

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

Özgün Çiçek<sup>\*1</sup>, Yassine Marrakchi<sup>\*1,2</sup>, Enoch Boasiako Antwi<sup>1,2,3</sup>,  
Barbara Di Ventura<sup>1,2</sup>, and Thomas Brox<sup>1,2</sup>

<sup>1</sup> University of Freiburg, Germany

<sup>2</sup> Signalling Research Centres BIOS and CIBSS, Freiburg, Germany

<sup>3</sup> Heidelberg Biosciences International Graduate School (HBIGS), Germany  
{cicek,marrakch}@cs.uni-freiburg.de

**Abstract.** Deploying off-the-shelf segmentation networks on biomedical data has become common practice, yet if structures of interest in an image sequence are visible only temporarily, existing frame-by-frame methods fail. In this paper, we provide a solution to segmentation of imperfect data through time based on temporal propagation and uncertainty estimation. We integrate uncertainty estimation into Mask R-CNN network and propagate motion-corrected segmentation masks from frames with low uncertainty to those frames with high uncertainty to handle temporary loss of signal for segmentation. We demonstrate the value of this approach over frame-by-frame segmentation and regular temporal propagation on data from human embryonic kidney (HEK293T) cells transiently transfected with a fluorescent protein that moves in and out of the nucleus over time. The method presented here will empower microscopic experiments aimed at understanding molecular and cellular function.

## 1 Introduction

In the past decades, it has become evident that proteins are very dynamic and their localization within the cell dictates which function they perform [3]. Cell biologists carry out time-lapse fluorescence microscopy experiments to study protein localization and unravel how this affects the protein’s function. Optogenetics can be used to control the localization of the protein of interest [18] to observe the effects of dynamic localization patterns. Take Figure 1 for an example: HEK293T cells expressing mCherry fused to the optogenetic tool LINuS [15] were observed over time to analyze the effect of blue light on the nuclear import of the protein. In the absence of blue light, the fusion protein localizes predominantly in the cytosol. Hence, the nucleus appears dark. However, in presence of blue light, the fusion protein enters the nucleus making it appear bright and difficult to distinguish from the cytosol. Giving light repeatedly creates oscillations of the protein in and out of the nucleus in time.

---

\* equal contribution




Fig. 1: (a) Exemplary time-lapse images of HEK293T cells depicting an oscillatory nuclear signal. (b) Oscillation at time  $t$  causing bad nuclei segmentation (up) and the corrected segmentation of it using our propagation method (down).

Learning-based methods, such as U-Net [19] and Mask R-CNN [6], succeed at segmenting structures in data with clearly visible patterns, but fail when the visibility deteriorates. When the signal in the nucleus is similar to that in the cytosol, nuclei segmentation from any network is not reliable. While in a single image it is not possible to improve these segmentations, past and future frames in a video provide additional information for refinement. An expert can play the video back and forth to infer the segmentation of ambiguous nuclei. An automated segmentation method, too, must (1) automatically identify critical frames and (2) propagate predictions from neighbouring frames.

In this paper, we address both challenges by (1) equipping Mask R-CNN with uncertainty estimation to identify erroneous predictions and (2) incorporating optical flow to improve the identified erroneous predictions by propagating certain predictions from neighboring frames. Doing so, we introduce the most recent uncertainty estimation methods in biomedical instance segmentation and solve a real task commonly experienced in signalling studies which is not yet addressed. So far, nuclear markers have been employed in the experiments so that the available automated segmentation tools can be used [4]. However, additional markers cause unreliable quantification since different proteins bleed-through and interfere with each other. They also limit the channel space needed for other proteins of interest. The presented method makes the use of a nuclear marker dispensable.

## 2 Related Work

We are the first to address instance segmentation of structures with an oscillatory fluorescent signal in biomedical videos. Our work is related to methods that benefit from temporal features in their design. Milan et al. [14] and Payer et al. [17] used Recurrent Neural Network (RNN) to aggregate temporal features. Paul et al. [16] incorporated temporal cues for segmentation via optical flow. Similarly, Jain et al. [9] propagated features instead of segmentation with flow

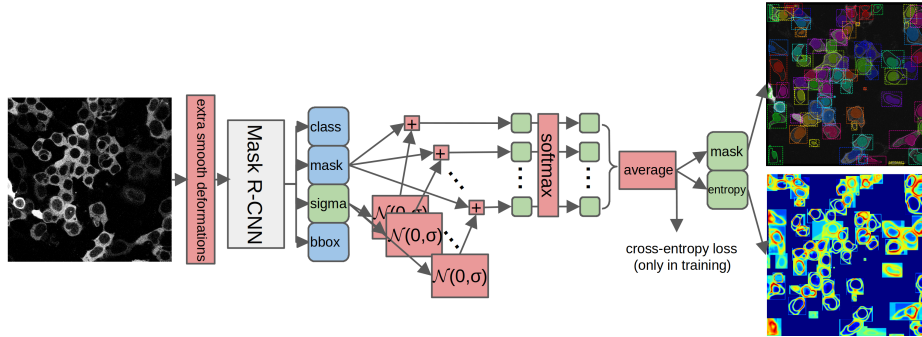


Fig. 2: Overview of Mask R-CNN with added data uncertainty. Changes to the original architecture are shown in red (operations) and green (outputs).

from key frames. A similar idea was applied to instance segmentation by Bertasius and Torresani [2]. Although these methods seem close, the task is different: (1) they heavily rely on dense annotation in time to learn interpolations explicitly while we cannot afford it due to high cost of expert annotation and difficulty in fine-grained annotation of imperfect frames and (2) they benchmark only on visible objects, while we are solely interested in objects with limited visibility.

### 3 Methods

#### 3.1 Instance-Aware Segmentation with Uncertainty Estimation

We base our model on Mask R-CNN [6] equipped with elastic deformations of U-Net [19] to create additional biomedically plausible images on-the-fly for better generalization. We incorporate uncertainty estimation in our Mask R-CNN architecture to detect erroneous predictions; see Figure 2. We consider data uncertainty (aleatoric), model uncertainty (epistemic), and their combination.

**Data uncertainty.** For data uncertainty, we use the modified cross-entropy loss [10] in the mask branch of Mask R-CNN. This models the data uncertainty as the learned noise scale from the data. To learn both the class scores and their noise scale, the negative expected log-likelihood is minimized for pixel  $i$  as:

$$L(s) = - \sum_i \log \left[ \frac{1}{T} \sum_t \exp(\hat{s}_{i,t,c'} - \log \sum_c \exp(\hat{s}_{i,t,c})) \right], \quad (1)$$

where  $c'$  is the correct class among all classes ( $c$ ) and  $\hat{s}$  are predicted logits corrupted by Gaussian noise with standard deviation  $\sigma$ .  $\sigma$  is also learned by the network alongside with the logits. Figure 2 shows where this loss is embedded in the network. At test time, the uncertainty is computed as the entropy of the class pseudo-probabilities for each pixel as:  $\mathcal{U}(p) = - \sum_c p(c) \log p(c)$ .

**Model uncertainty.** Model uncertainty requires sampling from the model. We experimented with the latest sampling strategies. **Dropout** is one of the common techniques to sample from networks [10]. **Ensemble** replaces the sampling

**Algorithm 1: Uncertainty-Based Nuclei Propagation**


---

**Input:** instance segmentation masks  $\mathcal{S}$ , uncertainty maps  $\mathcal{U}$ , motion per consecutive frame pairs  $\mathcal{M}$ , hard threshold  $\theta$ , relative thresholds:  $\alpha, \beta$

**Output:** tracks

```

1 tracks ← IoU-BasedCellTracker( $\mathcal{S}$ );
2 for track  $t$  in tracks do
3   for frame  $f$  in track  $t$  do
4      $\bar{u}_{t,f} \leftarrow \text{mean\_over\_nuclei}(\mathcal{U}_{t,f}, \mathcal{S})$ ;
5     order ← argsort( $\bar{u}_t$ );
6     for frame id  $f$  in order do
7       if  $\bar{u}_{t,f} \geq \theta$  and  $\beta \times \bar{u}_{t,f} \geq \bar{u}_{t,f-1}$  and  $\beta \times \bar{u}_{t,f} \geq \bar{u}_{t,f+1}$  then
8         maskprev ← warp( $\mathcal{S}_{t,f}, \mathcal{S}_{t,f-1}, \mathcal{M}$ );
9         masknext ← warp( $\mathcal{S}_{t,f}, \mathcal{S}_{t,f+1}, \mathcal{M}$ );
10         $\mathcal{S}_{t,f} \leftarrow \text{union}(\text{mask}_{\text{prev}}, \text{mask}_{\text{next}})$ ;
11      else if  $\bar{u}_{t,f} \geq \theta$  and  $\alpha \times \bar{u}_{t,f} \geq \bar{u}_{t,f-1}$  then
12         $\mathcal{S}_{t,f} \leftarrow \text{warp}(\mathcal{S}_{t,f}, \mathcal{S}_{t,f-1}, \mathcal{M})$ ;
13      else if  $\bar{u}_{t,f} \geq \theta$  and  $\alpha \times \bar{u}_{t,f} \geq \bar{u}_{t,f+1}$  then
14         $\mathcal{S}_{t,f} \leftarrow \text{warp}(\mathcal{S}_{t,f}, \mathcal{S}_{t,f+1}, \mathcal{M})$ ;

```

---

by dropout by individually trained networks [11]. **SGDR ensemble** replaces the individual trained networks by the pre-converged snapshots of the same model [7]. Pre-converged models are obtained at the end of each SGDR (SGD with warm restarts) [12] cycle. **WTA** (Winner-Takes-All) [8] aims for training a single network with multiple heads where at each iteration only the head with the best prediction gets penalized. **EWTA** (Evolving WTA) [13] is a variant of WTA which improves the trade-off between diversity and consistency.

### 3.2 Uncertainty-Based Nuclei Propagation

We base the nuclei propagation on the cell tracker<sup>4</sup> by Ronneberger et al. [19]. This overlapping-based tracker is well suited for the data at hand, which does not show large motion over time. Despite its simplicity, it is one of the top performing cell trackers in the challenge. After computing the overlapping-based tracks, we go over all the frames of each track with ascending average nuclei uncertainty order (without nuclei, uncertainty is set to infinity) and update the identified uncertain predictions by propagating masks from more certain neighbours as depicted in Figure 1. A prediction is marked uncertain if its average uncertainty is higher than a threshold  $\theta$ . Neighbours are more certain if they meet a relative threshold of  $\alpha$  in case of a single-side propagation and  $\beta$  in case of a two-sided propagation as detailed in Algorithm 1. This resembles the procedure that experts follow to find non-visible nuclei by sliding over time.

<sup>4</sup> <http://celltrackingchallenge.net/>

### 3.3 Motion Estimation for Biomedical Videos

Propagation of predictions over time requires motion estimation between frames to warp the certain predictions onto the less certain ones. One simple way is to warp by the shift and scaling parameters computed between the not yet updated and neighbouring nuclei predictions. This approach assumes the shape of the nuclei does not change over time; however, slight deformations can occur. Optical flow can provide fine-grained motion. Recent optical flow methods for natural images are networks [5] trained on synthetic datasets. These methods perform well on real images, but their performance deteriorates as the gap between real and synthetic images grows. Our data is very different from existing synthetic datasets and no synthetic data exists for biomedical data. Therefore, with the prior knowledge that expected flow in our videos can be well explained by smooth deformations, and being able to generate them on the fly via deformation augmentations of U-Net, we explicitly train a network to predict these deformations as explored by Sokooti et al. [20]. We warp an image ( $I_{t+1}$ ) backward in time with randomly generated smooth deformations ( $f_{t \rightarrow t+1}$ ) to obtain the previous image ( $I_t$ ). Then we train a FlowNet on the on-the-fly-generated image pairs and ground-truth flows. Even though the fluorescent signal in the nuclei regularly disappears making the optical flow challenging to estimate, the motion of the cytosol helps infer the motion of the nuclei.

## 4 Experiments

**Implementation details.** We based our implementation on the Mask R-CNN by Abdulla [1]. We incorporated publicly available elastic deformations<sup>5</sup> for U-Net [19] as additional augmentation. We used ResNet50-FPN as backbone. We used softmax cross entropy for the mask head with 3 classes (background, cytosol and nuclei). We trained the networks from scratch for 12k epochs except the SGDR ensemble, which was trained for 15k with 3k cycles. We used the last pre-converged models. The WTA and EWTA models were trained for 11k and then merged with a network trained for 4k. We used 4 as ensemble size for all methods. Ensembling was performed over bounding boxes. For nuclei propagation, we used  $\theta = 0.5$ ,  $\alpha = 0.7$  and  $\beta = 0.85$ . For Mask R-CNN and FlowNet (variant C) training we used deformations with 3 and 10 control points, respectively and deformation magnitude of 10. Please see the original papers for more details. Our code and exemplar dataset is publicly available<sup>6</sup>.

**Data and annotation.** While the challenge we address is common in signalling studies, there is no public dataset for this purpose. Our data was generated after 24 hours of transient transfection of human embryonic kidney (HEK293T) cells with a construct expressing the fusion protein mCherry-LINuS with a ZEISS LSM780 confocal microscope. MaskRCNN was trained on 82 cropped images from original 1800x1800 images with pixel size 0.11x0.11 and 6 cells on average

<sup>5</sup> [https://github.com/fcalvet/image\\_tools](https://github.com/fcalvet/image_tools)

<sup>6</sup> <https://lmb.informatik.uni-freiburg.de/Publications/2020/CMB20/>

in each and tested on 2 full-sized images with 73 cells in total to validate instance segmentation and uncertainty estimation. We evaluated our propagation algorithm on an unseen video with 35 full-sized frames with randomly selected 117 cells annotated by experts. In cases where the nuclei appeared ambiguous experts interpolated the annotations from past and future frames.

#### 4.1 Instance Segmentation and Uncertainty Estimation Evaluation

We evaluated our instance segmentation by mean average precision (mAP). In Table 1 we refer to standard mAP as mAP (sm) since it is based on softmax scores. To further evaluate the quality of the uncertainty estimation, we replaced the softmax scores by average entropy over the cell and re-computed mAP (mAP (ent)). This simulates our approach as we rely on averaged uncertainties in our nuclei propagation. This measure shows how reliable the predicted uncertainties are at ranking the prediction quality in the precision-recall curves, which are used commonly to evaluate uncertainty estimation for classification tasks. In Table 1 we see that WTA merged is the best at 0.5 IoU (Intersection over Union) threshold and at 0.75 threshold ensemble is the best. In the rest of our experiments we used the WTA merged with data uncertainty since it is computationally more efficient than ensemble.

Table 1: Quantitative evaluation for instance segmentation and uncertainty estimation in mAP (@0.5/@0.75 IoU).

	model uncertainty		combined uncertainty	
	mAP (sm)	mAP (ent)	mAP (sm)	mAP (ent)
Single	0.77/0.48	0.80/0.49	0.74/0.60	0.83/0.69
Dropout	0.74/0.61	0.78/0.65	0.77/0.61	0.83/0.67
Ensemble	0.82/ <b>0.64</b>	0.78/0.61	0.78/0.63	0.83/ <b>0.70</b>
SGDR Ensemble	0.75/0.54	0.72/0.51	0.71/0.49	0.63/0.44
WTA Merged	0.74/0.47	0.82/0.49	<b>0.83</b> /0.56	<b>0.85</b> /0.64
EWTA Merged	0.64/0.51	0.73/0.58	0.80/ <b>0.64</b>	0.77/0.59

#### 4.2 Nuclei Propagation Evaluation

We report the mean IoU of the nuclei segmentation for all experiments in Table 2. We use *interpolated* for nuclei completely missed before our improvement, *updated* for the nuclei that were segmented but our method decided to improve, *non-updated* for the nuclei that were chosen not to be improved by our method and *all* for all the mentioned cases. The first row shows the results before any propagation algorithm. We present 3 variants of motion used in propagation (column: *warped with*). We also explored using the certainties as pixel-wise weights

in fusing segmentations from candidate neighbors and computed the weighted average to find the final mask. To isolate the gain by our uncertainty-based error detection, we created a baseline which is identical to our mean-flow variant, but performs the propagation on all the nuclei over a track. The significant improvement obtained by all our variants, especially flow variants without fusion, shows that our method can effectively improve erroneous nuclei predictions. The baseline that propagates to all frames (*all*) had a lower performance on the *non-updated* frames, showing that propagation independent of an uncertainty measure harms nuclei with high confidence (approx. 90% of all nuclei). For a better understanding of the challenge we tackle and the usefulness of our proposed method, we provide more qualitative results in the supplementary material.

Table 2: Quantitative evaluation of uncertainty-based propagation in mean IoU for all/updated/interpolated/non-updated nuclei with respective nuclei counts.

update	warp with	mask fusion	all(117)	updated(51)	extrapolated(11)	non-updated(55)
none	none	no	0.62	0.55	0.00	<b>0.80</b>
uncertain	shift+scale	no	0.71	0.68	0.39	<b>0.80</b>
uncertain	mean nuclei flow	no	<b>0.73</b>	0.71	<b>0.45</b>	<b>0.80</b>
all	mean nuclei flow	no	0.69	0.70	0.40	0.74
uncertain	pixel-wise flow	no	<b>0.73</b>	<b>0.72</b>	0.44	<b>0.80</b>
uncertain	pixel-wise flow	yes	0.72	0.70	0.40	<b>0.80</b>

## 5 Conclusion

We addressed automated segmentation of image sequences, which cannot be analyzed frame-by-frame due to temporary uncertainties causing errors in predictions. First, we estimate uncertainty from Mask R-CNN to identify unreliable predictions. Second, we improve the less reliable predictions by propagating the more certain ones from neighbouring frames. We evaluated our method on HEK293T cells expressing a protein that oscillates in and out of the nucleus over time making the nucleus invisible temporarily. Our method improves nuclei segmentation over several baselines while keeping the cytosol and background segmentation untouched. We believe that our method will facilitate further time-series analysis for quantitative biology to understand the effect that the dynamic localization of the protein has on the cell without additional markers.

**Acknowledgments.** This project was funded by the German Research Foundation (DFG) and the German Ministry of Education and Science (BMBF). Gefördert durch die Deutsche Forschungsgemeinschaft (DFG) im Rahmen der Exzellenzstrategie des Bundes und der Länder – EXC-2189 – Projektnummer 390939984 und durch das Bundesministerium für Bildung und Forschung (BMBF) Projektnummer 01IS18042B und 031L0079.

## References

1. Abdulla, W.: Mask R-CNN for object detection and instance segmentation on keras and tensorflow. [https://github.com/matterport/Mask\\_RCNN](https://github.com/matterport/Mask_RCNN) (2017)
2. Bertasius, G., Torresani, L.: Classifying, segmenting, and tracking object instances in video with mask propagation. Tech. Rep. 1912.04573, arXiv (2019)
3. Brami-cherrier, K., Gervasi, N., Arsenieva, D., Boutterin, M.C., Ortega, A., Seantier, B., Gasmı, L., Bouceba, T., Noble, M., Girault, J.A., Arold, S.: Mechanisms of site-specific functions of focal adhesion kinase. *Biophysical Journal* **104**, 609 (2013)
4. Chen, S.Y., Osimiri, L.C., Chevalier, M., Bugaj, L.J., Ng, A.H., Stewart-Ornstein, J., Neves, L.T., El-Samad, H.: Optogenetic control reveals differential promoter interpretation of transcription factor nuclear translocation dynamics. *bioRxiv* (2019)
5. Dosovitskiy, A., Fischer, P., Ilg, E., Häusser, P., Hazırbağ, C., Golkov, V., v.d. Smagt, P., Cremers, D., Brox, T.: FlowNet: Learning optical flow with convolutional networks. In: *ICCV* (2015)
6. He, K., Gkioxari, G., Dollár, P., Girshick, R.B.: Mask R-CNN. In: *ICCV* (2017)
7. Huang, G., Li, Y., Pleiss, G.: Snapshot ensembles: Train 1, get M for free. In: *ICLR* (2017)
8. Ilg, E., Çiçek, Ö., Galesso, S., Klein, A., Makansi, O., Hutter, F., Brox, T.: Uncertainty estimates and multi-hypotheses networks for optical flow. In: *ECCV* (2018)
9. Jain, S., Wang, X., Gonzalez, J.E.: Accel: A corrective fusion network for efficient semantic segmentation on video. In: *CVPR* (2019)
10. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? In: *NIPS* (2017)
11. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: *NIPS Workshop* (2016)
12. Loshchilov, I., Hutter, F.: SGDR: Stochastic gradient descent with warm restarts. In: *ICLR* (2017)
13. Makansi, O., Ilg, E., Çiçek, Ö., Brox, T.: Overcoming limitations of mixture density networks: A sampling and fitting framework for multimodal future prediction. In: *CVPR* (2019)
14. Milan, A., Rezatofghi, S.H., Dick, A., Reid, I., Schindler, K.: Online multi-target tracking using recurrent neural networks. In: *AAAI* (2017)
15. Niopek, D., Benziger, D., Roensch, J., Draebing, T., Wehler, P., Eils, R., Di Ventura, B.: Engineering light-inducible nuclear localization signals for precise spatiotemporal control of protein dynamics in living cells. *Nature Communications* **5**, 4404 (July 2014)
16. Paul, M., Mayer, C., Gool, L.V., Timofte, R.: Efficient video semantic segmentation with labels propagation and refinement. Tech. Rep. 1912.11844, arXiv (2019)
17. Payer, C., Štern, D., Neff, T., Bischof, H., Urschler, M.: Instance segmentation and tracking with cosine embeddings and recurrent hourglass networks. In: *MICCAI*. pp. 3–11. LNCS (2018)
18. Repina, N.A., Rosenbloom, A., Mukherjee, A., Schaffer, D.V., Kane, R.S.: At light speed: Advances in optogenetic systems for regulating cell signaling and behavior. *Annual Review of Chemical and Biomolecular Engineering* **8**(1), 13–39 (2017)
19. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: *MICCAI* (2015)
20. Sokooti, H., de Vos, B., Berendsen, F., Lelieveldt, B.P.F., Išgum, I., Staring, M.: Nonrigid image registration using multi-scale 3d convolutional neural networks. In: *MICCAI* (2017)