Segmentation of Nuclei by Deep Regression of the Distance Map

Naylor et al., IEEE TMI 2019

Jan Kybic

2020

Nuclei instance segmentation

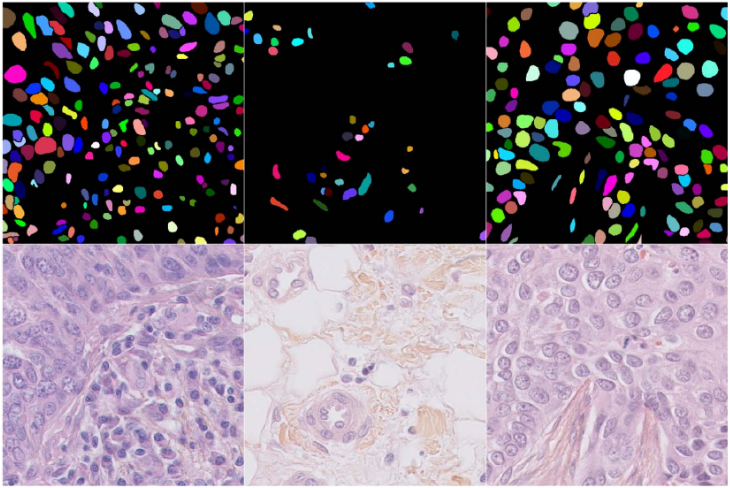

Datasets

- 50 images 512×512, breast 4022 annotated cell nuclei
- 30 images 1000×1000, different organs, 21623 annotated cell nuclei

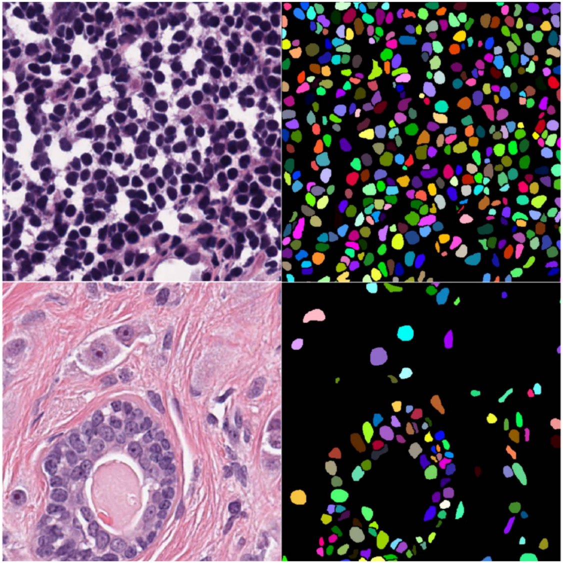

Classification network

- U-Net, softmax output
- logarithmic loss (cross-entropy)

$$\mathcal{L} = \frac{1}{N}\sum_{l=1}^{N} loss(B_l,\, f_2(A_l)) + \lambda ||w||_2^2$$

$$loss(B_l,\, f_2(A_l)) = \frac{1}{np}\sum_{i,j}\sum_{k} t_{i,j,k} \log(\widehat{p_{i,j,k}})$$

Distance map regression

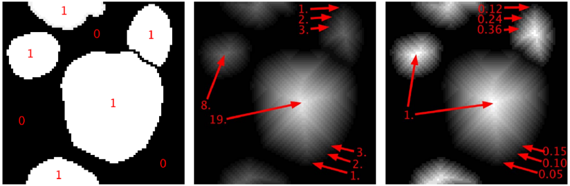

$$loss(B_{D,l}, \widehat{B_{D,l}}) = \frac{1}{np} \sum_{i,j} (B_{D,l}[i,j] - \widehat{B_{D,l}}[i,j])^2$$

- "chessboard" distance
- optionally: normalize to $[0, 1]$ for each component, softmax output

Distance map choice

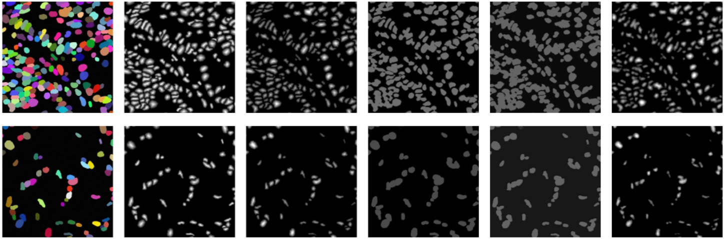

(a) $A \in \mathcal{A}$ (b) $B_{D,N} \in \mathcal{B}_{D,N}$ (c) $B_D \in \mathcal{B}_D$ (d) $\widehat{B_{D,N}} = f_1(A)$ (e) $\widehat{B_{D,N}}$ = softmax $\circ$ $f_2(A)$ (f) $\widehat{B_D} = f_1(A)$

Postprocessing

- threshold predicted distance map
- require significant drop from maximum$\rightarrow$seed

$$\min_{\substack{\mathcal{P}=(M,\ldots,M'),\\ y(M')>y(M)}} \{\max_{x\in\mathcal{P}} [y(M)-y(x)]\} > p_1$$

- watershed $+$ mask by thresholded distance map

Training

- **Augmentation:** rotation, mirroring, blurring, deformation, color perturbation (color deconvolution, affine histogram modification)
- Split dataset DS2 to training, validation (for hyperparameters), test
- **Networks:**
  - vgg16 *FCN* pretrained on ImageNet, fine tuning
  - *U-net,* batch size 10, learning rate decay
  - *Mask R-CNN (ResNet 101),* pretrained on COCO, batch size 4, learning rate decay

Mask R-CNN

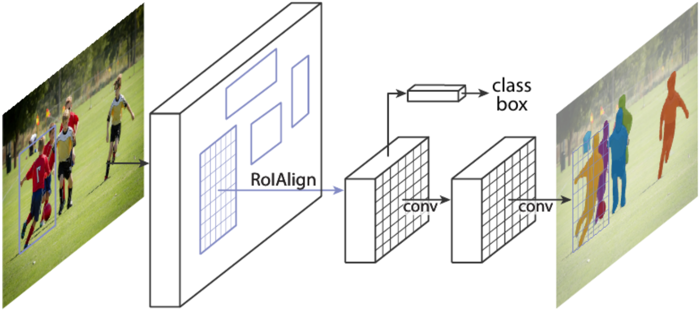

F1 score

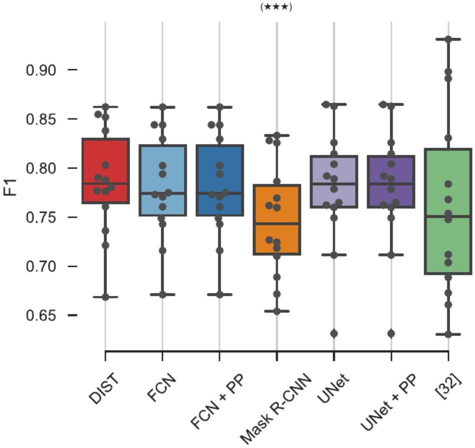
(★★★)
0.90
0.85
0.80
0.75
0.70
0.65
F1
DIST
FCN
FCN + PP
Mask R-CNN
UNet
UNet + PP
[32]

Average Jaccard index

- match GT objects to detected components

$$AJI = \frac{\sum_{i=1}^{L} |G_i \cap S_k^*(i)|}{\sum_{i=1}^{L} |G_i \cup S_k^*(i)| + \sum_{l \in U} |S_l|}$$

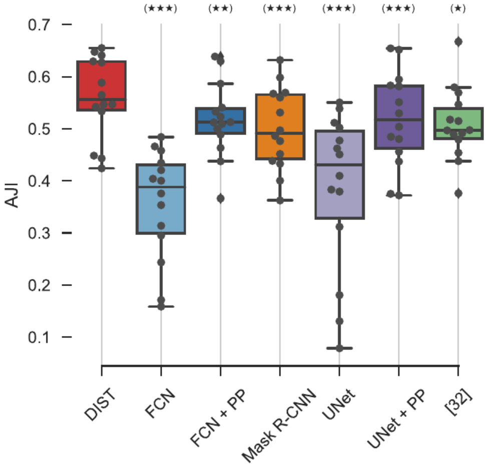

Example segmentations

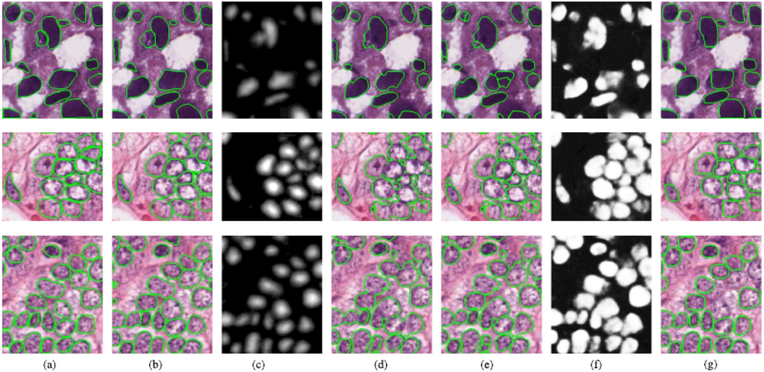

Fig. 8. Comparing segmentation results on cluttered cells. (a) Ground Truth. (b) DIST. (c) Distance regression output. (d) U-Net. (e) U-Net + PP. (f) U-Net probability map. (g) Mask R-CNN.

Conclusions

- CNNs give state of the art performance
- Give less importance to borders, more to objects
- Regressing distance better than pixel-level loss