, -1\} \text{ and } c = 0, \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

The weighted accuracy metric, defined for 3×3 confusion matrices using ternary labels, for a particular attribute a is

$$W_3 = \frac{1}{N} \sum_{i=1}^n f(T_i(a), C_i(a)). \quad (7)$$

Table 7 Ternary classification results for each attribute, using manual segmentation, using both shape-based and appearance-based features

Predicted	True		
	Non-spread (%)	Partially-spread (%)	Spread (%)
Non-spread	8.5	2.8	0.0
Partial-spread	1.4	11.1	3.1
Spread	0.2	3.5	69.4
	Non-polarized (%)	Partially-polarized (%)	Polarized (%)
	Non-polarized	17.7	3.1
Partial-polarized	2.4	30.2	9.7
Polarized	0.1	7.2	29.6
	Non-oriented (%)	Partially-oriented (%)	Oriented (%)
	Non-oriented	52.0	7.5
Partial-oriented	8.5	7.5	8.4
Oriented	1.8	2.8	5.7

Results were obtained using tenfold cross validation on 883 instances

If no examples exist with true or predicted labels belonging to the intermediate class, this metric is equivalent to “accuracy” as defined above for the binary case. The expected value of the metric is 0.5 if the truth labels and predicted labels are equally distributed.

For the full 27×27 confusion matrix, we compute the usual definition of accuracy as the sum of the diagonal of the confusion matrix. We also compute a weighted accuracy metric which is

$$W_{27} = \frac{1}{N \cdot L} \sum_{i=1}^N \sum_{l=1}^L f(T_i(l), C_i(l)). \quad (8)$$


When inspecting the scores computed with the A_3 and W_{27} metrics, we observed that the scores obtained using ternary labels were comparable to the scores obtained using binary labels, although the weighted accuracy metric is better (Table 5, column 3).

When we examined the confusion matrices, we found that there were fewer fully correct results, but far fewer fully incorrect results. To illustrate the differences in the nature of the classification errors, we present the 2×2 confusion matrices for each attribute for the binary case (Table 6) and the 3×3 confusion matrices for each attribute for the ternary case (Table 7). For example, for the polarized attribute, using binary labels, 14.3% of the examples were classified incorrectly, whereas using ternary labels, 0.1% were classified incorrectly.

References

- Al-Kofahi, O., Radke, R.J., Goderie, S.K., Shen, Q., Temple, S., Roysam, B.: Automated cell lineage construction: a rapid method to analyze clonal development established with murine neural progenitor cells. *Cell Cycle* **5**, 327–335 (2006)
- Boland, M., Markey, M., Murphy, R.: Classification of protein localization patterns obtained via fluorescence light microscopy. In: Proceedings of the 19th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, vol. 2, pp. 594–597 (1997)
- Bradhurst, C., Boles, W., Xiao, Y.: Segmentation of bone marrow stromal cells in phase contrast microscopy images. In: Proceedings of the 23rd International Conference Image and Vision Computing (IVCNZ), New Zealand, pp. 1–6 (2008)
- Dunn, G.A., Brown, A.F.: Alignment of fibroblasts on grooved surfaces described by a simple geometric transformation. *J. Cell Sci.* **83**, 313–340 (1986)
- Dzyubachyk, O., van Cappellen, W.A., Essers, J., Niessen, W., Meijering, E.: Advanced level-set based cell tracking in time-lapse fluorescence microscopy. *IEEE Trans. Med. Imaging* **29**(3), 852–867 (2010)
- Freund, Y., Schapire, R.E.: A decision-theoretic generalization of on-line learning and an application to boosting. In: Proceedings of the Second European Conference on Computational Learning Theory (EuroCOLT '95), London, UK, pp. 23–37. Springer-Verlag, Berlin (1995)
- Halabi, Y.S., SA, Z., Hamdan, F., Yousef, K.H.: Modeling adaptive degraded document image binarization and optical character system. *Eur. J. Sci. Res.* **28**(1), 14–32 (2009)
- Han, J.W., Breckon, T.P., Randell, D.A., Landini, G.: The application of support vector machine classification to detect cell nuclei. *Mach. Vis. Appl.* (2010) (Published online)
- Horn, B.K.P.: *Robot Vision*. The MIT Press, Cambridge (1986)
- House, D., Walker, M.L., Wu, Z., Wong, J.Y., Betke, M.: Tracking of cell populations to understand their spatio-temporal behavior in response to physical stimuli. In: Proceedings of the IEEE Computer Society Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA), Miami, FL pp. 186–193 (2009)
- Hu, M.-K.: Visual pattern recognition by moment invariants. *IRE Trans. Inf. Theory* **8**(2), 179–187 (1962)
- Huang, K., Murphy, R.: Boosting accuracy of automated classification of fluorescence microscope images for location proteomics. *BMC Bioinform.* **5**(1), 78 (2004)
- Jain, R., Kasturi, R., Schunk, B.: *Machine Vision*. McGraw Hill, (1995)
- Kachouie, N.N., Fieguth, P.W.: Extended-Hungarian-JPDA: Exact single-frame stem cell tracking. *IEEE Trans. Biomed. Eng.* **54**(11), 2011–2019 (2007)
- Li, K., Miller, E.D., Chen, M., Kanade, T., Weiss, L.E., Campbell, P.G.: Cell population tracking and lineage construction with spatiotemporal context. *Med. Image Anal.* **12**(5), 546–566 (2008)
- Mukherjee, D., Ray, N., Acton, S.: Level set analysis for leukocyte detection and tracking. *IEEE Trans. Image Process.* **13**(4), 562–572 (2004)
- Murphy, R.F.: Communicating subcellular distributions. *Cytometry Part A* **77A**(7), 686–692 (2010)
- Nanni, L., Lumini, A., Lin, Y.-S., Hsu, C.-N., Lin, C.-C.: Fusion of systems for automated cell phenotype image classification. *Expert Syst. Appl.* **37**(2), 1556–1562 (2010)
- Nath, S., Palaniappan, K., Bunyak F.: Cell segmentation using coupled level sets and graph-vertex coloring. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) *Medical Image Computing and Computer-Assisted Intervention MICCAI 2006*. Lecture Notes in Computer Science, vol. 4190, pp. 101–108. Springer, Berlin (2006)
- Olson, A.C., Larson, N.M., Heckman, C.A.: Classification of cultured mammalian cells by shape analysis and pattern recognition. *Proc. Nat. Acad. Sci.* **77**(3), 1516–1520 (1980)
- Ray, N., Acton, S.: Motion gradient vector flow: an external force for tracking rolling leukocytes with shape and size constrained active contours. *IEEE Trans. Med. Imaging* **23**(12), 1466–1478 (2004)
- Rodenacker, K., Bengtsson, E.: A feature set for cytometry on digitized microscopic images. *Anal. Cell. Pathol.* **25**, 1–36 (2003)
- Ruusuvuori, P., Seppala, J., Erkkila, T., Lehmussola, A., Puhakka, J., Yli-Harja, O.: Efficient automated method for image-based classification of microbial cells. In: Proceedings of the 19th International Conference on Pattern Recognition 2008 (ICPR 2008), pp. 1–4 (2008)
- Viola, P., Jones, M.J.: Robust real-time face detection. *Int. J. Comput. Vis.* **57**, 137–154 (2004)
- Wong, J.Y., Leach, J.B., Brown, X.Q.: Balance of chemistry, topography, and mechanics at the cell-biomaterial interface: Issues and challenges for assessing the role of substrate mechanics on cell response. *Surf. Sci.* **570**, 119–133 (2004)
- Yin, Z., Bise, R., Chen, M., Kanade T.: Cell segmentation in microscopy imagery using a bag of local Bayesian classifiers. In: *From Nano to Macro: The 2010 IEEE International Symposium on Biomedical Imaging*, pp. 125–128 (2010)

Author Biographies



Diane H. Theriault received a B.A. degree in Computer Science from Boston University in 2004, followed by an M.S. degree in Media Arts and Sciences from the Massachusetts Institute of Technology in 2006. From 2006 through 2009, she was a software engineer at BAE Systems, Advanced Information Technologies, working on tracking and applied computer vision. She is currently working towards a Ph.D. degree in Computer Science at Boston University as a member of the Image and Video Computing Group.



Matthew L. Walker received a B.A. degree in Biology from Brandeis University in 2001 and an M.S. degree from the Molecular Biology, Cell Biology and Biochemistry program at Boston University in 2006. Currently, he is a Ph.D. candidate in the Biomedical Engineering Department at Boston University, where he is studying cellular migration and materials science.



Joyce Y. Wong is an Associate Professor of Biomedical Engineering (BME) and a College of Engineering Distinguished Faculty Fellow at Boston University. Dr. Wong's research focuses on the development of biomaterials to probe how structure, material properties and composition of the cell-biomaterial interface affect fundamental cellular processes. Her current research interests include tissue engineering of small diameter blood vessels for bypass surgery and intravascular pharmacology (e.g. stents); development of targeted

nano- and micro-particle contrast agents for multi-modal (magnetic resonance, ultrasound, and optical) detection of atherosclerotic and vulnerable plaque; and engineering biomimetic systems to study restenosis and cancer. Her research has been funded by NIH, NASA, DOE, and industry. Awards she has received include an NSF CAREER Award, Clare Boothe Luce Assistant Professorship, Dupont Young Professor Award, and most recently she was named a College of Engineering Distinguished Faculty Fellow (at BU), was elected to AIMBE, and received a Hartwell Individual Biomedical Research Award. She was also selected for participation in the National Academy of Science Frontiers in Engineering, National Academies Keck Futures Initiative Conference, German-American Frontiers in Polymer Science, and Japan-America Frontiers in Engineering. Dr. Wong has served as Associate Director of the Center for Nanoscience and Nanobiotechnology at Boston University. Dr. Wong is also active in promoting women in science and engineering, and has co-founded a Women in BME Bookclub for graduate students at BU.



Margrit Betke is an Associate Professor of Computer Science at Boston University, where she co-leads the Image and Video Computing Research Group, and was a Research Scientist at the Massachusetts General Hospital and Harvard Medical School. She conducts research in computer vision, in particular, the development of methods for detection, segmentation, registration, and tracking of objects in visible-light, infrared, and x-ray image data. She has worked on gesture, vehicle, and animal

tracking, video-based human-computer interfaces, statistical object recognition, and medical imaging analysis. She has published over 80 original research papers. She earned her Ph.D. degree in Computer Science and Electrical Engineering at the Massachusetts Institute of Technology in 1995. Prof. Betke has received the National Science Foundation Faculty Early Career Development Award in 2001 for developing "Video-based Interfaces for People with Severe Disabilities." She co-invented the "Camera Mouse," an assistive technology used worldwide by children and adults with severe motion impairments. She co-developed the first patented algorithms for detecting and measuring pulmonary nodule growth in computed tomography. She was one of two academic honorees of the "Top 10 Women to Watch in New England Award" by Mass High Tech in 2005. She currently leads a 5-year research program to develop intelligent tracking systems that reason about group behavior of people, bats, birds, and cells.