Cell Detection with Star-convex Polygons

Uwe Schmidt*, Martin Weigert*, Coleman Broaddus & Gene Myers

Max Planck Institute of Molecular Cell Biology and Genetics, Dresden, Germany

Center for Systems Biology Dresden, Germany

{uschmidt|mweigert|broaddus|myers}@mpi-cbg.de

* Equal contribution

Abstract

Automatic detection and segmentation of cells and nuclei in microscopy images is important for many biological applications. Recent successful learning-based approaches include per-pixel cell segmentation with subsequent pixel grouping, or localization of bounding boxes with subsequent shape refinement. In situations of crowded cells, these can be prone to segmentation errors, such as falsely merging bordering cells or suppressing valid cell instances due to the poor approximation with bounding boxes. To overcome these issues, we propose to localize cell nuclei via *star-convex polygons*, which are a much better shape representation as compared to bounding boxes and thus do not need shape refinement. To that end, we train a convolutional neural network that predicts for every pixel a polygon for the cell instance at that position. We demonstrate the merits of our approach on two synthetic datasets and one challenging dataset of diverse fluorescence microscopy images.

1. Overview of Instance Segmentation

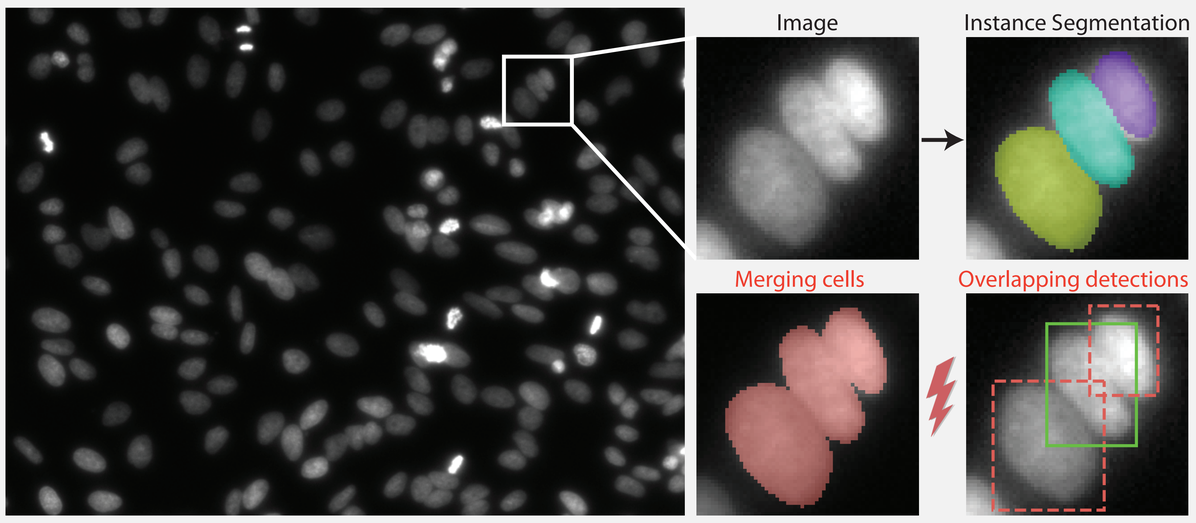

Instance segmentation of cells from microscopy images

- Automatic *detection and segmentation* of cells and nuclei in microscopy images is important for many biological applications [1].
- A common goal is to obtain an *instance segmentation*, which is the assignment of a cell instance identity to every pixel of the image
- Images often contain many *crowded* cells with touching borders, which makes segmentation difficult.

Common Approaches

- Classification of every pixel into semantic classes (*e.g. background*, *border*, *cell*) and subsequent grouping *e.g.* via connected components [2, 3]
- Localization of proposal cell instances with bounding boxes and subsequent mask refinement (*e.g.* Mask-RCNN [4])

Sources of segmentation errors

- Merging of touching cells
- Suppression of valid cell instances due to large overlap of bounding box localization

3. Quantitative Comparisons

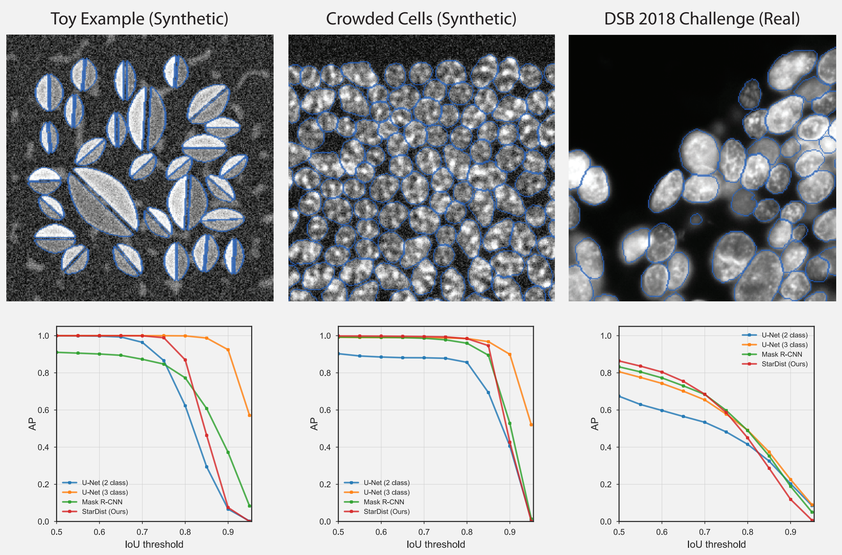

- *Datasets:* 2 synthetic and 1 real dataset that highlight different challenges.
- *Compared Methods:* U-Net with 2 and 3 classes [2, 3], Mask-RCNN [4].
- *Evaluation metric:* Average precision for several IoU thresholds τ

$$AP_\tau = \frac{TP_\tau}{TP_\tau + FN_\tau + FP_\tau}.$$

- StarDist achieves higher AP than all other compared methods (for $\tau < 0.7$)

5. Shape Completion

- StarDist can be trained to do shape completion for *partially visible* objects at the image boundary.
- Example Image: StarDist polygon candidate predictions for 200 random pixels (left) and for all pixels after non-maximum suppression (right); pixels and associated polygons are color-matched.

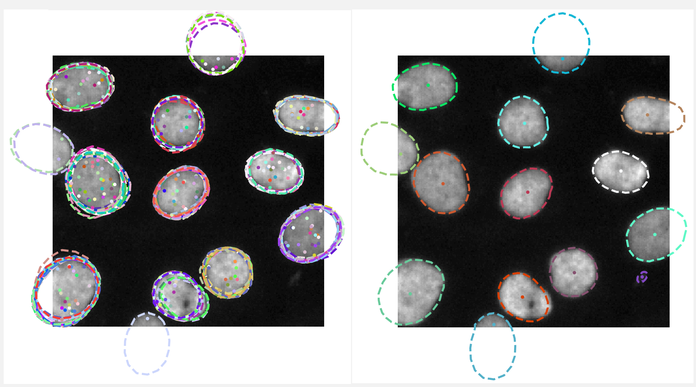

2. Our Approach – StarDist

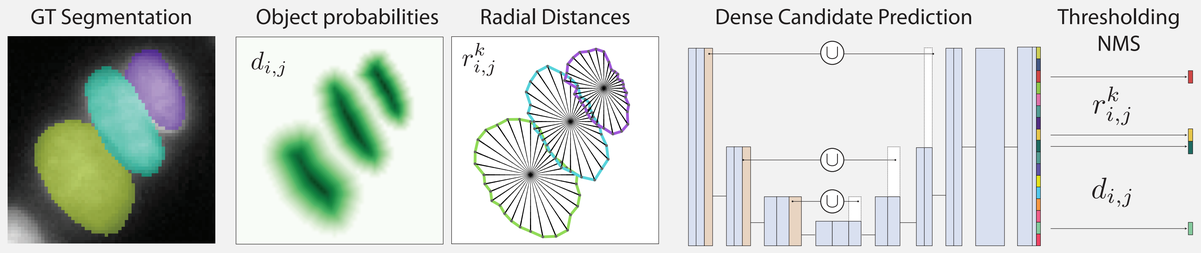

Training

Given training data with ground truth instances, we compute for each pixel $p_{i,j}$

- an *object probability* $d_{i,j} \in [0, 1]$ as the (normalized) Euclidean distance to the nearest background pixel, and
- the *radial distances* $r^k_{i,j} \in \mathbb{R}^+$ as the smallest distance from $p_{i,j}$ to the object boundary (if $p_{i,j}$ is inside an object) along $k = 1 \dots K$ equally spaced directions (we use $K = 32$).

We train a convolutional neural network (using a U-Net [2] as backbone) to densely predict both $d_{i,j}$ and $r^k_{i,j}$.

Inference

1. Predict *object probabilities* $d_{i,j}$ and *radial distances* $r^k_{i,j}$ from input image via trained neural network.
2. Identify *polygon candidates* from pixels with object probabilities above a threshold.
3. Perform *non-maximum suppression* of candidates to remove extraneous polygons.

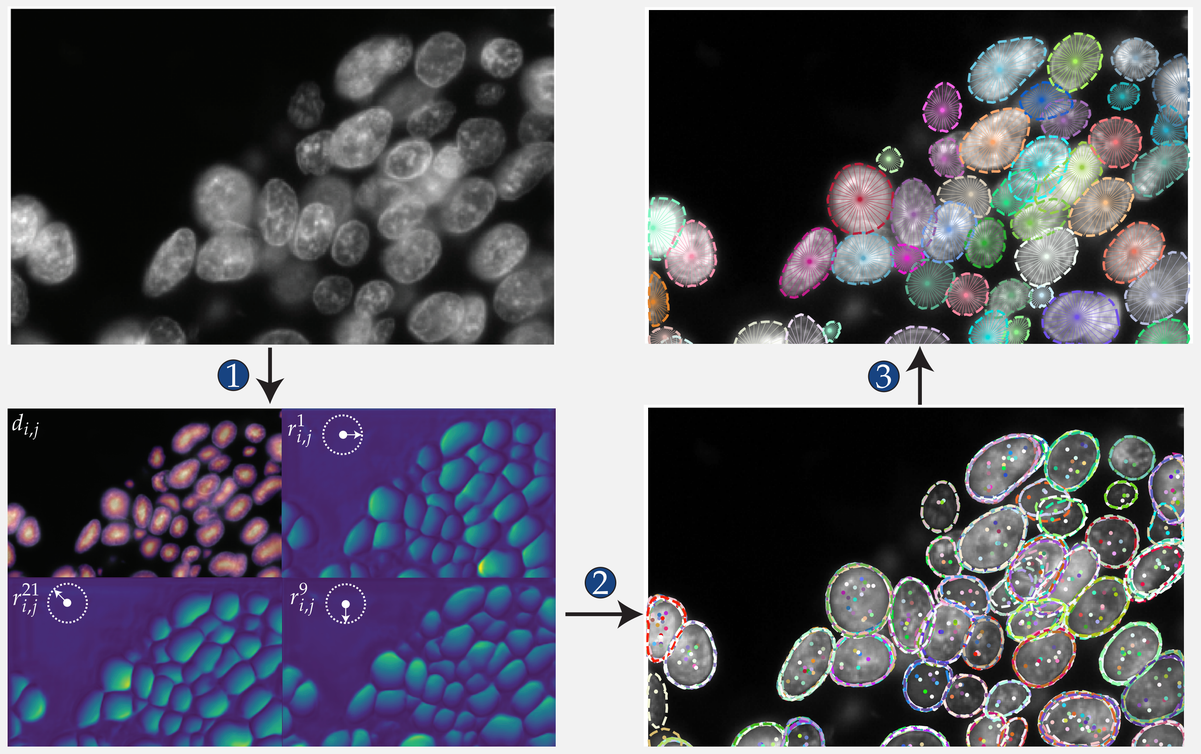

4. Segmentation Examples

TP Correct predictions (true positives).

FP Wrong predictions (false positives).

FN Missing predictions (false negatives) are indicated by blue outlines.

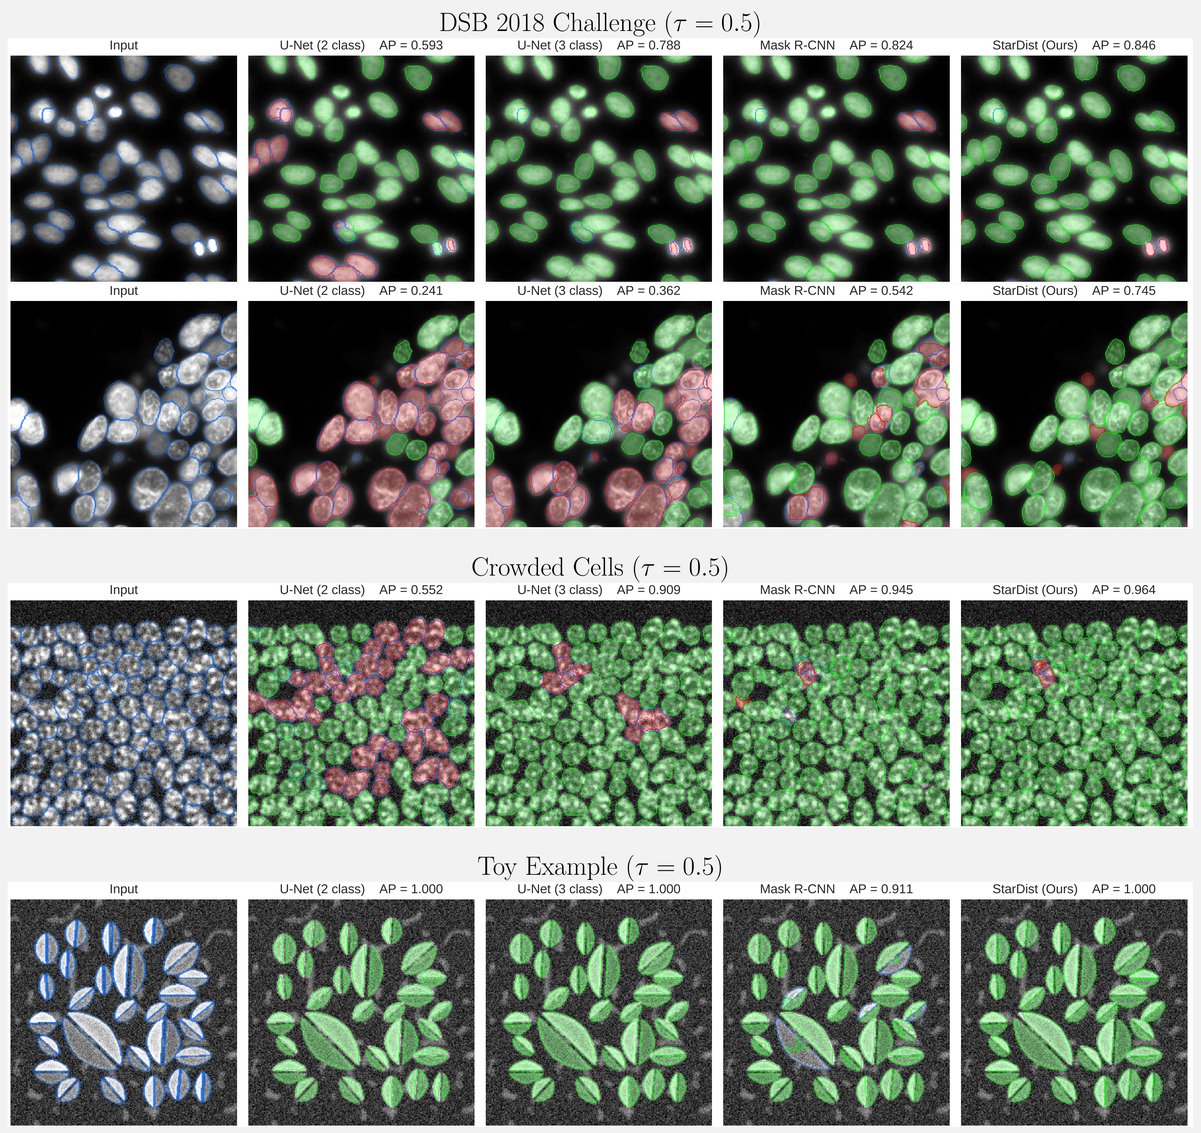

References

[1] Vladimír Ulman et al. "An objective comparison of cell-tracking algorithms". In: *Nature methods* 14.12 (2017), p. 1141.

[2] Olaf Ronneberger et al. "U-Net: Convolutional Networks for Biomedical Image Segmentation". In: *MICCAI*. 2015.

[3] Hao Chen et al. "DCAN: Deep contour-aware networks for accurate gland segmentation". In: *CVPR*. 2016.

[4] Kaiming He et al. "Mask R-CNN". In: *ICCV*. 2017.

Code available at

https://github.com/mpicbg-csbd/stardist