

Cell Detection and Segmentation using Correlation Clustering

Chong Zhang¹, Julian Yarkony², and Fred A. Hamprecht²

¹ CellNetworks, Heidelberg University, Germany

² HCI/IWR, Heidelberg University, Germany

Abstract. Cell detection and segmentation in microscopy images is important for quantitative high-throughput experiments. We present a learning-based method that is applicable to different modalities and cell types, in particular to cells that appear almost transparent in the images. We first train a classifier to detect (partial) cell boundaries. The resulting predictions are used to obtain superpixels and a weighted region adjacency graph. Here, edge weights can be either positive (attractive) or negative (repulsive). The graph partitioning problem is then solved using correlation clustering segmentation. One variant we newly propose here uses a length constraint that achieves state-of-art performance and improvements in some datasets. This constraint is approximated using non-planar correlation clustering. We demonstrate very good performance in various bright field and phase contrast microscopy experiments.

1 Introduction

With recent advances in microscope automation, long-term high-throughput imaging results in a vast amount of data in biological experiments. Consequently, the demand for computer aided microscopy image analysis is high. Cell detection and segmentation are fundamental tasks for further cell-level quantifications. The large diversity of cell lines and microscopy imaging techniques require the development of algorithms for these tasks to perform robustly and equally well in different scenarios. The technique described in this paper is applicable to images that have crowded cell regions acquired from different modalities and cell shapes, as long as they produce intensity changes at cell boundaries. Such patterns result from several microscopy imaging techniques, such as transillumination (e.g. bright field, dark field, phase contrast) and fluorescence (e.g. through membrane or cytoplasmic staining) images. Thus it is specifically suitable for images from which cells are almost transparent (Fig. 1).

Correlation clustering, or multicut, as an image segmentation method has attracted considerable interest in recent years [1,2,6,11]. It finds a partitioning of a weighted region adjacency graph into an arbitrary number of segments such that the set of edges that cohere different segments has minimum total weights. Finding a minimum weight partition is NP-hard for general and planar graphs [4]. Algorithms that can solve NP-hard correlation clustering problems to (near) optimality on instances of practically relevant size are only recently




Fig. 1: Example images of (*left-right*): Dataset-a, Bright field image of Diploid yeast cells (1K×1K pix; cell size 40-100 pix). Dataset-b, Bright field image of Fission cells [9] (1K×1K pix; cell size 30-140 pix). Dataset-c, Phase contrast image of cervical cancer cells of the HeLa cell line [3] (400×400 pix; cell size 10-40 pix)

developed [1,11], and their potential for biomedical applications has not yet been explored in depth. Adding cardinality or size constraints makes the problem even harder. Thus, approximate optimization strategies that can give an exact or close to exact solution in short turnaround times are desired. A mathematically sound way to incorporate long-range repulsive interactions while keeping the resulting model tractable has only been discovered very recently [2]. Our work builds on ideas from [2,11] and explores their applicability in cell detection and segmentation. In particular, we leverage this theoretical insight into practice in terms of a length constraint formulation, a sound cue for biological problems.

2 Method

Our method first computes a cell boundary probability map from a trained edge classifier, and then constructs superpixels and builds a weighted superpixel adjacency graph. Segmentation is then reflected from partitioning this graph using a correlation clustering procedure.

Extracting boundary evidence. In our case, we learn the boundary probability from a trained classifier using *ilastik* [10]. It is an open-source toolkit which relies on a family of generic image features and a robust nonlinear classifier, random forest, to estimate each pixel’s probability of belonging to cell boundary. This enables the flexibility of detecting edges from cells of interest (Fig. 2b). We only label very few pixels (<1%) to reduce annotation effort.

Computing superpixels. The obtained boundary probability map is then used to construct a set of initial regions as superpixels. While many computing strategies could be applied, we used the watershed transform. A slightly smoothed probability map helps avoiding tiny non-informative superpixels, while still keeping boundaries with low probability separated from background superpixels. In the test datasets, we relate the amount of smoothing with roughly average cell length l , i.e. a Gaussian filter with filter size $0.3l$ and standard deviation $0.1l$.




Fig. 2: (a) Example image. (b) Cell boundary pixel probabilities (the brighter the higher probability). The trained classifier gives low boundary probabilities to out-of-focus cells (*arrow*). (c) Superpixels (randomly color-mapped) with an overlay of the corresponding adjacency graph edges (*white*). (d) An illustration of the length constraint for a superpixel \mathbf{p} (*light gray*), where non-planar edges (*blue*) indicate the superpixels (*white*) to be separated from \mathbf{p} . They lie on the circle (*dashed-red*) of radius d centered at superpixel \mathbf{p} 's centroid.

Constructing a weighted region adjacency graph. A superpixel adjacency graph $G=(V, E)$ is built (Fig. 2c). Each graph edge e corresponds to an adjacent superpixel pair and is associated with a real valued potential θ_e . Negative potentials forces superpixels to be in separate regions, while positive ones favor to merge them. We first approximate the graph edge probability p_e by the average boundary probability over the pixels separating the corresponding adjacent superpixel pair. Then, the graph edge potential is given by

$$\theta_e = \log \left(\frac{1 - F(p_e + t)}{F(p_e + t)} \right),$$

where $F(x) = \max(\min(x, 0.9999), 0.0001)$, and t is a bias term that adjusts the transform from (non-negative) probabilities to negative/positive potentials. A small bias favors a single region of all superpixels, whereas a large one leads to regions of single superpixels.

Solving the graph partitioning problem. A partition is defined by a binary indicator variable X_e for each edge e . $X_e=1$ if edge e is to be cut and $X_e=0$ otherwise. A valid configuration [11] of labeling $X \in \{0, 1\}^E$ of edges gives a possible

segmentation. The correlation clustering problem can be expressed as:

$$\min_{X \in C} \theta_e X_e,$$

where C is the set of all possible segmentations. If edges are only between adjacent superpixels and have unrestricted potential values, this results in a planar graph. Given efficient search strategies it can be solved to optimality using an integer linear programming solver using a cutting planes approach [1]; or approximated using dedicated solvers such as PlanarCC [11], which often reaches the global optimum as well (see [5] for a benchmarking). We consider the latter for the cell segmentation problem, which belongs to instances of large sizes (Table 1). Also, given its closed contour property, the application of such models is valuable because cell boundary information is often only partially available and can be inconsistent.

Adding a length constraint. Restricting the size of segments is valid for cell segmentation problems as a specific cell type has a known size prior. This is particularly helpful in situations when cells are largely clustered together and their separation boundaries are missing. Thus, a constraint that no cell has a length (in its major axis) larger than a hyper-parameter d can be imposed. We formulate this to a graph structure that is amenable to a highly efficient approximate (and as it happens often exact) optimization: Semi-PlanarCC [2]. In this model, not only edges between adjacent superpixels have potentials, but also edges between distant non-adjacent superpixels have large negative potentials. To do so, it is sufficient to state that for any two superpixels, if separated by an approximated distance of d , then they must be in separate regions. For every superpixel \mathbf{p} , a circle of radius equal to d is drawn at the centroid of \mathbf{p} . All superpixels lying on that circle may not be in the same region as superpixel \mathbf{p} . An illustration is shown in Fig. 2d. The potentials between these superpixel pairs are set to a large negative value, e.g. $-\sum_e \theta_e - 0.001$. This means, it is preferable to cut all edges with positive potentials rather than missing cutting on edges with negative long-range potentials.

Finally, after removing background segments, each connected component enclosed by the cut edges is considered as an individual cell region. In our case, we currently use two heuristic size filtering steps to remove background: at segments level, consider those larger than $6l^2$ as background, and at superpixel level, consider segments containing superpixel(s) larger than $3l^2$ as background.

3 Experiments and Results

Our aim is to optimize both the detection and segmentation accuracy w.r.t. the ground truth (GT) in each image provided in the form of centroids for every cell and regions of a random subset of cells. Therefore, we evaluate the output of detection based on: A region is considered a true positive (TP) detection if

Table 1: Summary of experimental parameters.

Dataset	avg. cell length l	# superpixels	# planar edges	# non-planar edges
a	60	1225±242	3456±701	18025±6791
b [9]	100	3727±2450	10530±7010	147420±160420
c [3]	30	1081±364	3035±1038	26618±12683




Fig. 3: Precision-Recall plots of detection for the three datasets using various t values. (Results for Dataset-c using CellDetect is reproduced from [3]. Also for CellDetect, detection criteria was the same as in [3] due to the available output.)

the GT centroid is within this region; Regions that do not cover cell centroids are considered false positives (FP); Missed GT centroids are counted as false negatives (FN). The results are reported in terms of precision $P=TP/(TP+FP)$, recall $R=TP/(TP+FN)$, and F-measure $F=2 \times P \times R / (P+R)$. Due to the high cell density, we randomly selected 10–20% of cells from each image and manually segmented them to provide the GT. The segmentation accuracy is then computed from two area overlap measures between the TP detection regions R_{tpd} and the GT region R_{gt} : $M_1=(R_{tpd} \cap R_{gt}) / (R_{tpd} \cup R_{gt}) \times 100$, $M_2=(R_{tpd} \cap R_{gt}) / R_{gt} \times 100$.

Three datasets have been used to validate our method (Fig. 1). Dataset-a (15 images with 1768 in-focus cells in total) contains bright field Diploid yeast cells. It is required to detect in-focus cells while discriminating them from out-of-focus cells. The detection task is challenging since the cell boundary edges can be partially missing and often exhibit varying contrast patterns even in the same cell. Other challenging issues include poor contrast, partial or changing halo, overlap with out-of-focus cells, and imaging artifacts. Dataset-b (10 images with 2340 cells in total) is from bright field Fission yeast cells in [9], where cells are elongated cylinder like and can be very densely packed as well as in the Dataset-a but having different appearance. Dataset-c (11 images with 1073 cells in total) is from CellDetect [3], which contains phase contrast images of cervical cancer cell colonies of the HeLa cell line. It presents a high variability in cell shapes and sizes. Table 1 summarizes the parameter details.

Detection results are shown in Fig. 3. The detection differences between PlanarCC and Semi-PlanarCC lie on two sides: Semi-PlanarCC splits either large

Table 2: Results of segmentation in area overlap measures M_1 and M_2 on the TP detected cells (mean \pm standard deviation). The highest number (100 indicating the best) among the methods is indicated as bold.

Dataset	Meas.	PlanarCC	Semi-PlanarCC	CellX [8]	CellStat [7]
a	M_1	86.3 \pm 13.3	86.4 \pm 12.0	86.9 \pm 16.0	89.5\pm13.6
	M_2	97.0\pm1.6	96.8 \pm 3.1	94.8 \pm 14.2	91.4 \pm 13.9
b [9]	M_1	74.4\pm21.8	74.1 \pm 21.3	32.8 \pm 23.8	49.3 \pm 24.2
	M_2	86.9 \pm 23.8	90.2\pm21.7	32.8 \pm 23.8	55.7 \pm 29.6
c [3]	M_1	71.4\pm12.1	70.1 \pm 11.9	42.1 \pm 12.7	46.4 \pm 8.6
	M_2	93.9 \pm 13.3	94.4\pm11.4	42.2 \pm 12.7	60.6 \pm 14.0

groups of cells or cells from background that would be merged by PlanarCC; It might also split exceptionally large cells or background into multiple segments small enough to be rejected. This could be seen in Fig. 3 that Semi-PlanarCC has achieved higher recall values than PlanarCC for the same t value (i.e. seemingly compressed curve horizontally shifted rightwards). This may suggest that the former is less sensitive to the choice of t . And the lower precision values for Dataset-a are due to the multiple fragmented background segments whose sizes are similar to a cell. This is more likely to happen in microscopy images of cell cultures of medium confluence. This could be alleviated through a post-processing step using e.g. shape or texture features. Table 2 shows the area overlap measures on segmentations from a random subset of TP detected cells, both methods perform similarly in their best F measure cases (Fig. 3).

In Fig. 3 and Table 2, we also show detection and segmentation results of three methods: CellDetect [3], CellX [8], CellStat [7]. PombeX [9] has also been evaluated but the tool does not provide quantifiable output thus not reported. For each method we have tried to optimize their parameters to achieve the best possible results. Reported results for Dataset-c using CellDetect is reproduced from [3]. Also for CellDetect, detection criteria was the same as in [3], since its detection is represented as cell centroids. The overall comparison is favorable to our method over other methods. CellStat is primarily designed for bright field images of round cells, and it looks for contours with consistent profile pattern. This may in part account for their poorer performance for elongated cells and phase contrast images. Similarly, CellX also tries to match boundary profile pattern on cells without extreme shapes. Due to the image appearance variations within each dataset, we expect that these methods would perform better if parameters are tuned for each image individually rather than what we do here: the same set of parameters for the whole dataset. A visual inspection of our segmentation can be seen in Fig. 4 as red contours overlaid on example images. It clearly shows that for Dataset-a, the out-of-focus cells (highly-contrasted and blurred ones) are clustered with the background, i.e. they are excluded from the segmentation.




Fig. 4: (*top*) Example original image regions from Datasets a-c (*left to right*). Overlaid lines (*red*) on the original images are: superpixel lines (*middle top*), segmentation contours of these regions (*middle bottom*) and the entire example image (*bottom*), with the regions indicated as white frames.

4 Discussion and Conclusions

The technique for cell detection and segmentation presented here is able to achieve state-of-the-art performance across different scenarios. In general PlanarCC and Semi-PlanarCC perform similarly, but for overlapping cells with missing cell boundaries, the latter is more beneficial. The model does not have any

specifications about closed region shape and size, which allows for various cell types applications. It has demonstrated that it can also handle large number of superpixels and graphs, which is the case of microscopic cell images. Our method only requires a few sparse labels to train a cell boundary classifier for all images in an experiment. Apart from setting cell length (and bias), further steps are automatic and require no user interaction, thus well suited for high-throughput studies. In addition, the flexibility of training edges from cells of interest allows us to exclude edges from other structures. Deterministic post-processing steps should be applied to further reject segments with implausible characteristics of being a cell, according to e.g. image texture features, or cell morphology features. Here we only focus on evaluating the performance of our technique.

Acknowledgement

We thank F. Huber and M. Knop from ZMBH University of Heidelberg, Germany for sharing Dataset-a.

References

1. B. Andres, J. H. Kappes, T. Beier, U. Köthe, and F. A. Hamprecht. Probabilistic Image Segmentation with Closedness Constraints. In *ICCV*, 2011.
2. B. Andres, J. Yarkony, B. S. Manjunath, S. Kirchho, E. Turetken, C. C. Fowlkes, and H. Pfister. Segmenting Planar Superpixel Adjacency Graphs w.r.t. Non-planar Superpixel Affinity Graphs. In *EMMCVPR*, 2013.
3. C. Arteta, V. Lempitsky, J. A. Noble, and A. Zisserman. Learning to Detect Cells Using Non-overlapping Extremal Regions. In N. Ayache, H. Delingette, P. Golland, and K. Mori, editors, *MICCAI*, volume 7510, pages 348–356. Springer, Heidelberg, 2012.
4. Y. Bachrach, P. Kohli, V. Kolmogorov, and M. Zadimoghaddam. Optimal Coalition Structure Generation in Cooperative Graph Games. In *AAAI*, 2013.
5. J. H. Kappes, B. Andres, F. A. Hamprecht, C. Schnörr, S. Nowozin, D. Batra, S. Kim, B. X. Kausler, J. Lellmann, N. Komodakis, and C. Rother. A Comparative Study of Modern Inference Techniques for Discrete Energy Minimization Problems. In *CVPR*, 2013.
6. S. Kim, S. Nowozin, P. Kohli, and C. D. Yoo. Higher-Order Correlation Clustering for Image Segmentation. In *NIPS*, 2011.
7. M. Kvarnstrom, K. Logg, A. Diez, K. Bodvard, and M. Kall. Image Analysis Algorithms for Cell Contour Recognition in Budding Yeast. *Opt. Express*, 16(17):1035–1042, 2008.
8. C. Mayer, S. Dimopoulos, F. Rudolf, and J. Stelling. Using CellX to Quantify Intracellular Events. *Curr. Protoc. Mol. Biol.*, Chapter 14:Unit 14.22., 2013.
9. J. Y. Peng, Y. J. Chen, M. D. Green, S. A. Sabatino, S. L. Forsburg, and C. N. Hsu. PombeX: Robust Cell Segmentation for Fission Yeast Transillumination Images. *PLoS One*, 8(12):e81434, 2013.
10. C. Sommer, C. Straehle, U. Koethe, and F. A. Hamprecht. ilastik: Interactive Learning and Segmentation Toolkit. In *ISBI*, 2011.
11. J. Yarkony, A. Ihler, and C. C. Fowlkes. Fast Planar Correlation Clustering for Image Segmentation. In *ECCV*, 2012.