Segmentation with Residual Attention U-Net and an Edge-Enhancement Approach Preserves Cell Shape Features

Nanyan Zhu*1, Chen Liu*2, Britney Forsyth³, Zakary S. Singer³, Andrew F. Laine³, Tal Danino³, and Jia Guo⁴

¹Department of Biological Sciences, Columbia University ²Department of Electrical Engineering, Columbia University ³Department of Biomedical Engineering, Columbia University ⁴Department of Psychiatry, Mortimer B. Zuckerman Mind Brain Behavior Institute, Columbia University

Abstract-The ability to extrapolate gene expression dynamics in living single cells requires robust cell segmentation, and one of the challenges is the amorphous or irregularly shaped cell boundaries. To address this issue, we modified the U-Net architecture to segment cells in fluorescence widefield microscopy images and quantitatively evaluated its performance. We also proposed a novel loss function approach that emphasizes the segmentation accuracy on cell boundaries and encourages shape feature preservation. With a 97% sensitivity, 93% specificity, 91% Jaccard similarity, and 95% Dice coefficient, our proposed method called Residual Attention U-Net with edge-enhancement surpassed the state-of-the-art U-Net in segmentation performance as evaluated by the traditional metrics. More remarkably, the same proposed candidate also performed the best in terms of the preservation of valuable shape features, namely area, eccentricity, major axis length, solidity and orientation. These improvements on shape feature preservation can serve as useful assets for downstream cell tracking and quantification of changes in cell statistics or features over time.

I. INTRODUCTION

Observations of cell dynamics at the population level is inherently sub-optimal due to the heterogeneity among individual cells. In contrast, imaging approaches enable single cell analysis, but require robust delineation of cell boundaries (aka., cell segmentation) and, for temporal information, matching cell identities over time (aka., cell tracking). Simpler cases with desirable properties such as high foreground-background contrast, high signal-to-noise ratios, and similar cell appearances, can be well handled by automated algorithms. On the other hand, manual annotation by experts is necessary whenever complications arise. Since the manual task of labeling and tracking is both tedious and may be inefficient and inconsistent when annotated by different individuals, a robust and automated alternative is highly desired.

The most promising work comes from the field of computer vision, where the topics of instance segmentation (i.e. treating multiple objects of the same category as separable individuals, as opposed to semantic segmentation) and multiobject tracking have been heavily visited. As summarized

*These authors contributed equally.

in Ciaparrone et al [1], segmentation and tracking tasks are typically addressed with four sequential steps: object detection, feature extraction, similarity calculation and identity association. While a lot of studies has been place on each of these four steps, the necessity of preserving cell shape features is underrepresented in the literature. Additionally, traditional evaluation metrics to quantify segmentation performance, such as accuracy, specificity, and Jaccard similarity, do not directly reflect how well a cell boundary is delineated. For example, if a star-shaped cell whose long processes constitutes 3% of its total area is approximated with a simple elliptical segmentation, the scores given by the aforementioned metrics can go beyond 95%, even though essential shape features are completely discarded.

Morphological details carry important information about the cell. For instance, the various unusual shapes of red blood cells reflect different developmental stages and/or pathologies [2]. Efforts spent on preserving cell shape features not only benefit the performance of various tracking algorithms that primarily rely on shape features such as Dewan et al [3], but also enable more accurate quantification of cell characteristics since the segmentation result is more reliable.

II. RELATED WORK

Computational techniques adopted for cell segmentation range from thresholding [4] on the simpler end along the spectrum of methods to deep learning and graph theory [5] toward the more sophisticated end. The method should be selected on a case-specific basis, i.e., a straightforward and effective algorithm for one case may completely fail for a second, while a rigorous and robust algorithm for the second may be overly complicated for another, etc. For instance, Al-Hafiz et al [4] could achieve an 87.9% accuracy rate by merely applying a dynamic thresholding algorithm on the image gradient and refining with morphological operations because their red blood cells are mostly homogeneous in intensity, regularly shaped, and clearly distinguishable from the background.

A few benchmarks for cell segmentation, or instance segmentation in general, include the U-Net family [6] and the R-CNN family of neural networks. The former is known for wide application in biomedical image segmentation while the latter excels at proposing likely objects and evaluating such detection proposals.

Datasets complicated by bias fields, visual artifacts, ambiguous cell boundaries, irregular cell shapes, overlapping or touching cells and other factors collectively render the problem significantly more challenging. One of the best performing algorithms in a similarly complicated dataset is introduced by Falk's team [6]Therefore we implemented the same deep learning architectures as they did and made performance-boosting adjustments on a variety of aspects from architecture and loss function to image processing techniques.

III. METHODS

The data we used was acquired at Columbia University in the department of Biomedical Engineering, where images were collected on the same baby hamster kidney (BHK) cell cultures every 20 minutes over roughly 15 hours. These images were acquired using fluorescence widefield microscopy images of nuclear and cytoplasmic live cell stains, differential interference contrast (DIC), and fluorescent proteins from cytoplasmically localized reporters. How the data is subsequently processed is described in the upcoming subsections as well as illustrated in Figure 1.



Fig. 1. Demonstration of the segmentation and tracking pipeline. a.&f. Bias field correction and intensity normalization. b.&g. Channelcombination to form pseudo-RGB images. c. Manual annotation. d.&h. Provide as input. e. Use the manual annotation as the ground truth to tune the deep learning network. Only provided during training and validation but not during testing. i. Post-processing. j. Future tracking algorithms.

A. Pre-processing and Data Augmentation

The nuclear and cytoplasmic stains, along with the DIC images, are treated as three channels of information for the localization of cells. Additional fluorescent reporters are treated as cell parameters that can be extracted once the segmentation and tracking is complete. The three channels are bias-field corrected and intensity normalized. A histogram-based intensity normalization is then performed on each of the three images. A selected percentage of the brightest and darkest intensities are clipped while the remaining values are linearly mapped to the range of [0, 1]. Next, the 3 channels are concatenated to form pseudo-RGB images (red: nuclei, green: cytoplasm, blue: DIC). Manual annotation of the cells are fed into the deep learning network as the input

representations with the manual annotations as the ground truths that guide the training process.

Data augmentation is done by random horizontal and vertical flipping, as well as random patch selection on the $4,792 \times 3,200$ regions yielding patches of dimension $1,600 \times 1,600$. Scaling and affine deformations are not considered here as they may alter naturally-occurring shape features. Also, data augmentation is only implemented for the training and validation sets. In total, we used approximately 4,600 cells in the training set, 1,000 in the validation set and 500 in the test set.

B. Deep Learning Architecture for Segmentation

The deep learning architecture is a five-layered Residual Attention U-Net (ResAttU-Net), inspired by an open-source GitHub repository [7]. The major differences between this architecture and the popular U-Net is the addition of two pieces in the former: Residual Blocks introduced in He at el [8] and Attention Mechanism introduced in Oktay et al [9].

C. The Edge-Enhancement Approach

Due to the desire to retain the long cellular extensions characteristic to our cell-line of interest, BHK-21, we implemented is a novel loss function approach that called edgeenhancement (EE). A loss function is defined as the binary cross-entropy loss between the deep-learning-predicted segmentation and the manually-annotated segmentation ground truth after flattening them into vectors (Figure 2a). For edge-enhancement, we emphasize the weighting on the accurate prediction of the cell edges (Figure 2b). We find the Laplacian-of-Gaussian of the ground truth segmentation, take two binary-thresholded versions of the resulting image that respectively corresponds to the inner and outer cell boundaries, vectorize the foreground regions as well as their corresponding regions in the deep learning prediction, and append them to their respective segmentation vectors for loss calculation.

D. Post-processing of the Segmented Cells

Once the deep learning cell segmentation is generated, it is refined by watershed splitting of adjacent touching cells. The cleaned up segmentation result can be directly provided as input to any suitable downstream tracking algorithm.

E. Evaluation of Cell Segmentation: Traditional Metrics

The segmentation performance of the candidates being compared are listed in Table I. Among all candidates, ResAttU-Net with our edge-enhancement approach achieved the best performance in all metrics evaluated. With 97% sensitivity and 93% specificity, the method we propose outperformed a segmentation benchmark, U-Net [10].

F. Evaluation of Cell Segmentation: Shape Feature Preservation

We first grouped a total number of 85 randomly selected cells from the test set into three shape categories based on their convexity, as calculated by the area of the cell divided by the area of its convex hull. Highly-convex cells



Fig. 2. Illustration of the Edge-Enhancement Approach on an Example Patch of Dimension 8×8 . a. Typical calculation of binary cross entropy loss. b. Proposed calculation for binary cross entropy loss by emphasizing edge information.

(convexity > 0.95) are categorized as round, least-convex cells (convexity < 0.5) as star-shaped, while the ones in between as spindle-shaped. Afterwards, five shape features were extracted from the respective segmentation mask of each cell. From Figure 3 and Figure II, it can be seen that our proposed method performed the best in terms of cell shape feature preservation. As a result, the mean intensity values of the three information-providing channels within the segmentation mask given by our proposed method also best resembled those in the manual annotation ground truth masks. It is reasonable to believe that the same will hold true for other information to be extracted from these segmentation results.

TABLE I Evaluation of the segmentation using traditional Quantitative metrics.

Architecture	Sensitivity	Specificity	F1-score	Jaccard Similarity	Dice Coefficient
U-Net	0.91	0.91	0.91	0.84	0.91
U-Net EE	0.94	0.88	0.91	0.84	0.91
ResAtt U-Net	0.90	0.76	0.82	0.70	0.82
ResAtt U-Net EE	0.97	0.93	0.95	0.91	0.95

IV. DISCUSSIONS

A. Irreplaceability of Shape Feature Metrics

Traditional evaluation metrics (Table I) imply barely distinguishable performance of the two U-Net candidates regardless of edge-enhancement. However, the shape feature metrics (Figure II) reveals their advantages and disadvantages in a more comprehensive manner. We can thus conclude that shape feature metrics provide insights that are not reflected in the traditional metrics.

B. Effect of Edge-enhancement on U-Net

The effect of edge-enhancement on U-Net is better explained by the shape features. The version with edgeenhancement performed slightly better for the spindle-shaped cells but slightly worse for cells in the other two categories.

TABLE II Evaluation of the segmentation using traditional quantitative metrics.

Cell type	Metrics	Manual (GT)	ResAttU-Net-EE	U-Net-EE	U-Net	ResAttU-Net
Spheroid	Area	1380 ± 90	6.6 ± 6.1 %	17±10%	14.6 ± 6.5 % (****)	28 ± 11 %
(N = 33)	Eccentricity	0.638 ± 0.037	1.10 ± 0.64 %	5.9 ± 1.3 % (**)	4.95 ± 0.93 %	17.2 ± 4.2 %
	Maior Axis Length	49.8 ± 2.0	4.4 ± 4.3 %	13 ± 10 %	7.9 ± 4.6 % (****)	19.9 ± 8.9 % (*)
	Solidity	0.9721 ± 0.0017	$1.2 \pm 1.1\%$	2.1 ± 1.5 %	1.82 ± 0.97 %	4.65 ± 0.92 % (****)
	Orientation	0.814 ± 0.080	0.51 ± 0.34 %	4.08 ± 0.84 % (***)	4.9 ± 1.4 % (**)	11.5 ± 2.5 % (****)
	Mean Nuc Intensity	0.664 ± 0.014	1.5 ± 1.5 %	5.6 ± 1.7 %	6.2 ± 1.4 % (****)	15.3 ± 3.9 % (*)
	Mean Cyto Intensity	0.575 ± 0.013	1.5 ± 1.3 %	$5.5 \pm 1.6 \%$	6.1 ± 1.3 % (****)	14.9 ± 3.6 % (*)
	Mean DIC Intensity	0.4950 ± 0.0076	$0.38 \pm 0.29 \%$	0.93 ± 0.34 %	$0.93 \pm 0.28 \%$	$1.8 \pm 0.68 \%$
Spindle-shaped (N = 23)	Area	2380 ± 210	46 ± 19 %	47 ± 17 %	54 ± 19 % (*)	42.5 ± 11.3 % (*)
	Eccentricity	83.266 ± 0.026	$10.0 \pm 3.6 \%$	10.7 ± 3.4 %	13.6 ± 3.7 % (*)	$13.6 \pm 3.0 \%$
	Major Axis Length	84.1 ± 5.6	$13.8 \pm 7.3 \%$	$14.8 \pm 4.8 \%$	$20.2 \pm 9.1 \%$	15.5 ± 4.3 % (*)
	Solidity	$0.8610\ \pm\ 0.0056$	$2.62 \pm 0.77 \%$	3.96 ± 0.67 %	$4.7 \pm 1.1 \%$	5.6 ± 1.2 %
	Orientation	0.88 ± 0.11	21.7 ± 7.4 %	22.3 ± 7.3 %	$22.4 \pm 7.0 \%$	27.4 ± 6.4 %
	Mean Nuc Intensity	0.439 ± 0.041	6.8 ± 2.8 %	9.4 ± 2.7 % (*)	$10.3 \pm 3.0\% (*)$	$14.0 \pm 3.4\%$
	Mean Cyto Intensity	0.401 ± 0.026	3.1 ± 1.1 %	5.6 ± 1.2 % (*)	6.4 ± 1.7 % (*)	$10.8 \pm 2.6 \%$
	Mean DIC Intensity	$0.4790 \pm \ 0.0038$	$0.62\ \pm 0.23\ \%$	$0.83\ \pm 0.22\ \%$	1.2 ± 1.7 % (*)	$1.63 \pm \ 0.39 \ \%$
Stellate (N = 5)	Area	7290 ± 120	$10.5 \pm 6.6 \%$	16.1 ± 6.8 % (*)	$12.3 \pm 5.8 \%$	35.0 ± 9.4 % (*)
	Eccentricity	0.84 ± 0.05	5.6 ± 2.4 %	$6.3 \pm 2.2 \%$	$6.8 \pm 2.7 \%$	14.3 ± 5.9 % (*)
	Major Axis Length	180 ± 16	$8.1 \pm 6.1 \%$	16.0 ± 4.8 % (*)	11.4 ± 5.2 % (*)	40.7 ± 7.9 % (*)
	Solidity	0.502 ± 0.066	$39 \pm 20 \%$	44.2 ± 19.8 % (*)	$38 \pm 18\%$	77 ± 24 % (*)
	Orientation	0.27 ± 0.13	$11.5 \pm 6.8 \%$	$11.8 \pm 7.4 \%$	$12.8 \pm 7.9 \%$	$11.4 \ \pm 7.9 \ \%$
	Mean Nuc Intensity	0.251 ± 0.015	$10.2 \pm 9.1 \%$	17.1 ± 9.1 % (*)	$11.6 \pm 8.7 \%$	29.3 ± 7.4 % (*)
	Mean Cyto Intensity	$0.279\ \pm\ 0.018$	$7.7 \pm 6.7 \%$	13.6 ± 6.7 % (*)	$8.8~\pm~6.3~\%$	27.3 ± 7.5 % (*)
	Mean DIC Intensity	$0.47565 \ \pm \ 0.00098$	$2.66\pm\ 0.12\ \%$	$0.24 \pm 0.11 \%$	$0.32~\pm~0.13~\%$	$0.345~\pm~0.077~\%$
Other (N = 24)	Area	$2430\pm\ 230$	$6.1 \pm 2.3 \%$	$9.7 \pm 2.5 \%$	12.2 ± 2.8 % (*)	21.1 ± 3.8 % (*)
	Eccentricity	$0.830 \ \pm \ 0.029$	$4.2 \pm 1.6 \%$	5.5 ± 1.7 % (*)	$7.6 \pm 2.0 \%$	$14.1 \pm 6.7 \%$
	Major Axis Length	86.6 ± 8.1	7.4 ± 2.7 %	9.8 ± 2.7 % (*)	12.8 ± 2.8 %	$17.1 \pm 3.2 \ \%$
	Solidity	0.878 ± 0.018	$4.1\pm1.4~\%$	$4.4\pm1.0~\%$	$4.8\pm1.3~\%$	5.8 ± 1.5 %
	Orientation	0.857 ± 0.093	$3.0\pm1.1~\%$	$3.77 \pm 0.91~\%$	$4.14 \pm 0.88~\%$	$10.9 \pm 4.5 \ \%$
	Mean Nuc Intensity	0.527 ± 0.043	$3.1\pm1.0~\%$	5.5 ± 1.0 %	6.8 ± 1.0 % (*)	$17.5 \pm 4.5 \ \%$
	Mean Cyto Intensity	0.446 ± 0.026	$3.0\pm1.0~\%$	5.6 ± 1.1 %	$6.85 \pm 0.99 \ \% \ (*)$	$17.6 \pm 4.0 \ \%$
	Mean DIC Intensity	0.4737 ± 0.0041	$0.38 \pm 0.12~\%$	0.67 ± 0.13 %	$0.65 \pm 0.13~\%$	1.61 ± 0.74 %

However, more spindle shapes were more prevalent in the training data. It is possible that U-Net does gain a better grasp of shape features under the help of edge-enhancement, but the improvement is somewhat coarse and is tailored better towards more frequently occurring cell shapes. The deep learning model, despite having difficulty learning the nuances of cells from all shape categories, is nevertheless able to provide a better segmentation for cells from the most prevalent shape category.

C. Effect of Edge-enhancement on ResAttU-Net

Unlike the more ambiguous case with U-Net, the edgeenhancement approach boosts the segmentation performance of ResAttU-Net on almost every aspect. We hypothesize that a network with the attention mechanism would benefit more from edge-enhancement. Indeed this was the case, since the attention mechanism enabled the network to better



Fig. 3. Visualization of the segmentation and demonstration of the cell features. The original pseudo-RGB images (R: nuclei, G: cytoplasm, B: DIC) are overlayed with the respective segmentation masks in light green. The cytoplasm extensions are best captured by the ResAttU-Net with edge enhancement (see Section 3.4). The percentage change of each shape feature with respect to the manual ground truth also indicates a better shape feature preservation in that candidate. Note that for orientation, 100 percent difference represents a mismatch of 90 degrees.

understand / encode the regional contribution to the loss, and therefore the network was able to know which parameters to adjust in response to the training feedback. In this way, it was better able to adjust itself to preserve the shape features for all cell shapes.

D. Potential Utility of our Proposed Candidate

The ability to track cells over time provides insight into processes ranging from bacterial dynamics to mammalian development [11]. A central challenge, however, is robust cell segmentation. This is especially challenging for identifying the entirety of the cytoplasm of complex and motile cells who may have irregular shapes and/or long and thin cellular processes. This contrasts with simple alternative methods that rely exclusively on identifying the nucleus of a cell, but which necessarily exclude the study of cytoplasmic processes or fluorescent reporters not amenable to nuclear reporters. The method described here helps overcome the limitations of alternative segmentation approaches and preserves the unusual processes and extensions observed in certain primary cells or cell lines, and will likely find relevance beyond the BHK-21 line studied here.

ACKNOWLEDGMENT

This study was initiated by Zakary S. Singer and Tal Danino from Columbia Biomedical Engineering. Computational resources are provided by Jia Guo and Tommy Vaughan from Mortimer B. Zuckerman Mind Brain Behavior Institute.

REFERENCES

[1] Mhalla, Ala, et al. "Spatio-Temporal Object Detection by Deep Learning: Video-Interlacing to Improve Multi-Object Tracking." Image and Vision Computing, vol. 88, 2019, pp. 120–131., https://doi.org/10.1016/j.imavis.2019.03.002.

- [2] Bessis, Marcel, et al. Red Cell Shape Physiology, Pathology, Ultrastructure. Springer, 1973.
- [3] M. A. A. Dewan, M. O. Ahmad and M. N. S. Swamy, "Tracking Biological Cells in Time-Lapse Microscopy: An Adaptive Technique Combining Motion and Topological Features," in IEEE Transactions on Biomedical Engineering, vol. 58, no. 6, pp. 1637-1647, June 2011, doi: 10.1109/TBME.2011.2109001.
- [4] Al-Hafiz, Fatimah, et al. "Red Blood Cell Segmentation by Thresholding and Canny Detector." Proceedia Computer Science, vol. 141, 2018, pp. 327–334., https://doi.org/10.1016/j.procs.2018.10.193.
- [5] Schiegg, Martin et al. "Graphical model for joint segmentation and tracking of multiple dividing cells." Bioinformatics (Oxford, England) vol. 31,6 (2015): 948-56. doi:10.1093/bioinformatics/btu764
- [6] Falk, Thorsten et al. "U-Net: deep learning for cell counting, detection, and morphometry." Nature methods vol. 16,1 (2019): 67-70. doi:10.1038/s41592-018-0261-2
- [7] LeeJunHyun. "LeeJunHyun/imagesegmentation: Pytorch Implementation of U-Net, R2U-Net, Attention U-Net, and Attention R2U-Net." GitHub, https://github.com/LeeJunHyun/ImageSegmentation.
- [8] He, Kaiming, et al. "Deep Residual Learning for Image Recognition." ArXiv.org, 10 Dec. 2015, https://arxiv.org/abs/1512.03385.
- [9] Oktay, Ozan, et al. "Attention U-Net: Learning Where to Look for the Pancreas." ArXiv.org, 20 May 2018, https://arxiv.org/abs/1804.03999.
- [10] Falk, Thorsten et al. "U-Net: deep learning for cell counting, detection, and morphometry." Nature methods vol. 16,1 (2019): 67-70. doi:10.1038/s41592-018-0261-2
- [11] Muzzey, Dale, and Alexander van Oudenaarden. "Quantitative time-lapse fluorescence microscopy in single cells." Annual review of cell and developmental biology vol. 25 (2009): 301-27. doi:10.1146/annurev.cellbio.042308.113408