Supplementary: Recovering the Imperfect: Cell Segmentation in the Presence of Dynamically Localized Proteins

Özgün Çiçek[*1], Yassine Marrakchi[*1,2], Enoch Boasiako Antwi[1,2,3], Barbara Di Ventura[1,2], and Thomas Brox[1,2]

[1] University of Freiburg, Germany
[2] Signalling Research Centres BIOSS and CIBSS, Freiburg, Germany
[3] Heidelberg Biosciences International Graduate School (HBIGS), Germany
{cicek,marrakch}@cs.uni-freiburg.de

1 Qualitative Results

1.1 Slight Signal Loss

Since the network has been trained on a very clean training data, even small shifts in the signal distribution might disturb its performance and cause failures as in Figure 1. Notice that in the first frame, we have two connected components wrongly building together the nucleus mask. As each cell can only have a single nucleus, we consider propagating one which is the most certain. The remaining connected component can be filtered with a simple post-processing step. We see that our method improves upon the initial segmentation masks significantly.

1.2 Extreme Signal Loss

The scenario in Figure 2 shows the full potential of our method in the case of extreme signal loss. Uncertainty estimation allows to rank segmentation masks and infer more reasonable masks for the less certain frames. In the last frame the network misses the nucleus completely in a very confident way. Since we set the uncertainty in such cases to infinity to ensure that it gets updated (biologically there is certainly a nucleus in each cell), we are able to recover the missing nucleus. For a biologist taking measurements from the results of our method is more reliable compared to the non-updated results.

1.3 No Signal Loss - Segmentation Failure

If the training data is not rich with all variations of the real task at hand, which is often the case in biomedical settings, networks can also fail at generalization. Our method is not only limited to refining failures due to signal loss, on the contrary, it is generic to all failures caused by the uncertainties inherited in the network or the data. Our predicted uncertainty estimations provide a good indicator for these scenarios and allows to recover them as we show in Figure 3.

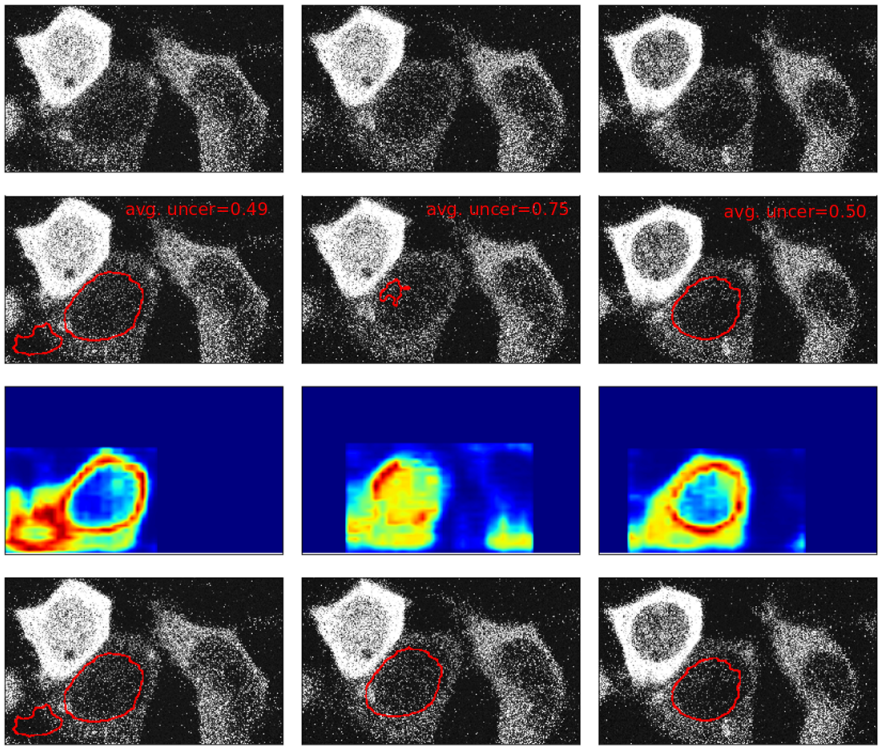

Fig. 1: Example of a nucleus getting slightly less visible due to increased noise at time t. Columns correspond to time points $t-1$, t and $t+1$, from left to right respectively. In the first row we show the raw images. In the second row we overlay MaskRCNN outputs with the input image and report the average uncertainty. Pixel-wise uncertainty maps are presented in the third row while updated masks are overlaid to the original images in the last row.

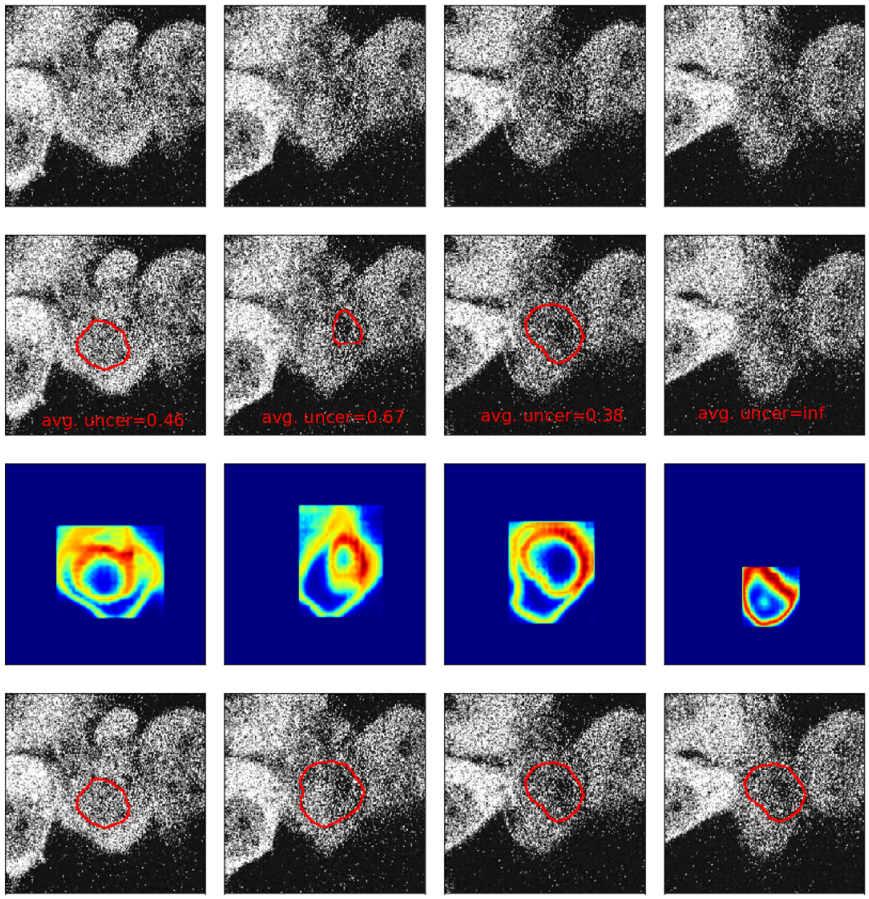

Fig. 2: Example of a nucleus becoming completely indistinguishable from the surrounding cytosol in all time points. Columns correspond to time points $t-1$, t, $t+1$ and $t+2$, from left to right respectively. In the first row we show the raw images. In the second row we overlay MaskRCNN outputs with the input image and report the average uncertainty. Pixel-wise uncertainty maps are presented in the third row while updated masks are overlaid to the original images in the last row.

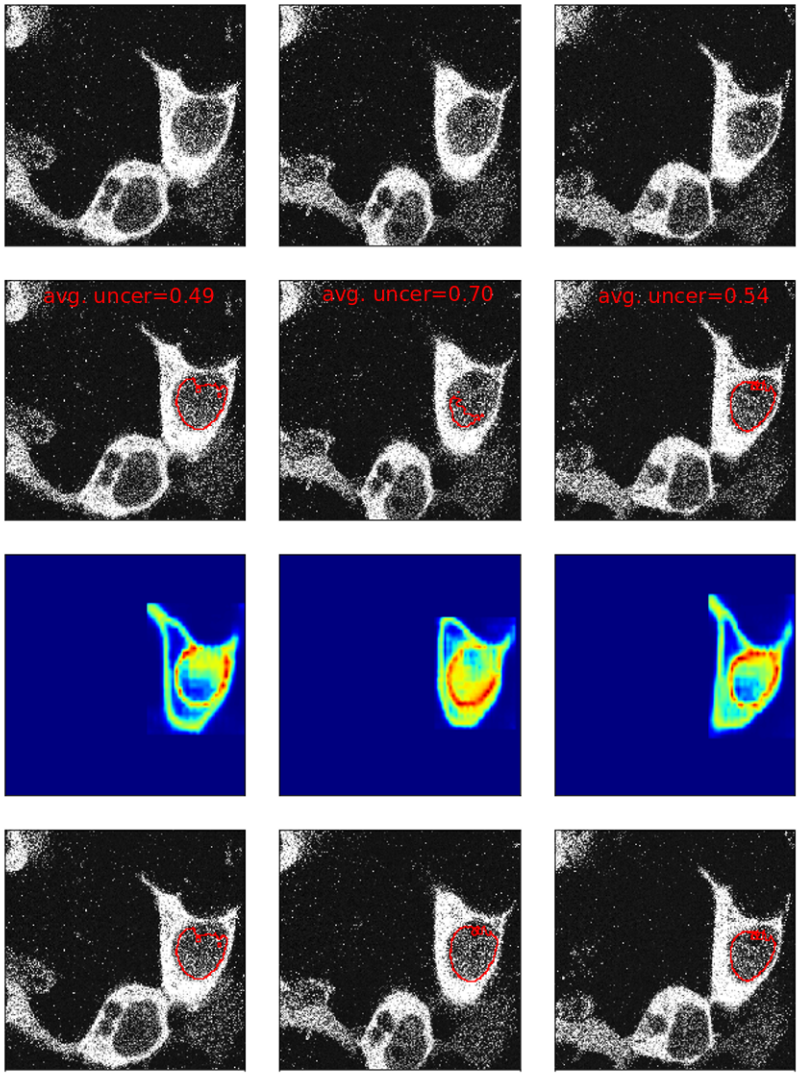

Fig. 3: Example of recovery from an under-segmentation failure at time stamp t. Columns correspond to time points $t-1$, t and $t+1$, from left to right respectively. In the first row we show the raw images. In the second row we overlay MaskRCNN outputs with the input image and report the average uncertainty. Pixel-wise uncertainty maps are presented in the third row while updated masks are overlaid to the original images in the last row.