Star-convex Polyhedra for 3D Object Detection and Segmentation in Microscopy

Martin Weigert[1,2,3,⋆] Uwe Schmidt[2,3,⋆] Robert Haase[2,3] Ko Sugawara[4,5] Gene Myers[2,3]

[1]Institute of Bioengineering, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland [2]Max Planck Institute of Molecular Cell Biology and Genetics (MPI-CBG), Dresden, Germany [3]Center for Systems Biology Dresden (CSBD), Germany [4]Institut de Génomique Fonctionnelle de Lyon (IGFL), École Normale Supérieure de Lyon, France [5]Centre National de la Recherche Scientifique (CNRS), Paris, France ⋆*Equal contribution*

martin.weigert@epfl.ch, uschmidt@mpi-cbg.de @martweig, @uschmidt83

Motivation & Contributions

- Instance segmentation of cell nuclei is important in many biomedical projects
- Common approaches struggle with *noisy* images and *dense packing* of nuclei
- *StarDist* [1] alleviates these problems by using *star-convex polygons* to describe the typically roundish shapes of cell nuclei, but only for 2D images

Contributions (extension of *StarDist* [1] from 2D to 3D):

- Faithful *star-convex polyhedra* representation of 3D cell nuclei via judicious selection of radial directions, also for typical anisotropic voxels in microscopy
- Efficient intersection computation/bounds for pairs of star-convex polyhedra, necessary to make non-maximum suppression practical for large 3D volumes
- Superior results on two challenging datasets, especially with little training data

Instance Segmentation of Cell Nuclei

Common approaches:

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

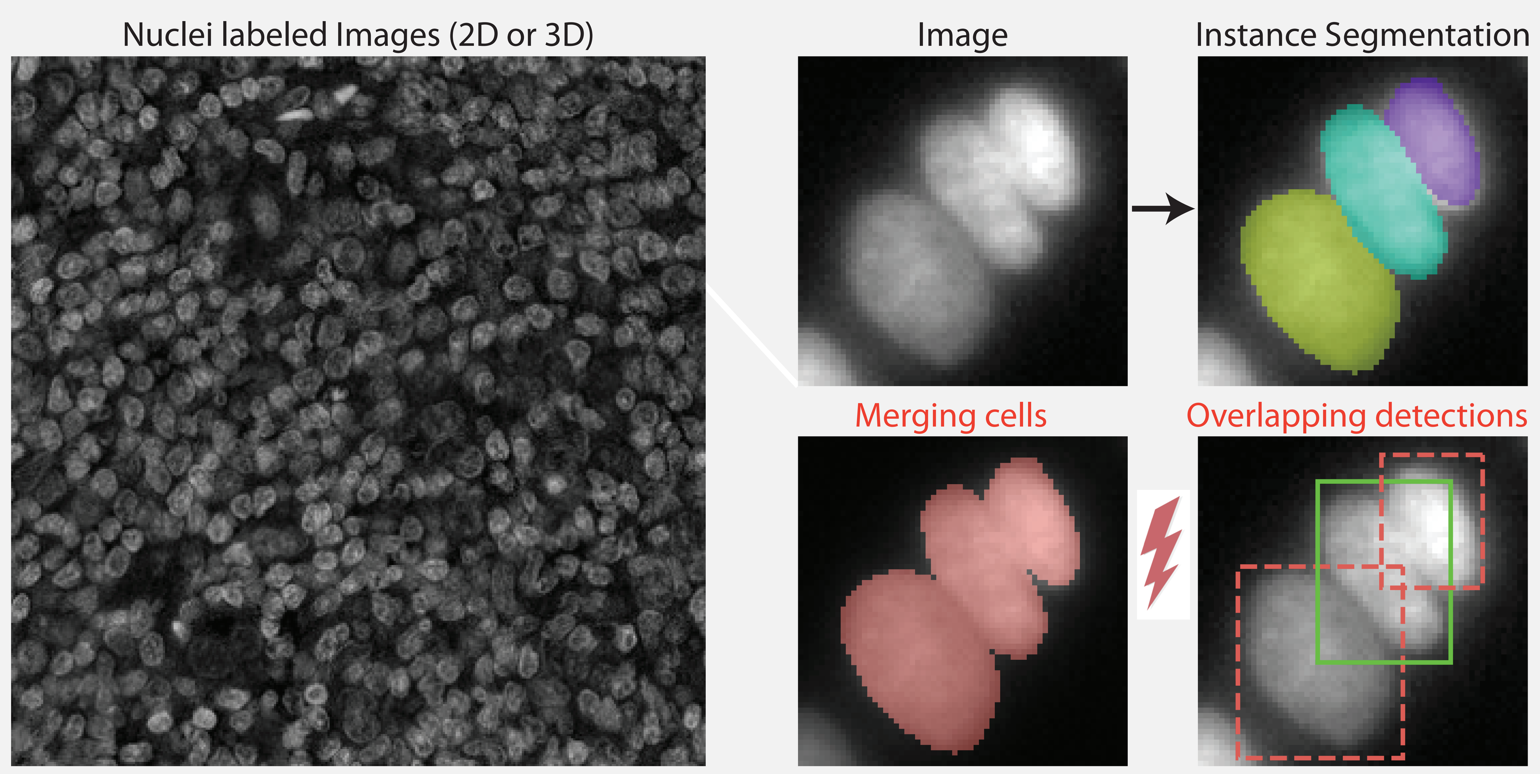

Segmentation errors: Images often contain many *densely-packed* cells with touching borders. Segmentation results often exhibit *merging of touching cells* and *suppression of valid cell instances* due to large overlap of bounding box localizations.

StarDist Examples

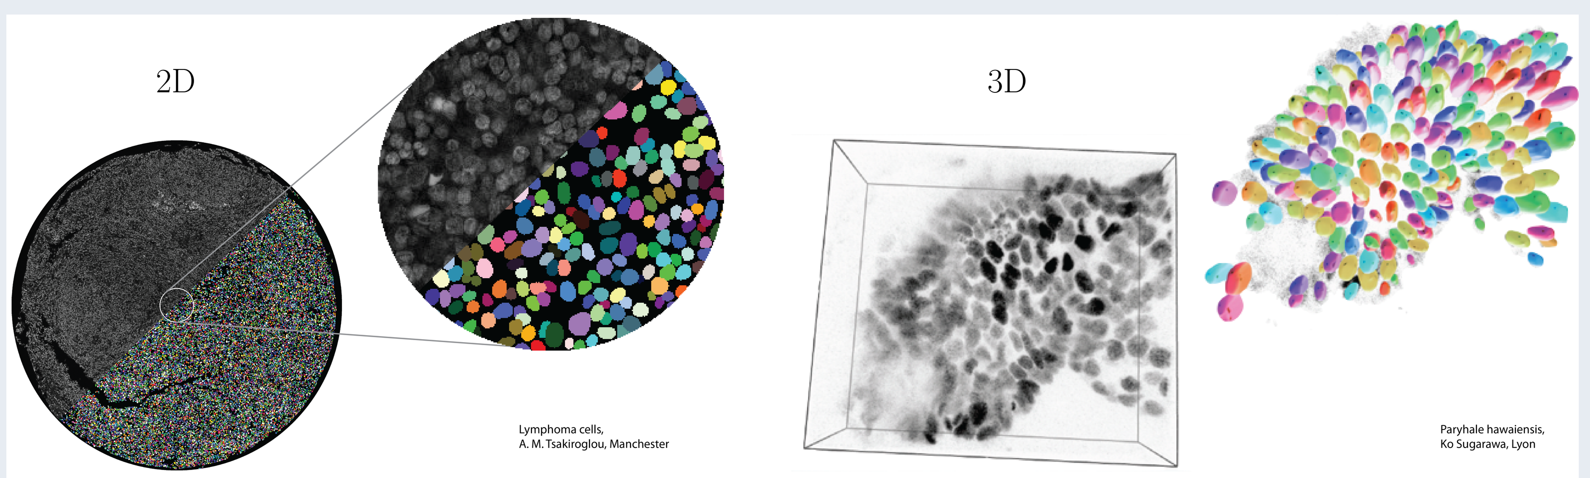

StarDist – Object Detection with Star-convex Shapes

Training

Given training data with ground truth instances, we compute for each pixel

- an *object probability* p as the (normalized) distance to the nearest background pixel, and
- the *radial distances* d_k to the boundary of the object that the pixel belongs to.

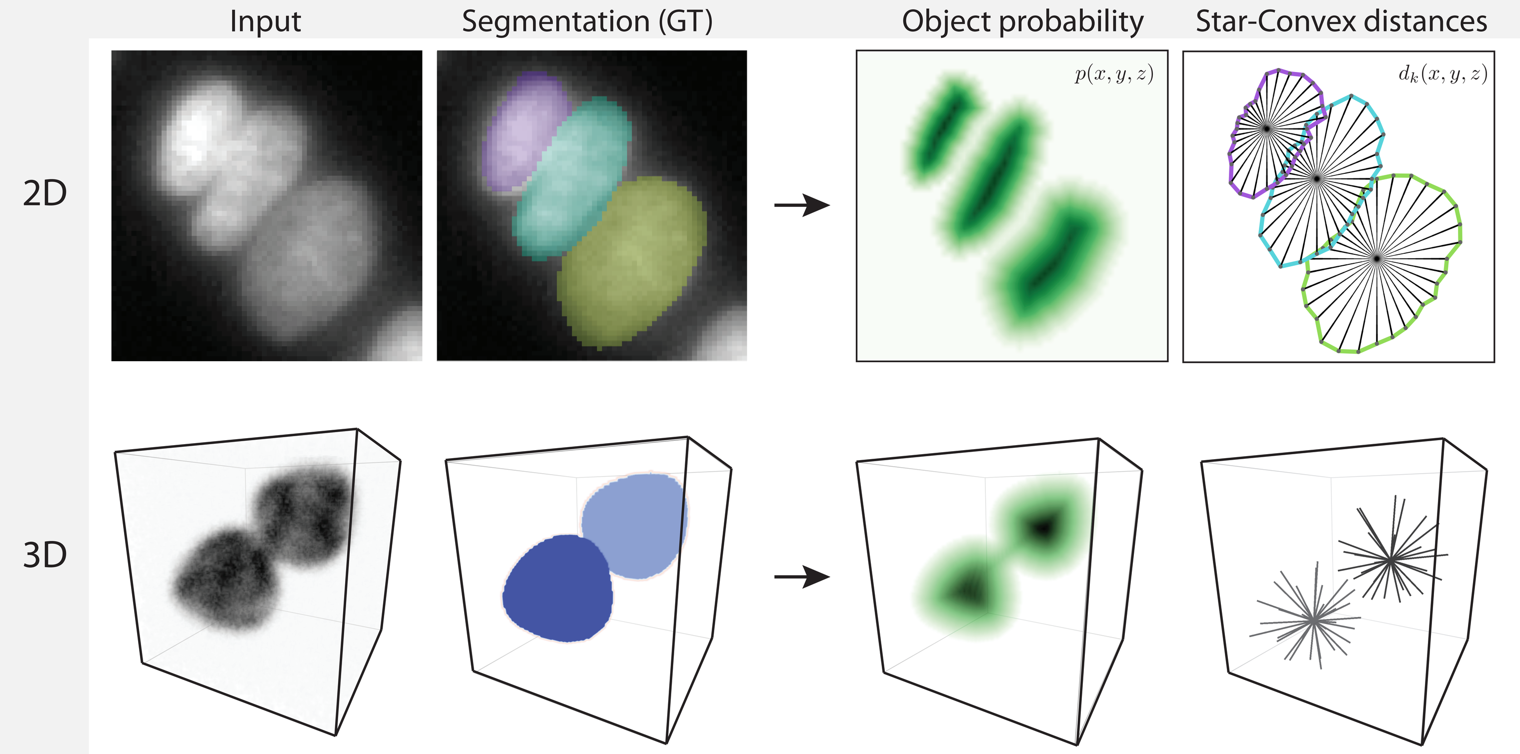

We train a CNN (with U-Net [2] or ResNet [3] backbone) to densely predict both p and d_k.

Choice of radial directions:

- In 2D, we select equidistant radial directions in polar coordinates.
- In 3D, we choose radial directions corresponding to a *Fibonacci lattice* [4] of approximately equally distributed points on a sphere (or ellipsoid for anisotropic data).

Inference

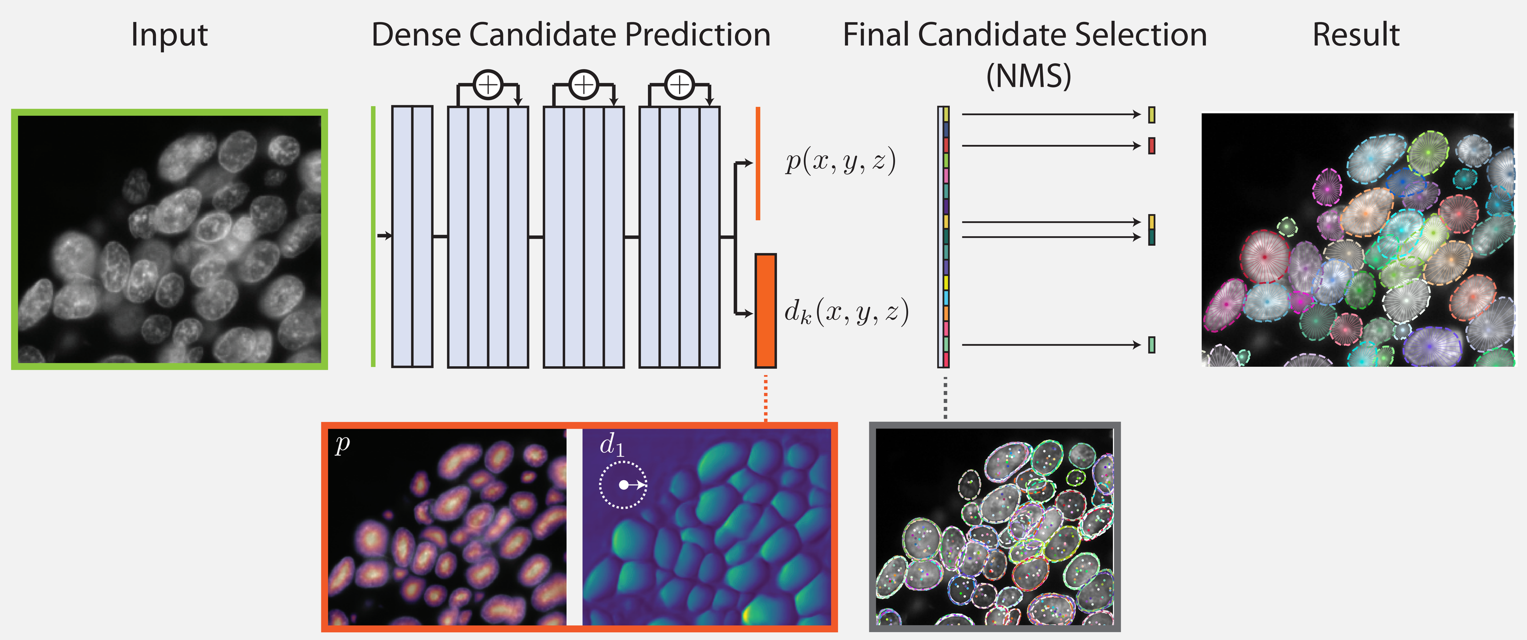

- *Dense candidate prediction*: Predict *object probabilities* p and *radial distances* d_k from input image. Identify polygon/polyhedra object candidates from pixels with p above a threshold.
- *Final candidate selection*: Perform typical overlap-based *non-maximum suppression* (NMS) of candidates (sorted by their probabilities p) to remove redundant object proposals.

2D: Computing the overlap of two polygons is rather easy, and there are good implementations.
3D: Efficiently computing the overlap of two star-convex polyhedra is challenging, therefore we use a series of bounds (see below) to check for overlap. In practice, exact but expensive computation via rasterization is only rarely necessary.

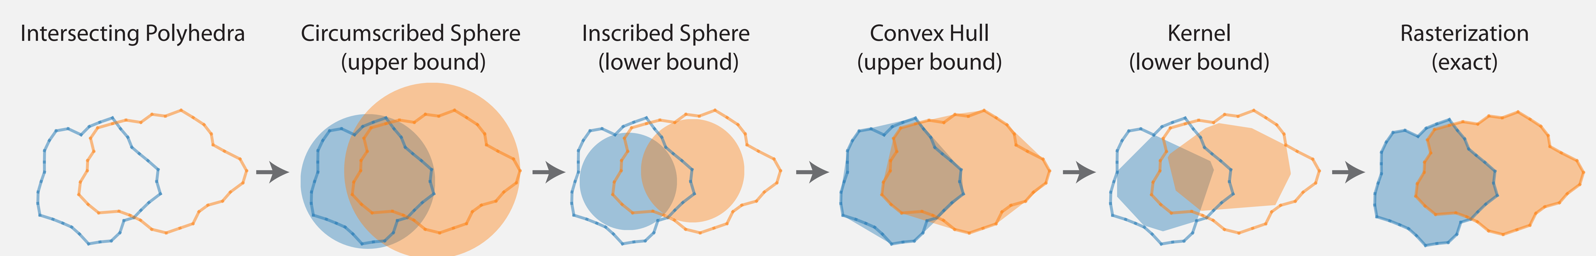

Code, Documentation, Examples

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

Comparison & Examples

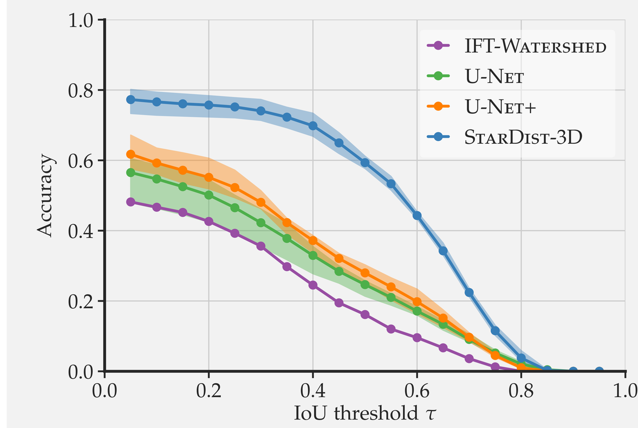

We compare our StarDist-3D approach against

- a classical method (IFT-Watershed [5])
- 3-class 3D U-Net [6] (backgr., cell interior, cell boundary)
- 3D U-Net with watershed postprocessing (U-Net+)

Below: Colors denote instances, *i.e.* correct predictions have the same color as GT. False positives in red hues, false negatives not highlighted.

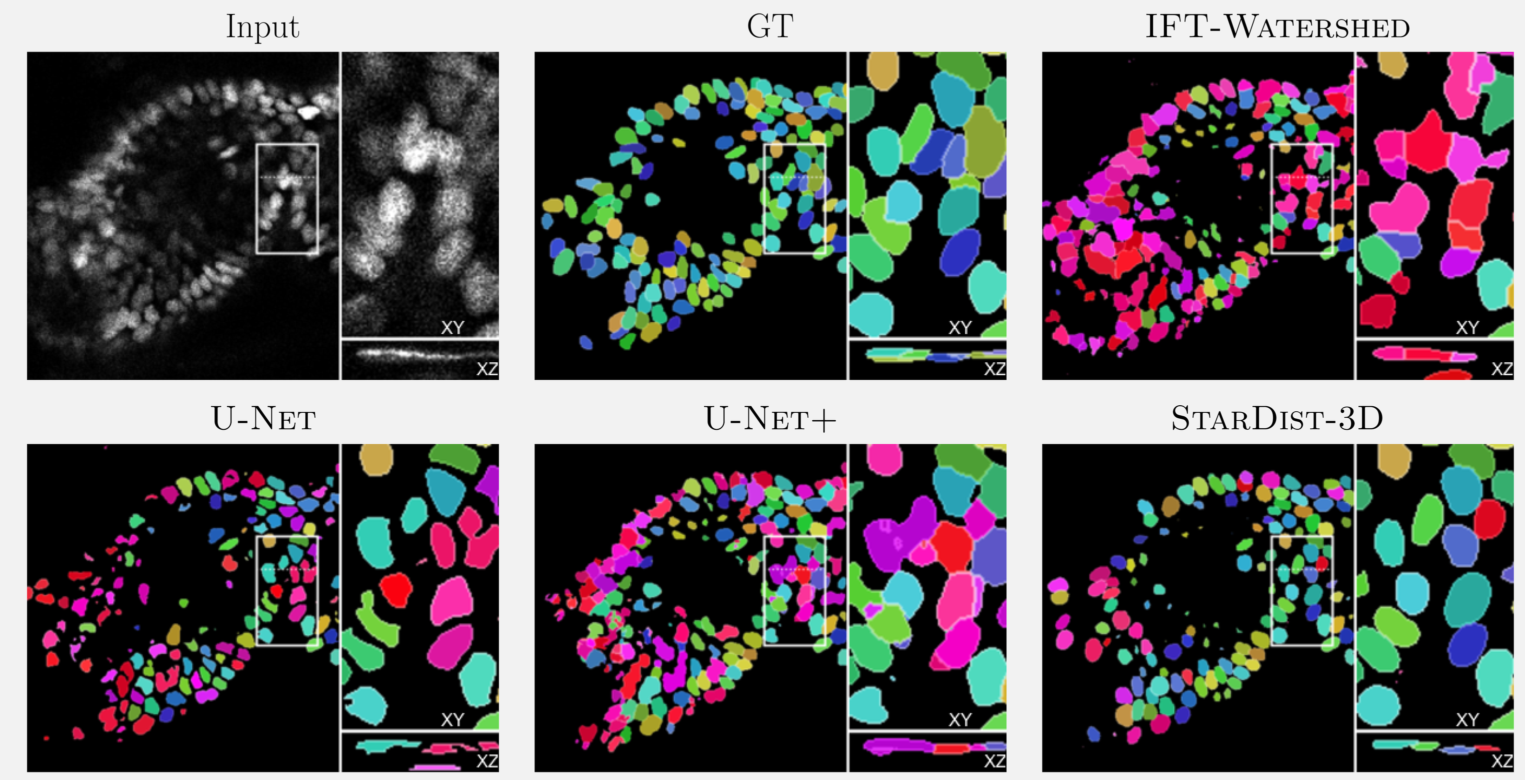

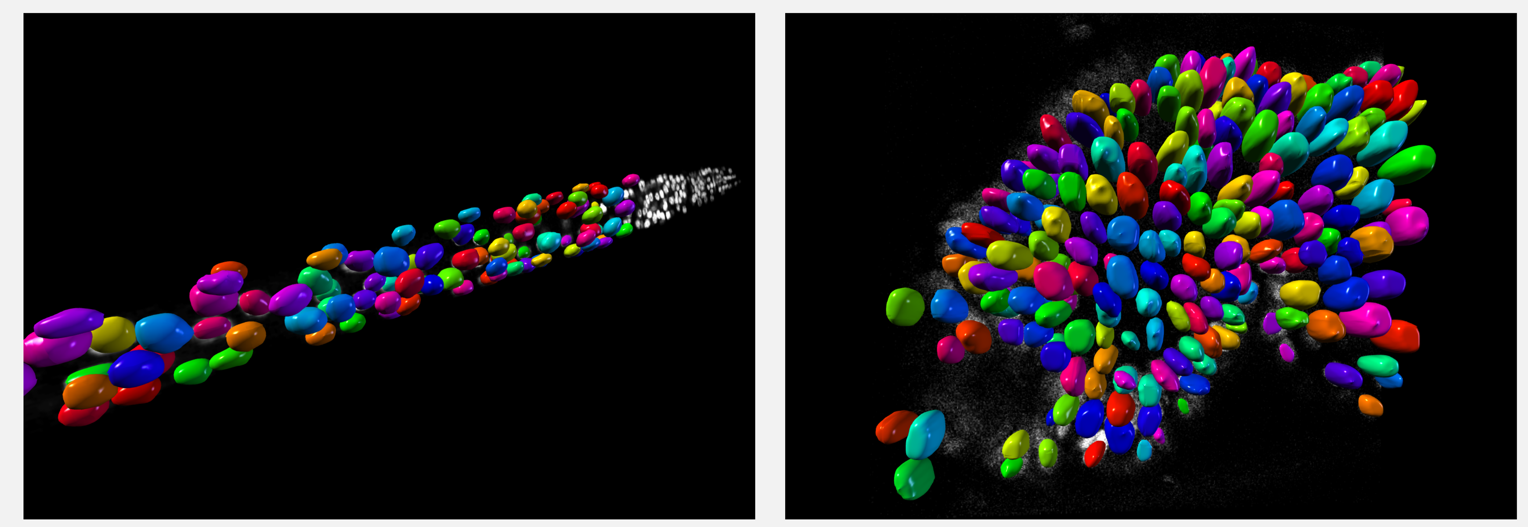

Example results of StarDist-3D for two challenging 3D fluorescence microscopy datasets. Each instance of a predicted cell nucleus is assigned a random color (not all shown on the left).

References

[1] Uwe Schmidt et al. "Cell Detection with Star-Convex Polygons". In: *MICCAI*. 2018.
[2] Olaf Ronneberger et al. "U-Net: Convolutional Networks for Biomedical Image Segmentation". In: *MICCAI*. 2015.
[3] Kaiming He et al. "Mask R-CNN". In: *ICCV*. 2017.
[4] Álvaro González. "Measurement of areas on a sphere using Fibonacci and latitude–longitude lattices". In: *Mathematical Geosciences* 42.1 (2010).
[5] R. A. Lotufo et al. "IFT-Watershed from gray-scale marker". In: *XV Brazilian Symp. on Comp. Graphics and Image Processing*. Oct. 2002.
[6] Juan C. Caicedo et al. "Evaluation of Deep Learning Strategies for Nucleus Segmentation in Fluorescence Images". In: *Cytometry Part A* 95.9 (2019), pp. 952–965.