

RESEARCH ARTICLE

Cell segmentation and tracking using CNN-based distance predictions and a graph-based matching strategy

Tim Scherr^{1*}, Katharina Löffler^{1,2}, Moritz Böhland¹, Ralf Mikut¹

1 Institute for Automation and Applied Informatics, Karlsruhe Institute of Technology, Eggenstein-Leopoldshafen, Germany, **2** Institute of Biological and Chemical Systems - Biological Information Processing, Karlsruhe Institute of Technology, Eggenstein-Leopoldshafen, Germany

* tim.scherr@kit.edu

OPEN ACCESS

Citation: Scherr T, Löffler K, Böhland M, Mikut R (2020) Cell segmentation and tracking using CNN-based distance predictions and a graph-based matching strategy. PLoS ONE 15(12): e0243219. <https://doi.org/10.1371/journal.pone.0243219>

Editor: Konradin Metze, University of Campinas, BRAZIL

Received: July 21, 2020

Accepted: November 17, 2020

Published: December 8, 2020

Copyright: © 2020 Scherr et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the paper and its [Supporting information](#) files.

Funding: We are grateful for funding by the Helmholtz Association in the programs BioInterfaces in Technology and Medicine (TS, RM) and the Helmholtz Information & Data Science School for Health (KL, RM). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. We acknowledge support by the KIT-

Abstract

The accurate segmentation and tracking of cells in microscopy image sequences is an important task in biomedical research, e.g., for studying the development of tissues, organs or entire organisms. However, the segmentation of touching cells in images with a low signal-to-noise-ratio is still a challenging problem. In this paper, we present a method for the segmentation of touching cells in microscopy images. By using a novel representation of cell borders, inspired by distance maps, our method is capable to utilize not only touching cells but also close cells in the training process. Furthermore, this representation is notably robust to annotation errors and shows promising results for the segmentation of microscopy images containing in the training data underrepresented or not included cell types. For the prediction of the proposed neighbor distances, an adapted U-Net convolutional neural network (CNN) with two decoder paths is used. In addition, we adapt a graph-based cell tracking algorithm to evaluate our proposed method on the task of cell tracking. The adapted tracking algorithm includes a movement estimation in the cost function to re-link tracks with missing segmentation masks over a short sequence of frames. Our combined tracking by detection method has proven its potential in the IEEE ISBI 2020 Cell Tracking Challenge (<http://celltrackingchallenge.net/>) where we achieved as team KIT-Sch-GE multiple top three rankings including two top performances using a single segmentation model for the diverse data sets.

Introduction

State-of-the-art microscopy imaging techniques such as light-sheet fluorescence microscopy imaging enable to investigate cell dynamics with single-cell resolution [1, 2]. This allows to study cell migration and proliferation in tissue development and organ formation at early embryonic stages. Establishing the required complete lineage of each cell, however, requires a virtually error-free segmentation and tracking of individual cells over time [2, 3]. A manual data analysis is unfeasible, due to the large amount of data acquired with modern imaging

Publication Fund of the Karlsruhe Institute of Technology.

Competing interests: The authors have declared that no competing interests exist.

techniques. In addition, low-resolution objects are very difficult to detect even for human experts. Deep learning-based cell segmentation methods have proven to outperform traditional methods even on very diverse 2D data sets [4]. However, state-of-the-art cell tracking methods often still need a time-consuming manual cell track curation, e.g., using EmbryoMiner [5] or the Massive Multi-view Tracker (MaMuT) [6]. Especially for low signal-to-noise ratio and 3D data, further method development is required for both cell segmentation and cell tracking [7].

Traditional segmentation methods, such as TWANG for the segmentation of roundish objects [8], are often designed for a specific application. These methods commonly consist of sophisticated combinations of pre-processing filters, e.g., Gaussian or median filters, and segmentation operations, e.g., a region adaptive thresholding followed by a watershed transform [9]. To reach a reasonable segmentation quality, such traditional methods need to be carefully adapted to the cell type and imaging conditions. Therefore, expert knowledge is needed. In contrast, deep learning-based segmentation methods shift the expert knowledge needed to the model design and to the training process. Thus, less expert knowledge is needed for the application of a trained model and to fine-tune the post-processing which is often kept very simple. A review of cell segmentation methods is provided in [10].

To improve the generalization ability of a trained deep learning model, a preferably diverse and large annotated data set is needed. This fact is especially problematic when dealing with touching cells since this case is usually underrepresented in training data sets. Therefore, models for cell boundary or border prediction (see Fig 1) are often not able to handle touching cells well. The result are merged cells, due to gaps in predicted cell boundaries and borders between touching cells [11, 12]. To overcome this problem, several approaches have been proposed. In [11], models are trained to predict adapted thicker borders and smaller cells, which can decrease the amount of merged cells. [12] utilizes new gap and touching classes with J regularization. [13] combines distance transforms for single cell nuclei with discrete boundaries. A center vector encoding which is aimed to be more robust to label inconsistencies is proposed in [14], whereas in [15], horizontal and vertical gradient maps are used. To improve the generalization ability of a model for cell types with only few or no annotated images, a generative adversarial network-based image style transfer to generate augmented training samples of that cell types has been used in [16]. An advantage of border-based approaches is that a deep learning model is enforced to focus on touching cells that are underrepresented in the training data. However, border-based approaches still have the shortcoming that only touching cells can be used to train the border prediction.

Although deep learning methods have been successfully applied to multi object tracking on natural images [17, 18], there are only few deep learning approaches for cell tracking [19, 20]. In [19], cells are simultaneously segmented and tracked by combining a recurrent hourglass network with a pixel-wise metric embedding learning. [20] proposes a particle-filter-based motion model in combination with a CNN-based observation model. However, cell tracking is still dominated by traditional tracking approaches [7, 21]. One reason is the lack of high quality annotations as provided in natural image tracking benchmarks [22–24]. Thus, training data are often not available. Another aspect that complicates the task of cell tracking are cell death and division events, which do not occur in natural image tracking data. Therefore, traditional tracking algorithms with comparably few parameters and explicit modeling of cell division events still dominate cell tracking benchmarks [7]. The comparison of cell tracking algorithms in [7] shows that the majority of tracking approaches uses an adapted version of nearest neighbors, a graph-based linking or multi hypothesis tracking. In [21], the Viterbi algorithm is applied to track cells. A joint model for segmentation and tracking is proposed in [25] where model parameters are learned based on Bayes risk minimization.

In this paper, we propose a novel representation of cell borders, the neighbor distances, to solve the challenging problem of segmenting touching cells of various types in the absence of large training data sets. Thus far, problems of border prediction approaches are the sensitivity to annotation inconsistencies, and that only touching cells provide border information in the training. The neighbor distances are aimed to be less sensitive to annotation inconsistencies, and enable to learn also from close cells. This additional information in the training process results in a more robust border prediction. Similar to [13], we combine our border predictions with cell distances to further prevent the erroneous merging of close cells. However, in contrast to [13], we remove the bottlenecks of non-robust discrete boundaries and of the feature fusion layers. This results in a simplified architecture and training process and in less merged cells. For the cell tracking, we adapt a coupled minimum-cost flow algorithm to include an object movement estimation. In addition, our formulation is able to link fragmented tracks due to missing segmentation masks in a short sequence of frames. The remainder of this paper covers the methodology we use to detect and segment cells and the subsequent cell tracking. In the results section, we demonstrate the quality of our introduced method on data from the Cell Tracking Challenge [7, 9].

Materials and methods

Cell segmentation using CNN-based distance predictions

For cell segmentation, we train a deep learning model to predict cells and cell borders, followed by a post-processing with a seed extraction and a seed-based watershed segmentation. A key for the successful application of supervised deep learning methods in the absence of large training data sets is to introduce representations that allow to use as much information as possible. Thus, instead of discrete cell boundary (Fig 1c) and cell border representations (Fig 1d), we combine cell distances (Fig 1e, [13]) with novel neighbor distances (Fig 1f). These representations allow incorporating the regional information not only from touching cells but also from close cells resulting in more robust deep learning models. A segmentation network based on the U-Net architecture [27], modified similar to [13], is utilized as the backbone of the method. An overview of the proposed method provides Fig 2.

Cell distances and neighbor distances. The cell distances, as shown in Fig 1e, are generated from ground truth data by computing the Euclidean distance transform for each cell independently. Adjacent cells are treated as background in this step and the distance transform is normalized into the range [0, 1]. Thus, each pixel of a cell represents the normalized distance to the nearest pixel not belonging to this cell. The cell distance prediction alone is sufficient to obtain seeds for the post-processing. However, a precondition is that the CNN has learned to

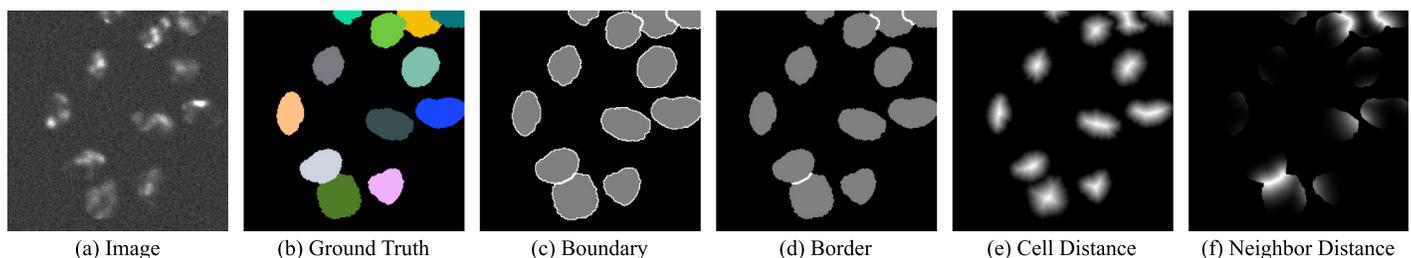


Fig 1. Training data representations for the training of deep learning models. Image (a) and ground truth (b) show a crop of the simulated Cell Tracking Challenge data set Fluo-N2DH-SIM+ [7, 9]. Generated boundaries (c) and borders (d) can be used to split touching cells. Many training data sets contain only few touching cells resulting in few training samples for borders and boundaries between cells. The combination of cell distances (e) with neighbor distances (f) is aimed to solve this problem since models can also learn from close cells.

<https://doi.org/10.1371/journal.pone.0243219.g001>

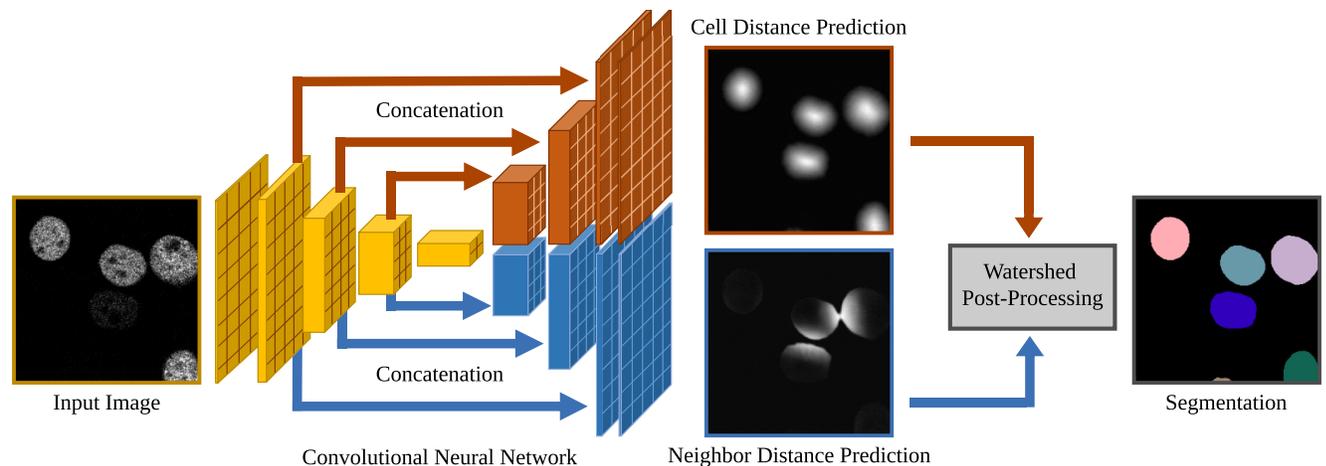


Fig 2. Overview of the proposed segmentation method using distance predictions (adapted from [13]). The CNN consists of a single encoder that is connected to both decoder paths. The network is trained to predict cell distances and neighbor distances that are used for the watershed-based post-processing. The input image shows a crop of the Cell Tracking Challenge data set Fluo-N2DH-GOWT1 [7, 9].

<https://doi.org/10.1371/journal.pone.0243219.g002>

deal with cell distances of touching cells. By combining cell distances with the novel neighbor distances the erroneous merging of touching cells is prevented.

Fig 3 shows the generation of the neighbor distances in which each pixel of a cell represents the inverse normalized distance to the nearest pixel of the closest neighboring cell. Therefore, a background-foreground conversion step is applied for each cell independently (Fig 3b) and the Euclidean distance transform (Fig 3c) is calculated. The distance transform is cut to the cell size and normalized (Fig 3d) followed by an inversion (Fig 3e). The normalization to the range [0, 1] is required to suppress neighbor distances for cells without close neighbors. To further reduce the erroneously merging of cells, gaps between close cells are closed by applying a grayscale closing (Fig 3g). Finally, to get a steeper decline within cells, a scaling is applied by taking the closed neighbor distances (Fig 3g) to the power of three (Fig 3h). This confines the neighbor distances to the outer cell area and therefore eases the seed extraction in the post-processing. An advantage of the neighbor distances is that they also provide information in the training process when cells are close but do not touch. This can be seen in Fig 3h (bottom right cell and bottom left cell) and is especially advantageous for training data sets with few touching cells providing only little border information in the training process.

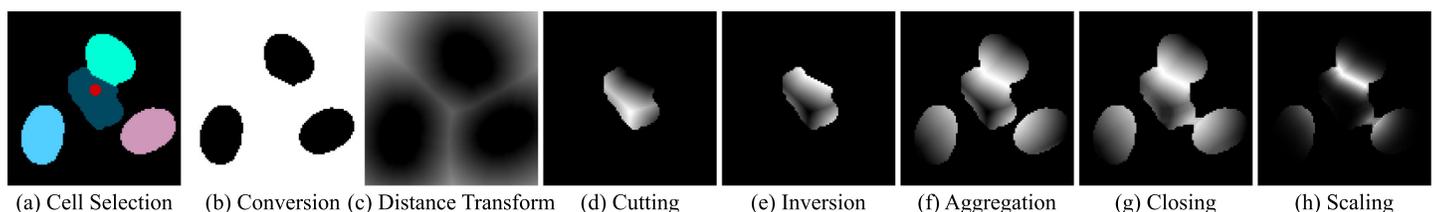


Fig 3. Main steps of the neighbor distance creation. After the automated selection of a cell (a), indicated with red, the selected cell and the background are converted to foreground (white in b) while the other cells are converted to background (black in b). Then, the distance transform is calculated (c), cut to the cell region and normalized (d). After inversion (e), the steps are repeated for the remaining cells (f). Finally, the grayscale closed neighbor distances (g) are scaled (h). Shown is a crop of the Broad Bioimage Benchmark Collection data set BBBC039v1 [26].

<https://doi.org/10.1371/journal.pone.0243219.g003>

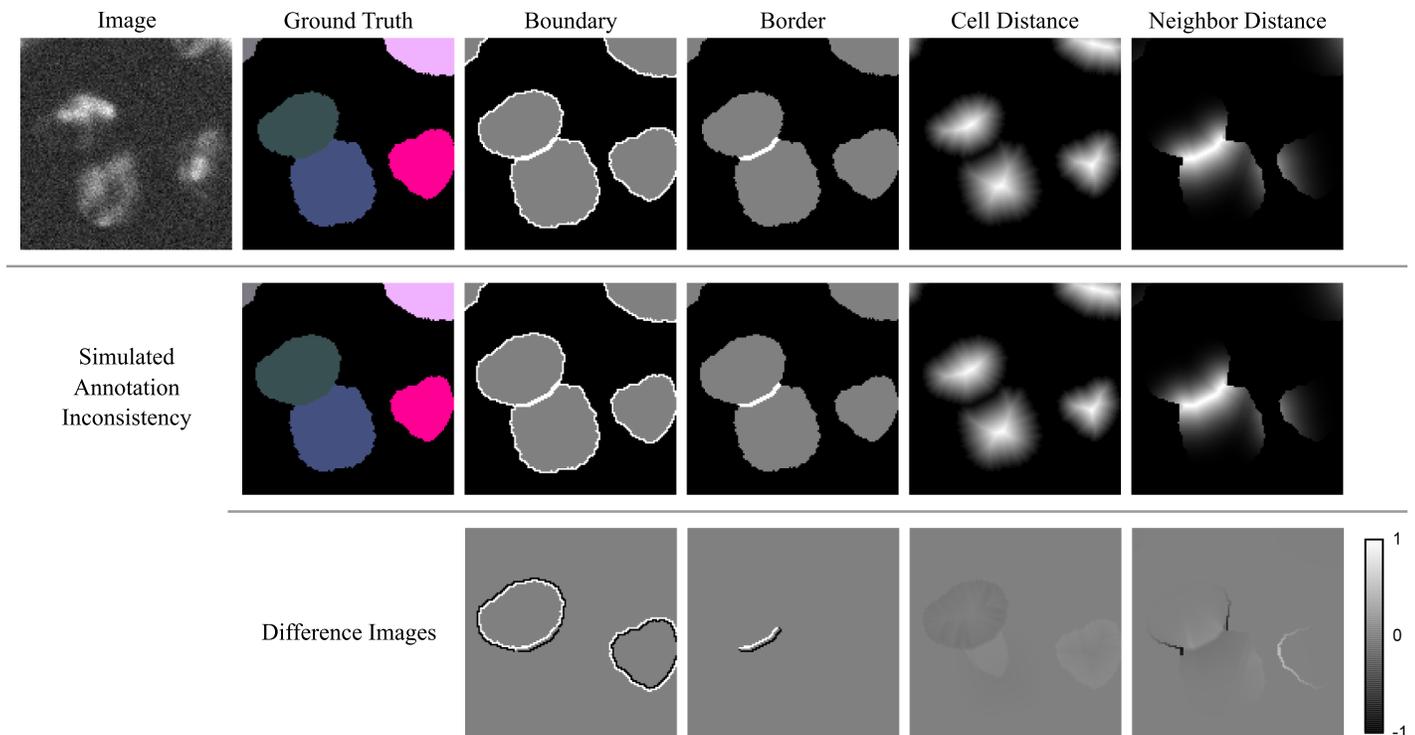


Fig 4. Robustness of training data representations to annotation inconsistencies. Small changes in the ground truth, simulated with morphological erosions and dilations, result in different boundaries and borders (first and second row). The difference images between the first row and the second row show that the changes for the distance labels are smoother. Shown is a crop of the Cell Tracking Challenge data set Fluo-N2DH-SIM+ [7, 9].

<https://doi.org/10.1371/journal.pone.0243219.g004>

Robustness of neighbor distances to annotation inconsistencies. Fig 4 shows that the neighbor distances are more robust to annotation inconsistencies than boundaries and borders, i.e., a cell was morphologically eroded and another cell dilated resulting in masks that only differ in single pixels. The location of the discrete boundaries and borders change, meaning that a prediction of the initial border is considered incorrect and penalized in the training. This makes it difficult to train models well on small data sets. In contrast, the proposed continuous neighbor distance shows a smooth change. Therefore, the influence of annotation inconsistencies on the training process is reduced resulting in a more robust training.

Architecture. The architecture is based on the U-Net architecture [27]. Instead of a single decoder path, two parallel decoder paths are used allowing each path to focus on features related to the desired output. In addition, the feature detection in the shared encoder branch of the network is trained using backpropagated information from both decoder branches. The maximum pooling layers are replaced with 2D convolutional layers with stride 2 and kernel size 3. Additionally, batch normalization layers are added. The number of feature maps is doubled from 64 feature maps to a maximum of 1024 in the encoder path and halved in each decoder path correspondingly. To avoid the need of cropping before concatenation, zero padding is applied in the convolutional layers to keep the feature map size consistent. The rectified linear unit activation function is used within the network and a linear activation for the output layers. Fig 2 provides an overview of the architecture, i.e., convolutional and downsampling layers are summarized into blocks. A more detailed description of the architecture is provided in S1 Fig.

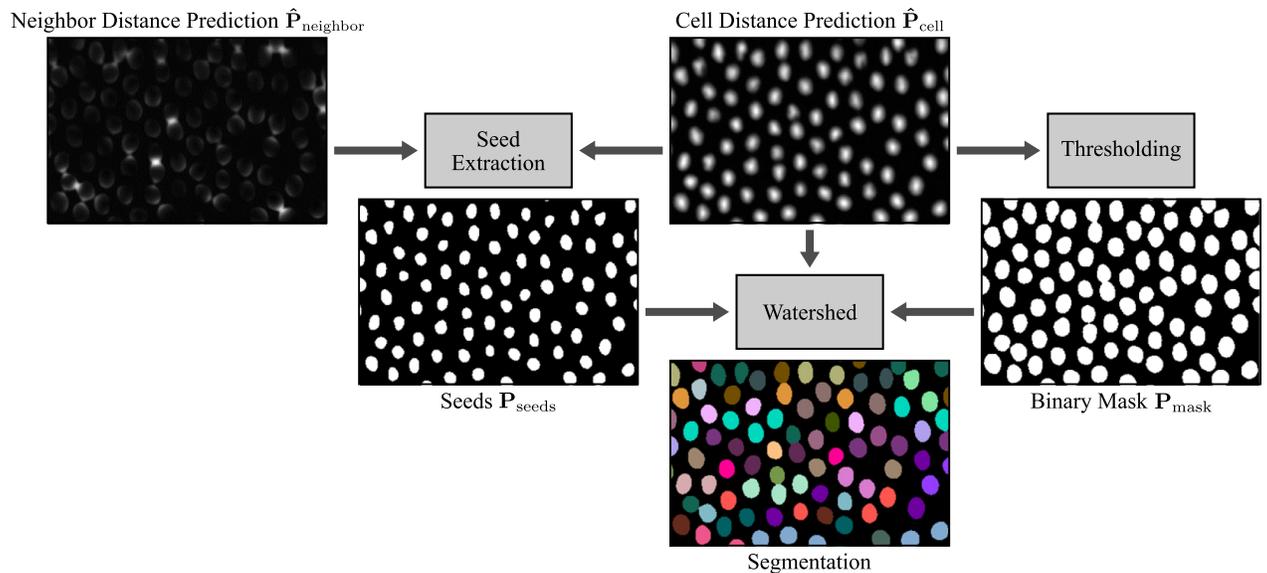


Fig 5. Overview of the watershed post-processing for segmentation. The post-processing consists of a threshold-based seed extraction and mask creation, and a watershed. The predictions show a 2D crop of the Cell Tracking Challenge data set Fluo-N3DL-TRIC [7, 9].

<https://doi.org/10.1371/journal.pone.0243219.g005>

Watershed post-processing. Fig 5 shows the main steps of the post-processing. The cell distance prediction P_{cell} and the neighbor distance prediction $P_{neighbor}$ are smoothed to avoid the erroneous splitting of cells in the seed extraction step:

$$\hat{P}_{cell} = P_{cell} * G(\sigma) , \tag{1}$$

$$\hat{P}_{neighbor} = P_{neighbor} * G(\sigma) , \tag{2}$$

with $G(\sigma)$ representing a Gaussian kernel with standard deviation σ and $*$ being the convolution operator. Then, the region to flood P_{mask} with a seed-based watershed is extracted from the smoothed cell distance prediction by applying a threshold ϱ_{mask} :

$$P_{mask} = \hat{P}_{cell} > \varrho_{mask} . \tag{3}$$

To obtain the seeds, the smoothed and squared neighbor distance prediction is subtracted from the cell prediction and the threshold ϱ_{seed} is applied:

$$P_{seeds} = (\hat{P}_{cell} - \hat{P}_{neighbor}^2) > \varrho_{seed} . \tag{4}$$

Depending on the cell size, the squaring can be omitted or replaced by an even steeper function to fine-tune the seed extraction. Seeds with an area smaller than $3px^2$ are removed. For 3D and 3D+t data, detected merged cells in z -direction can be split by increasing the seed extraction threshold ϱ_{seed} till multiple seeds are found for the merged cells. For the detection of merged cells, a priori knowledge about cell sizes or an outlier detection can be used.

Inspired by the Dual U-Net architecture [13], we first attempted to enforce the CNN to predict an additional seed output from the cell distances and the neighbor distances. However, our traditional post-processing provided better results in tests and enables a fine-tuning to cell

types not included in the training data. In addition, it simplifies the architecture and the training process.

Cell tracking

Cell tracking aims to reconstruct the lineage of cells, by linking related cells over time. This linking task is trivial in case of low cell density, error-free segmentation, and high temporal resolution resulting in negligible cell movements between adjacent frames. However, especially for low signal-to-noise-ratio images with touching and dividing cells, fragmented tracks can occur. To re-link fragmented tracks, we match tracks without assigned cells over a short sequence of frames and add a coarse position estimation to the cost function. The proposed algorithm is capable of tracking all segmented cells in an image sequence as well as tracking only a subset, e.g., a selection of manually marked cells.

Initialization. The tracking algorithm traverses the image sequence $\mathcal{I} = \{\mathbf{I}_t \mid 0 \leq t \leq T\}$ forwards, where \mathbf{I}_t is the image at time point t and T the number of time points. A track is initialized for each segmented object in the first frame. For data sets with marked objects in the first frame, tracks are only initialized for marked objects. It is assumed that the object movement between successive frames is small compared to the overall image size. Therefore, for each tracked object a rectangular region of interest (ROI) is defined as a search space for the same object in the next frame. The initial center of each ROI is set to the median position calculated from the first assigned segmentation mask of each track.

Movement estimation. The tracking step consists of a movement estimation followed by a graph-based matching strategy. To estimate the movement of an object, the image frames \mathbf{I}_t and \mathbf{I}_{t+1} are cropped to the object ROI. Then, a phase correlation [28] is calculated between the image crops to estimate a shift between those. The object movement is the shift between the image crops which is given by the position of the maximum peak of the phase correlation. Based on the estimated object movement, the ROI at time point $t + 1$ is adapted for each object individually.

Graph-based matching strategy. All tracks with no successors and their last assigned segmentation mask within time span $\{t - \Delta t, \dots, t\}$ are considered active. Therefore, tracks with missing segmentation masks over at most Δt time points can be re-linked. Next, for each active track a set of potential matching candidates is selected based on its ROI at time point $t + 1$. Active tracks and potential matching candidates are matched by using an adapted version of the coupled minimum-cost flow algorithm proposed in [29]. The algorithm minimizes the overall cost by selecting edges in the graph with minimal cost subject to a set of constraints. The constraints model flow, appearance/disappearance of objects, and splitting/merging of objects. For an in depth introduction please see [29].

Fig 6 shows our adapted graph structure of the coupled minimum-cost flow algorithm. The following adaptations are applied: the appearance cost of objects is set to 0, as spurious tracks will be filtered out by the subsequent post-processing. This appears to be advantageous in scenarios with the objective to track only a few selected objects. The disappearance cost is set to the length of the largest edge of the ROI instead of using appearance-based features. Therefore, tracks with missing segmentation masks can be assigned to the disappearance node as well. The merging node proposed in [29] is removed, as it only models the merging of two objects per time point. The matching cost $c_{s,n}$ between track s and potential matching candidate n is adapted to:

$$c_{s,n} = \|\hat{p}_{t+1}^s - p_{t+1}^n\|_2, \quad (5)$$

where \hat{p}_{t+1}^s is the estimated position of the tracked object s at time point $t + 1$ and p_{t+1}^n is the

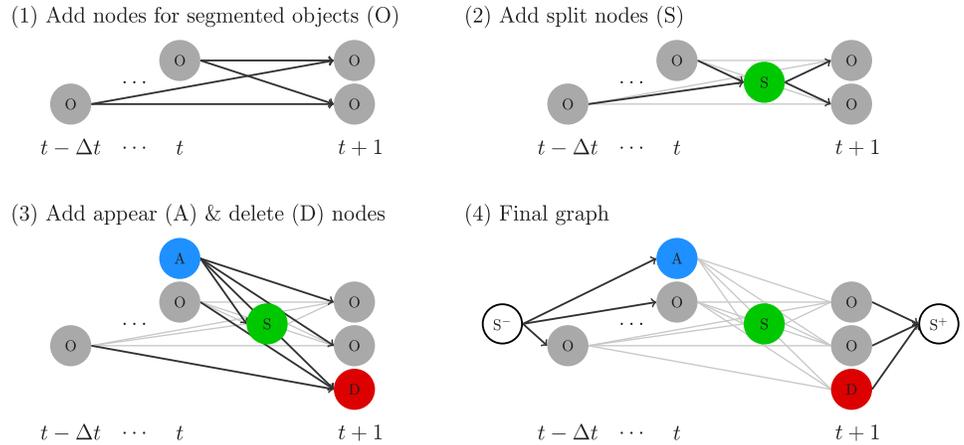


Fig 6. Graph construction steps exemplary for four segmented objects. Edges added in a construction step are black, edges added in previous steps are gray. The gray nodes (O) correspond to segmented objects. The segmented objects from $\{t - \Delta t, \dots, t\}$ are the last matched objects of all active tracks, whereas the segmented objects of $t + 1$ are not matched to tracks yet. The blue node models the appearance of objects (A), the red node the disappearance of objects (D), and the green node split events (S). Split event nodes (S) are added for each pair of objects at $t + 1$. Therefore, a split event node (S) has exactly two outgoing edges but can have several ingoing edges from object nodes (O). Source (S^-) and sink nodes (S^+) are added for the formulation as coupled minimum cost flow problem.

<https://doi.org/10.1371/journal.pone.0243219.g006>

position of the potential matching candidate. The estimated position \hat{p}_{t+1}^s is given by:

$$\hat{p}_{t+1}^s = p_t^s + d_{t,t+1}^s, \tag{6}$$

where $d_{t,t+1}^s$ is the estimated shift of the ROI of track s between time points t and $t + 1$. p_t^s is the position of the tracked object at time point t . The cost of split events are computed based on the size and position of the tracked object s and its potential successor candidates n and k :

$$c_{s,(n,k)} = \begin{cases} \|\hat{p}_{t+1}^s - \frac{1}{2}(p_{t+1}^n + p_{t+1}^k)\|_2 & \text{if } C_s = 1, \\ \rho & \text{else,} \end{cases} \tag{7}$$

where $c_{s,(n,k)}$ are the split costs and C_s the split condition. In practice, we set ρ to ten times the disappearance cost. The split condition C_s is given by:

$$C_s = \begin{cases} 1 & \text{if } \frac{V_{t+1}^n}{V_{t+1}^k} > \alpha, \frac{V_{t+1}^n + V_{t+1}^k}{V_{t_{last}}^s} < \beta, \|p_{t+1}^n - p_{t+1}^k\|_2 < \gamma, \\ 0 & \text{else.} \end{cases} \tag{8}$$

V_{t+1}^n , V_{t+1}^k and $V_{t_{last}}^s$ are the sizes of the segmentation masks of successor candidates n and k , and of the last assigned object to the track s at time point t_{last} , respectively. The successor candidates are sorted so $V_{t+1}^k \geq V_{t+1}^n$ holds. α , β and γ are hyper-parameters. A possible parametrization of those hyper-parameters is provided in the results section. The split condition ensures that successors are of similar size, have a combined size similar to the size of the predecessor object, and should be reasonably close to each other.

Each active track is only linked to segmented objects overlapping with its ROI, reducing the number of edges in the graph. As all active tracks are added to the graph and not only segmented objects between successive time points, tracks with missing segmentation masks over

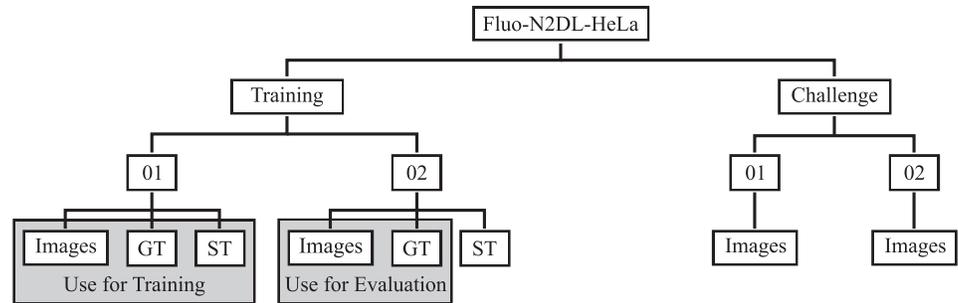


Fig 7. Cell tracking challenge data set structure. Since no ground truths are publicly available for the challenge sets, the two provided training sets need to be split into a set used for training and a set for evaluation.

<https://doi.org/10.1371/journal.pone.0243219.g007>

a short sequence of frames can be linked. A solution of the matching problem is then found using integer linear programming.

For data sets with the aim to track all segmented objects, each non-matched object at time point $t+1$ is initialized as a new track.

Post-processing. In the post-processing step, missing segmentation masks are added by placing masks at the linearly interpolated positions between t_{last} and $t+1$. Furthermore, trajectories of length one without any predecessor and successors are removed.

Results

Data set

We conduct our experiments with data from the Cell Tracking Challenge [7, 9]. For each cell type, the provided data sets are split into two training sets with publicly available ground truths, and two challenge sets (see Fig 7). For our experiments, we use selected data from one training set to train models and evaluate on the other. The provided annotations consist of gold truth (GT) instance segmentation masks, interlinked GT cell seeds for cell detection and tracking, and computer-generated instance segmentation masks, referred to as silver truth (ST). The ST annotations, computed from a majority vote of submitted algorithms of former challenge participants, can include segmentation errors. The GT segmentation masks not necessarily include all cells in a frame.

For four data sets, we manually selected segmentation GTs where all cells in a frame are annotated, and STs that do not show obvious segmentation errors. 27 GTs of the data set BF-C2DL-HSC (Mouse hematopoietic stem cells), 15 GTs of the data set BF-C2DL-MuSC (Mouse muscle stem cells), 16 STs of the data set Fluo-N2DL-HeLa (HeLa cells), and 3 GT slices of the 3D data set Fluo-N3DH-CE (*C. elegans* developing embryo) fulfilled our requirements. This heterogeneous data set will be referred to as CTC training set and consists of 268 crops of size 256px×256px including 52 crops for validation. A difficulty is that the CTC training set contains comparatively few touching cells, whereas in the evaluation the segmentation of touching cells is important, especially for late time points after many cell divisions. Each cell type is evaluated separately using all detection and segmentation GTs of the second set (see S2 Fig).

Evaluation criteria

For evaluation, we use the performance measures of the Cell Tracking Challenge. The normalized acyclic oriented graph matching measure for detection DET is used to evaluate object

level segmentation errors [30]. Pixel level segmentation errors are evaluated with the Jaccard similarity index based measure SEG. The normalized acyclic oriented graph matching measure TRA is used to evaluate the tracking [30]. The overall performances for the Cell Segmentation Benchmark (CSB) and Cell Tracking Benchmark (CTB) are calculated as follows:

$$OP_{CSB} = 0.5 \cdot (DET + SEG), \quad (9)$$

$$OP_{CTB} = 0.5 \cdot (SEG + TRA). \quad (10)$$

Parameter selection

The proposed segmentation method has three adjustable post-processing parameters: the mask threshold ϱ_{mask} , the seed threshold ϱ_{seed} and the standard deviation σ of the Gaussian smoothing. We fix them to: $\varrho_{\text{mask}} = 0.09$, $\varrho_{\text{seed}} = 0.5$, $\sigma = (1.5, 1.5)$ for 2D/2D+t data, and $\sigma = (1.5, 1.5, 0.5)$ for 3D/3D+t data. In practice, a fine-tuning of these parameters is only needed if cells are too small or too large (ρ_{mask}) and if multiple splits or merges occur (ρ_{seed}). For the tracking, we computed cell division and movement statistics from tracking ground truth data and chose the following parameters experimentally: $\Delta t = 3$ (dimensionless difference of frames), $\alpha = 0.5$, $\beta = 1.2$ and $\gamma = 2 \cdot \sqrt[D]{V_{\text{last}}^s}$ with the number of image dimensions $D \in \{2, 3\}$. The ROI is set to $150\text{px} \times 150\text{px}$ for 2D data sets, and to $(100\text{px})^3$ for 3D data sets. For some large 3D+t data sets, e.g., Fluo-N3DL-TRIC and Fluo-N3DL-TRIF of the Cell Tracking Challenge, the ROI is reduced to $(60\text{px})^3$. Due to the observed variety of the cell division and movement statistics over the different data sets, we expect improved tracking results by fine-tuning the tracking parameters to each data set individually.

Compared segmentation methods

The proposed segmentation method is compared with boundary and border prediction methods (Fig 1c and 1d), adapted borders [11], the Dual U-Net [13], and the J4 method proposed in [12].

For the boundary and border prediction methods, we adapt our proposed architecture and use a single decoder path with a three channel output: background, cell, and boundary/border. Instead of the linear activation, the softmax activation is applied in the output layer.

For the adapted borders [11], we use our proposed network architecture with two decoder paths. One decoder path is trained to predict binary cell masks (sigmoid activation), the other to predict background, eroded cells and adapted borders (softmax activation).

The Dual U-Net method [13] uses a similar architecture compared to ours but max-pooling layers and a feature fusion block. Intermediate predictions of discrete boundaries (sigmoid activation) and cell distances (linear activation) are forwarded to the feature fusion block which predicts the final segmentation map (sigmoid activation). We removed in our comparison the dropout layer since none of the other compared methods use dropout.

The last method in our comparison is the J4 method [12] which uses J regularization to tackle the class imbalance problem. The J4 method predicts a four channel output: background, cell, touching, and gap. We use the same architecture with softmax activation in the last layer as for the boundary and border method.

Detailed information about the post-processing of the compared methods is provided in S1 File. Similar to the proposed method, seeds with an area smaller than 3px^2 are removed for all methods. Table 1 shows the boundary and border information in the CTC training set. The

Table 1. Boundary and border information in the CTC training set.

	Boundary	Border	Adapted Border	J4	Proposed
Pixel fraction [%]	10.51	0.11	0.33	0.56 / 0.96	1.51

Stated are the ratios of boundary/border pixels to all pixels. For the proposed neighbor distances only pixels with a value greater than 0.5 are counted in this comparison. Nevertheless, pixels with smaller neighbor distance values can provide information in the training process as well. For the J4 method [12], ratios of the touching and of the gap class are provided.

<https://doi.org/10.1371/journal.pone.0243219.t001>

proposed method can utilize more border information in the training process. For the boundary method the most information is resulting from non-touching cells.

Training settings, inference and experimental environment

For each method, eleven models are trained. This allows to evaluate the robustness of the training. Models are trained with a batch size of 8 using the Adam optimizer [31] and the learning rate is initialized with 8×10^{-4} . After 12 subsequent epochs without validation loss improvement, the learning rate is multiplied by 0.25 till a minimum learning rate of 6×10^{-5} is reached. The training is stopped when 28 subsequent epochs without improvement occurred or 200 epochs are reached. To learn the cell distances and the neighbor distances, PyTorch's SmoothL1Loss is used and both losses are added. The loss functions used to train the compared methods are provided in [S1 File](#). During training the augmentations flipping (probability: 75%), scaling (30%), rotation (30%), contrast changing (30%), blurring (30%), and noise (30%) are applied randomly in this order, and the training images are min-max normalized into the range [-1, 1].

For inference, each frame of a time series is min-max normalized independently into the range [-1, 1], whereas the whole volume is normalized for 3D data. The normalized data are processed frame-by-frame with 3D data being processed slice-wise. The CNN model inputs are zero-padded if necessary.

We performed the experiments using a system with two NVIDIA TITAN RTX GPUs, Ubuntu 18.04, and a Intel Core i9-9900K CPU with 64 GB RAM. The methods are implemented in Python and PyTorch is used as deep learning framework. Implementations of the proposed method and of the compared methods are available at <https://bitbucket.org/tscherr/cell-segmentation-and-tracking/>.

Segmentation results

BF-C2DL-HSC. The segmentation results of the Mouse hematopoietic stem cells in [Fig 8](#) show that the proposed segmentation method provides the best cell detection. The SEG score, which evaluates pixel level errors, is mainly limited due to the fact that the predicted cells in the proposed method are slightly too large as indicated in [Fig 8g](#). These results can be even further improved by fine-tuning the mask threshold. Surprisingly, boundaries can be learned almost as good as adapted borders and better than simple borders. A possible explanation is the small amount of touching cells in the CTC training set which prevents from learning simple borders ([Fig 8c](#)). The Dual U-Net method suffers from some uncertain regions in the final segmentation map prediction ([Fig 8e](#) top) resulting in false negatives and split cells. The limitation of the J4 method is that the touching and the gap class are quite similar for this data set. This results in an oversegmentation and imperfect cell shapes since the gap class is considered to be background. The latter limits mainly the SEG score.

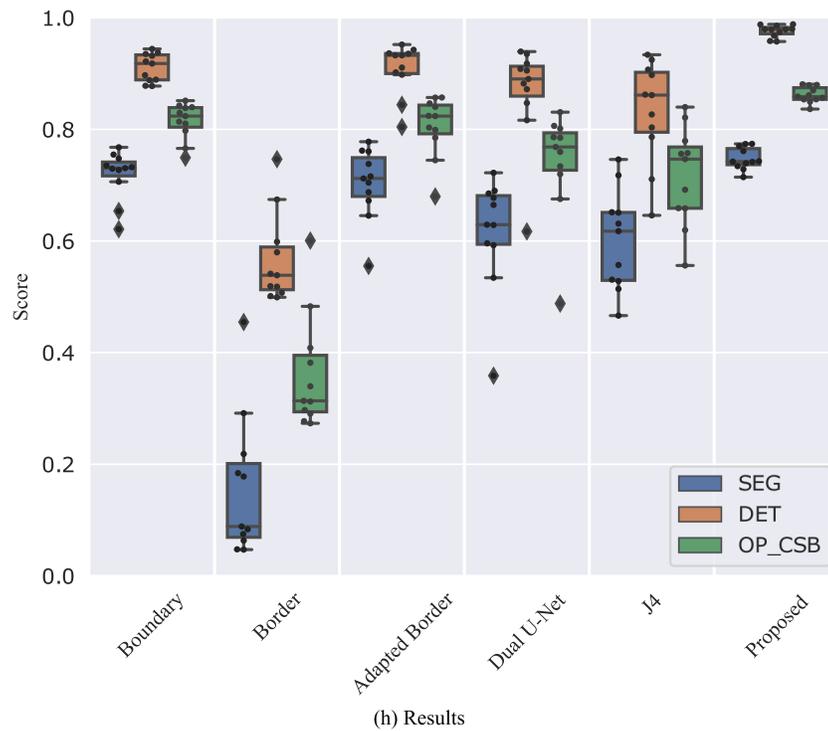
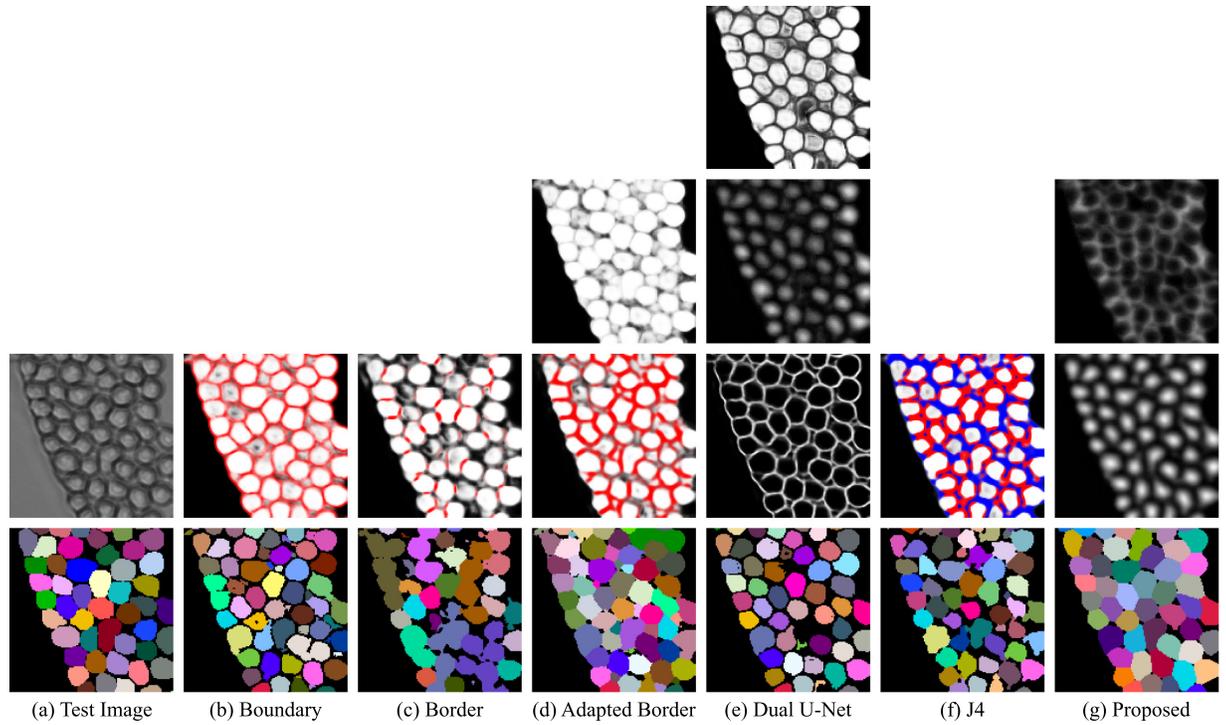


Fig 8. Segmentation results on the BF-C2DL-HSC test set. Shown are raw predictions and segmentations of a 140 px×140 px test image crop (a-g, best OP_{CSB} models). For multi-channel outputs, channels are color-coded (cell/seed class: white, boundary/border/touching class: red, gap class: blue). The plot at the bottom shows the evaluation on the test set (h).

<https://doi.org/10.1371/journal.pone.0243219.g008>

Fluo-N3DH-CE. For the segmentation of the 3D data set Fluo-N3DH-CE, we do not apply the mentioned splitting of cells that are detected as merged. This enables a better comparison of the methods since the almost binary predictions of the other methods do not allow such a simple splitting post-processing. In addition to the 3D nature of this data set, the low signal-to-noise-ratio and the use of only 3 slices of that cell type in the training set makes the segmentation difficult. Again, the proposed method shows the best results whereas boundaries are often unsharp and not closed resulting in merged cells as shown in Fig 9. Borders and adapted borders do not appear anymore and cannot be used to split cells. Especially in late frames after many cell divisions, the boundary and border segmentation methods break down and predict only a few very large objects. In contrast, the J4 method shows an improved splitting of touching cells. However, also the J4 method and the Dual U-Net method decrease in segmentation performance in late frames whereas the proposed method still provides a reasonable segmentation (see S3 Fig).

Fluo-N2DL-HeLa. HeLa cells provide the largest quantity of cells from a specific cell type in the CTC training set resulting in the methods performing more similar, as shown in Fig 10. For this cell type, the adapted border method shows its advantages over the boundary method by predicting robust borders. The models trained with the J4 method learned to predict and differentiate gap and touching class for this cell type very well resulting in the best performance of all methods. However, the proposed method performs also well for this cell type. The Dual U-Net method suffers from merged objects. This is probably due to non-closed boundaries which induce the merging of cells in the feature fusion layer. Our approach with more robust neighbor distances and a traditional post-processing avoids this. S1 Table shows that the neighbor distances can prevent from the erroneous merging of cells.

BF-C2DL-MuSC. Mouse muscle stem cells are difficult to segment since they change their shape from small roundish objects to elongated objects. Both cell states are shown in Fig 11. The Dual U-Net method provides the best segmentation of elongated cells, however, roundish cells are sometimes merged. Nevertheless, the better segmentation of the elongated cells compensates this. The J4 method in contrast suffers from oversegmentation on this cell type, resulting in lower scores. As for the data set BF-C2DL-HSC, the proposed method can handle the small roundish cells well resulting in the second best method for this cell type. The cell distance predictions of the elongated cells are sometimes uncertain and below the seed threshold. This results in missing cells. In contrast, the feature fusion block of the Dual U-Net is able to detect such cells, but has the drawback of merging cells. The segmentation problem of the elongated cells can be solved using the training data of both BF-C2DL-MuSC training data subsets. This is shown in the next section.

Cell tracking challenge

For our submission to the 5th IEEE ISBI 2020 Cell Tracking Challenge as team KIT-Sch-GE, we combined our segmentation method with our adapted tracking approach. We selected a training data set similarly to the CTC training set. The data set consists of 997 crops of size 256px×256px of carefully selected Cell Tracking Challenge data, CBIA HL60 cell line data [32], BBBC038 drosophila images [4], and generated semi-synthetic data [33, 34]. A more detailed description of the data set, data set specific segmentation and tracking parameters, and executables can be found on the challenge website. For our submission we manually selected a segmentation model from three trained models. To avoid issues with the TRA measure, frames without any tracked object are replaced by the tracking result of the temporally closest frame. For the Fluo-N3DH-CE data set cells are split if their volume is bigger than $\frac{4}{3}$ times the mean object volume at that time point by iteratively increasing the seed extraction threshold ρ_{seed} .

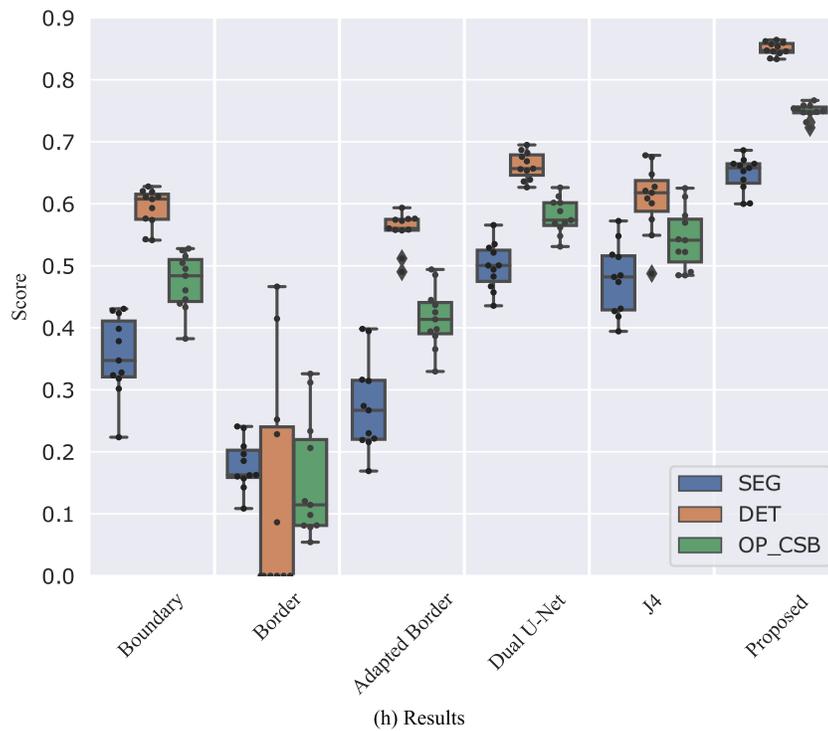
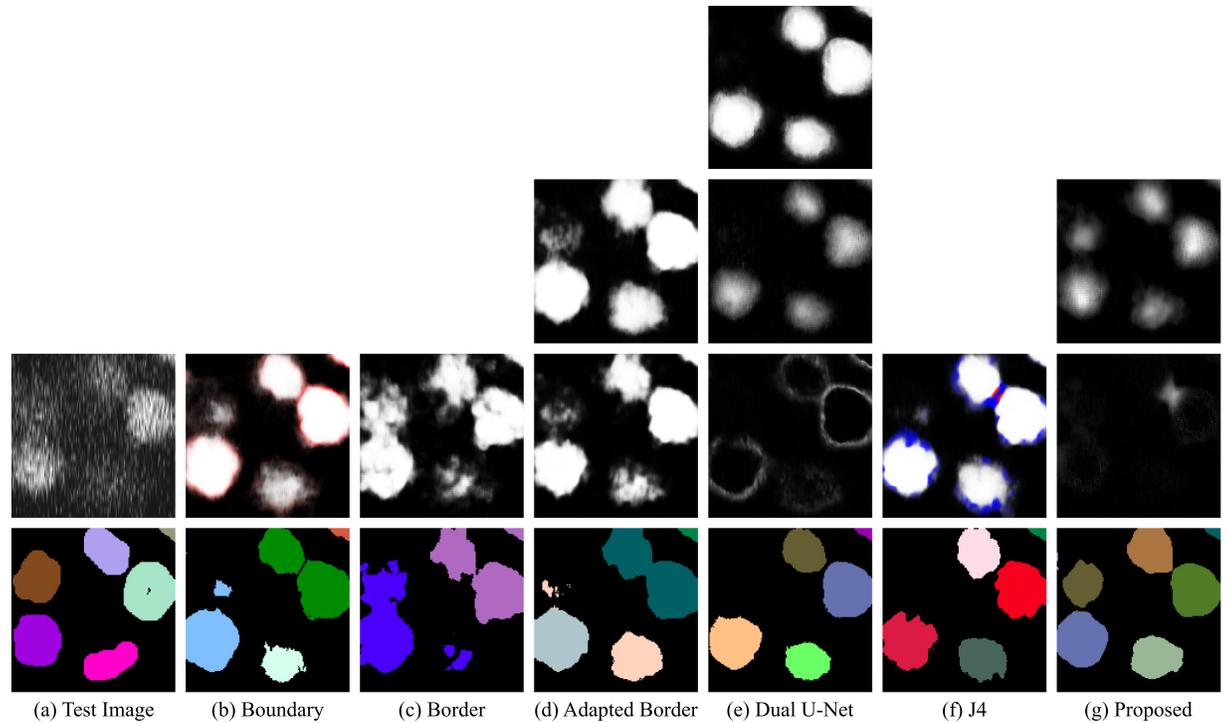


Fig 9. Segmentation results on the Fluo-N3DH-CE test set. Shown are raw predictions and segmentations of a 140 px×140 px test image crop (a-g, best OP_{CSB} models). For multi-channel outputs, channels are color-coded (cell/seed class: white, boundary/border/touching class: red, gap class: blue). The plot at the bottom shows the evaluation on the test set (h). Note: this is a 3D data set and the erroneous merging of cells can result from any of the slices a cell appears.

<https://doi.org/10.1371/journal.pone.0243219.g009>

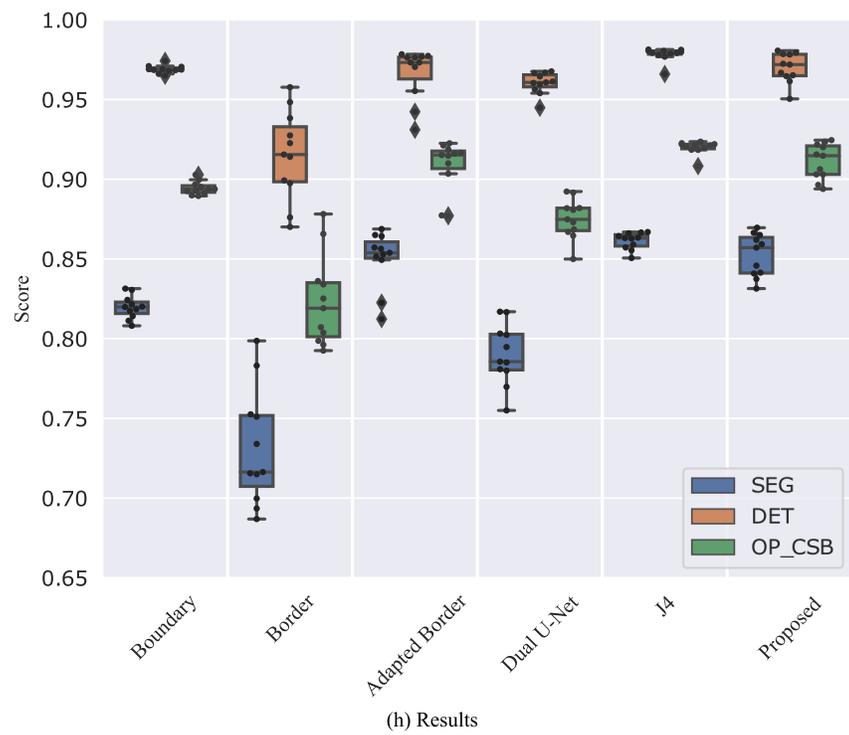
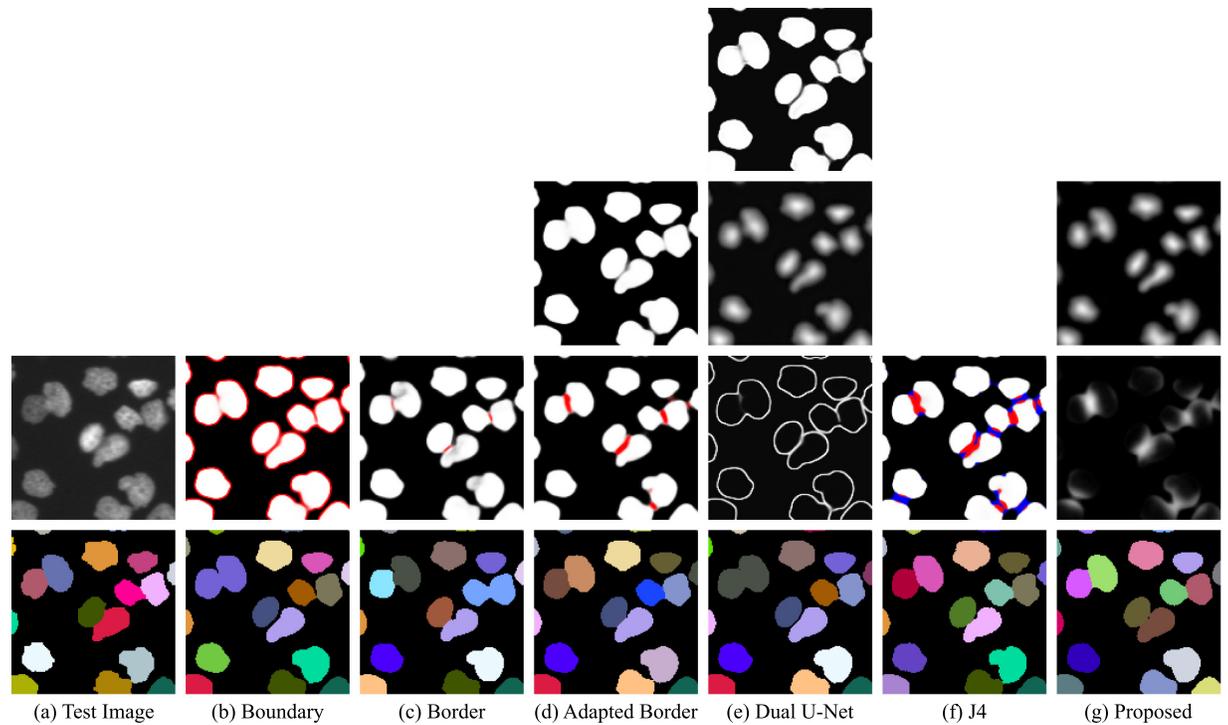


Fig 10. Segmentation results on the Fluo-N2DL-HeLa test set. Shown are raw predictions and segmentations of a 140 px×140 px test image crop (a-g, best OP_{CSB} models). For multi-channel outputs, channels are color-coded (cell/seed class: white, boundary/border/touching class: red, gap class: blue). The plot at the bottom shows the evaluation on the test set (h).

<https://doi.org/10.1371/journal.pone.0243219.g010>

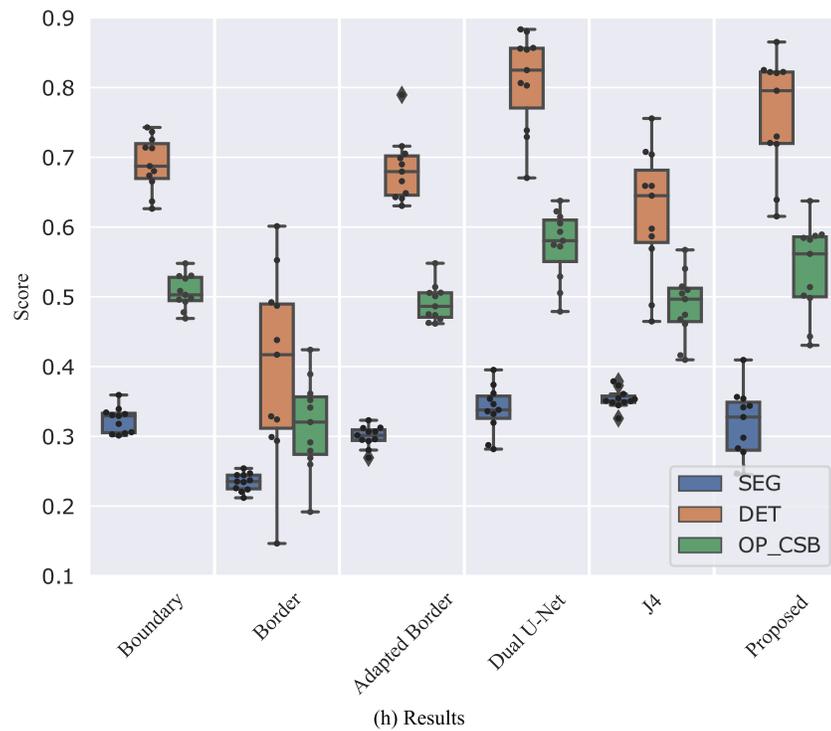
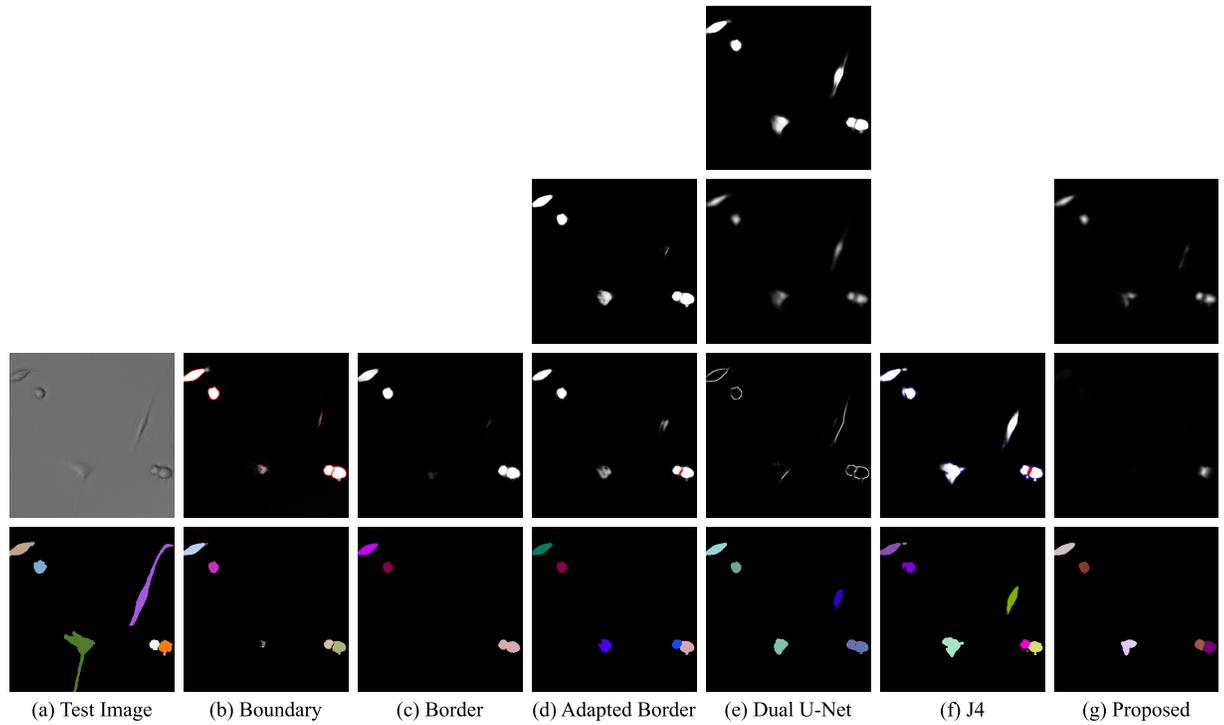


Fig 11. Segmentation results on the BF-C2DL-MuSC test set. Shown are raw predictions and segmentations of a 360 px×360 px test image crop (a-g, best OP_{CSB} models). For multi-channel outputs, channels are color-coded (cell/seed class: white, boundary/border/touching class: red, gap class: blue). The plot at the bottom shows the evaluation on the test set (h).

<https://doi.org/10.1371/journal.pone.0243219.g011>

Table 2. Cell tracking benchmark and cell segmentation benchmark results (5th edition). Top 3 rankings in the overall performances OP_{CSB} and OP_{CTB} are written in bold. The corresponding [Cell Tracking Benchmark](#) and [Cell Segmentation Benchmark](#) leaderboards are available on the Cell Tracking Challenge website.

Data Set	SEG	DET	TRA	OP_{CSB}	OP_{CTB}	Ranking OP_{CSB}	Ranking OP_{CTB}
BF-C2DL-HSC	0.750	0.974	0.929	0.862	0.840	2nd	3rd
BF-C2DL-MuSC	0.702	0.977	0.967	0.839	0.835	1st	1st
Fluo-C3DH-H157 [†]	0.789	0.949	0.948	0.869	0.869	4th	4th
Fluo-C3DL-MDA231 ^{††}	0.616	0.851	0.820	0.733	0.718	3rd	3rd
Fluo-N2DH-GOWT1	0.828	0.950	0.949	0.889	0.889	14th	12th
Fluo-N2DL-HeLa	0.895	0.992	0.989	0.944	0.942	3rd	3rd
Fluo-N3DH-CE ^{††}	0.729	0.930	0.886	0.830	0.808	1st	1st
Fluo-N3DH-CHO	0.871	0.945	0.948	0.908	0.909	3rd	3rd
Fluo-N3DL-DRO	0.562	0.761	-	0.661	-	4th	-
Fluo-N3DL-TRIC [†]	0.821 / 0.766 *	0.961	0.809	0.891	0.787	2nd	2nd
Fluo-N3DL-TRIF [†]	0.601 / 0.573 *	0.926	0.788	0.763	0.680	3rd	3rd
Fluo-N2DH-SIM+	0.800	0.949	0.945	0.875	0.873	9th	7th
Fluo-N3DH-SIM+	0.668	0.937	0.933	0.802	0.800	4th	2nd

[†] No data of that cell type used to train the segmentation model.

^{††} ≤ 5 slices of that cell type used to train the segmentation model.

* Two scores for each benchmark (CSB/CTB) due to a different treatment of additionally detected and segmented cells.

<https://doi.org/10.1371/journal.pone.0243219.t002>

Cell segmentation benchmark. In the Cell Segmentation Benchmark, we achieved eight top three rankings, including two first places, and three fourth places, all of them with the same model (see [Table 2](#)). The performances of the highlighted data sets with no or almost no training data used imply a good utilization of training data and a good generalization ability of our model. A comparison of the scores for the data set BF-C2DL-MuSC on the CSB benchmark and the previous experiment in [Fig 11](#) shows an improved performance. We assume this is due to the larger amount of elongated cells in the training data. Furthermore, the results show that our proposed 2D segmentation with 3D post-processing approach performs well on 3D data. An exemplary segmentation result is shown in [Fig 12](#).

Cell tracking benchmark. In the Cell Tracking Benchmark, we achieved nine top three rankings, including two first places, and a fourth place (see [Table 2](#)). Exemplarily, some

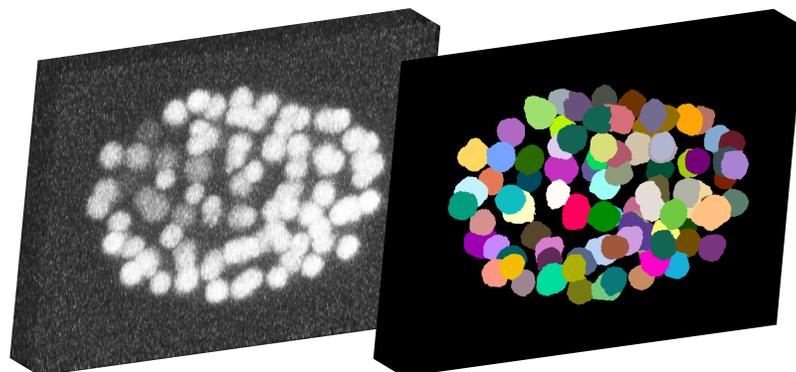


Fig 12. Segmentation result of the Fluo-N3DH-CE challenge data. The maximum intensity projection of the raw data (left) and of the segmentation (right) show that cells can be segmented well even on this challenging data set. [S1 Video](#) shows a video of a tracked developing embryo.

<https://doi.org/10.1371/journal.pone.0243219.g012>

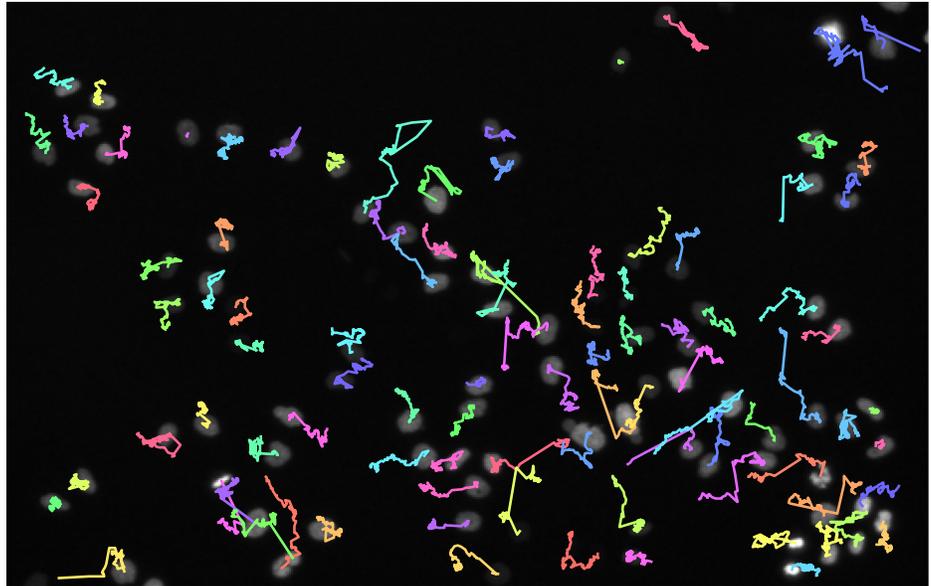


Fig 13. Tracking results on the Fluo-N2DL-HeLa challenge data. The first raw image is overlaid with the tracks starting in the first frame. For better visibility, tracks starting in later frames are excluded. [S2 Video](#) shows a video of the tracked cells.

<https://doi.org/10.1371/journal.pone.0243219.g013>

tracking results of our approach are shown in [Fig 13](#). Some tracks show jumps, visible as long straight lines, possibly due to some remaining linking errors in our adapted tracking approach. However, none of the competing tracking approaches yields perfect tracking results for all cells. The multiple top performances in the Cell Tracking Benchmark show that our tracking approach, which combines a movement estimation and a graph-based matching strategy, belongs to the best performing approaches.

Discussion

Trained on a data set with few touching samples, our proposed segmentation method outperforms all compared methods for at least three of the four cell types evaluated. This is due to the fact that the neighbor distances enable our method to learn from close cells which results in additional information in the training process (see [Table 1](#)) and the fact that this information can be easily combined with the cell distances. The differences between the segmentation results of both discrete border methods show how important the utilization of border information is. We want to emphasize that for the proposed method the seed and the mask thresholds can be adjusted for each cell type and for each trained model separately. This improves the segmentation results shown in [Figs 8–11](#). The other discrete methods do not allow to do so since the needed sigmoid or softmax activation functions prevent a major fine-tuning of the post-processing. However, to allow a better comparison, we fixed the post-processing parameters of our proposed method in our experiments. The results of the J4 method on the HeLa data set show that specialized loss functions work very well, at least for the dominating cell type in the training set. So far, our approach only uses standard loss functions.

Our successful participation in the Cell Segmentation Benchmark and the Cell Tracking Benchmark show that our proposed tracking by detection method yields excellent results in cell segmentation and cell tracking. Especially the success on data sets with only little or very

sparse annotated training data, i.e., data with only very few cells in a frame annotated, shows the advantages of our method.

Conclusion

The segmentation and tracking of touching and dividing cells of different types is a challenging task. In this work, a new cell segmentation method using a combination of cell distances and neighbor distances is proposed. The segmentation method utilizes information from touching and close cells in the training process. Therefore, it shows an improved generalization ability for cell types underrepresented or absent in the training data set compared to border and boundary prediction methods. This advantage enables to segment even cell types with no or almost no annotated training data available. Our success in the Cell Segmentation Benchmark emphasizes the strengths of our segmentation method. Our adapted tracking algorithm, which uses a movement estimation with a graph-based matching strategy, can handle cell divisions and missing segmentation masks over a short sequence of frames. The combination of the tracking with our proposed segmentation method resulted in top performances at the Cell Tracking Benchmark.

As future research, we plan to further improve the segmentation performance using a larger and on ImageNet pre-trained encoder or mixed convolution blocks [35], test-time augmentation [36], and the synthetic generation of new training samples [16]. In addition, studies about how cell features, e.g., size, shape and texture, influence the generalization ability to new cell types are needed. A long-term goal is to develop a user-friendly software for the segmentation and tracking of a large variety of cell types using a well-trained segmentation model. Including tunable post-processing parameters facilitates an adaptation of the cell and neighbor distances to new data.

Supporting information

S1 Fig. CNN architecture and data flow for a 256px×256px test image.

(PDF)

S2 Fig. Training and test data set collection. The ground truths of the challenge sets from the Cell Tracking Challenge are not publicly available. Thus, the two training data sets are split into a training set and cell type specific test sets for our segmentation experiments. For the training data set only fully annotated segmentation GTs and good quality STs can be used to train models well. For evaluation, all segmentation and detection GTs can be used.

(PDF)

S3 Fig. Segmentation results on the Fluo-N3DH-CE test set (late frame). Only the proposed method is able to prevent from merging close cells in late frames after many cell divisions.

Note: this is a low-resolution 3D data set and the erroneous merging of cells can result from any of the slices a cell appears.

(PDF)

S1 File. Post-processing and loss functions of the compared methods.

(PDF)

S1 Table. Resolved merges due to the use of neighbor distances exemplary for the best proposed OP_{CSB} models. Especially for the HeLa cells less erroneously merged cells occur compared to using only the predicted cell distance information. This enables the proposed method to be a good generalist in our comparison.

(PDF)

S1 Video. Tracking results on the Fluo-N3DH-CE challenge data.
(MP4)

S2 Video. Tracking results on the Fluo-N2DL-HeLa challenge data.
(MP4)

Acknowledgments

The authors would like to thank Andreas Bartschat for proofreading, the organizers of the Cell Tracking Challenge, and all data set providers.

Author Contributions

Conceptualization: Tim Scherr, Katharina Löffler, Moritz Böhland, Ralf Mikut.

Formal analysis: Tim Scherr, Katharina Löffler.

Funding acquisition: Ralf Mikut.

Investigation: Tim Scherr, Katharina Löffler.

Methodology: Tim Scherr, Katharina Löffler.

Project administration: Ralf Mikut.

Software: Tim Scherr, Katharina Löffler.

Supervision: Ralf Mikut.

Writing – original draft: Tim Scherr, Katharina Löffler.

Writing – review & editing: Tim Scherr, Katharina Löffler, Moritz Böhland, Ralf Mikut.

References

1. Chhetri RK, Amat F, Wan Y, Höckendorf B, Lemon WC, Keller PJ. Whole-animal functional and developmental imaging with isotropic spatial resolution. *Nat Methods*. 2015; 12:1171–1178. <https://doi.org/10.1038/nmeth.3632>
2. Kobitski AY, Otte JC, Takamiya M, Schäfer B, Mertes J, Stegmaier J, et al. An ensemble-averaged, cell density-based digital model of zebrafish embryo development derived from light-sheet microscopy data with single-cell resolution. *Sci Rep*. 2015; 5(8601):1–10. <https://doi.org/10.1038/srep08601> PMID: 25712513
3. Khairy K, Keller PJ. Reconstructing embryonic development. *Genesis*. 2011; 49:488–513. <https://doi.org/10.1002/dvg.20698>
4. Caicedo JC, Goodman A, Karhohs KW, Cimini BA, Ackerman J, Haghighi M, et al. Nucleus segmentation across imaging experiments: the 2018 Data Science Bowl. *Nat Methods*. 2019; 16:1247–1253. <https://doi.org/10.1038/s41592-019-0612-7> PMID: 31636459
5. Schott B, Traub M, Schlagenhauf C, Takamiya M, Anritter T, Bartschat A, et al. EmbryoMiner: a new framework for interactive knowledge discovery in large-scale cell tracking data of developing embryos. *PLOS Comput Biol*. 2018; 14(4):1–18. <https://doi.org/10.1371/journal.pcbi.1006128> PMID: 29672531
6. Wolff C, Tinevez JY, Pietzsch T, Stamatakis E, Harich B, Guignard L, et al. Multi-view light-sheet imaging and tracking with the MaMuT software reveals the cell lineage of a direct developing arthropod limb. *eLife*. 2018; 7:e34410. <https://doi.org/10.7554/eLife.34410> PMID: 29595475
7. Ulman V, Maška M, Magnusson KE, Ronneberger O, Haubold C, Harder N, et al. An objective comparison of cell-tracking algorithms. *Nat Methods*. 2017; 14:1141–1152. <https://doi.org/10.1038/nmeth.4473> PMID: 29083403
8. Stegmaier J, Otte JC, Kobitski A, Bartschat A, Garcia A, Nienhaus GU, et al. Fast segmentation of stained nuclei in terabyte-scale, time resolved 3d microscopy image stacks. *PLOS ONE*. 2014; 9(2):1–11. <https://doi.org/10.1371/journal.pone.0090036> PMID: 24587204

9. Maška M, Ulman V, Svoboda D, Matula P, Matula P, Ederra C, et al. A benchmark for comparison of cell tracking algorithms. *Bioinformatics*. 2014; 30(11):1609–1617. <https://doi.org/10.1093/bioinformatics/btu080> PMID: 24526711
10. Vicar T, Balvan J, Jaros J, Jug F, Kolar R, Masarik M, et al. Cell segmentation methods for label-free contrast microscopy: review and comprehensive comparison. *BMC Bioinformatics*. 2019; 20:360. <https://doi.org/10.1186/s12859-019-2880-8>
11. Scherr T, Bartschat A, Reischl M, Stegmaier J, Mikut R. Best practices in deep learning-based segmentation of microscopy images. In: *Proceedings 28. Workshop Computational Intelligence, Dortmund, Germany*; 2018. p. 175–195.
12. Peña FAG, Fernandez PDM, Tarr PT, Ren TI, Meyerowitz EM, Cunha A. J regularization improves imbalanced multiclass segmentation. In: *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*; 2020. p. 1–5.
13. Li X, Wang Y, Tang Q, Fan Z, Yu J. Dual U-Net for the segmentation of overlapping glioma nuclei. *IEEE Access*. 2019; 7:84040–84052. <https://doi.org/10.1109/ACCESS.2019.2924744>
14. Li J, Hu Z, Yang S. Accurate nuclear segmentation with center vector encoding. In: *Inf Process Med Imaging*. Springer International Publishing; 2019. p. 394–404.
15. Graham S, Vu QD, Raza SEA, Azam A, Tsang YW, Kwak JT, et al. Hover-Net: simultaneous segmentation and classification of nuclei in multi-tissue histology images. *Med Image Anal*. 2019; 58:101563. <https://doi.org/10.1016/j.media.2019.101563> PMID: 31561183
16. Hollandi R, Szkalicity A, Toth T, Tasnadi E, Molnar C, Mathe B, et al. nucleAlzer: a parameter-free deep learning framework for nucleus segmentation using image style transfer. *Cell Syst*. 2020; 10(5):453–458.e6. <https://doi.org/10.1016/j.cels.2020.04.003>
17. Ciaparrone G, Sánchez FL, Tabik S, Troiano L, Tagliaferri R, Herrera F. Deep learning in video multi-object tracking: a survey. *Neurocomputing*. 2020; 381:61–88. <https://doi.org/10.1016/j.neucom.2019.11.023>
18. Yao R, Lin G, Xia S, Zhao J, Zhou Y. Video object segmentation and tracking: a survey. *ACM Trans Intell Syst Technol*. 2020; 11(4). <https://doi.org/10.1145/3391743>
19. Payer C, Štern D, Feiner M, Bischof H, Urschler M. Segmenting and tracking cell instances with cosine embeddings and recurrent hourglass networks. *Med Image Anal*. 2019; 57:106–119. <https://doi.org/10.1016/j.media.2019.06.015>
20. He T, Mao H, Guo J, Yi Z. Cell tracking using deep neural networks with multi-task learning. *Image Vis Comput*. 2017; 60:142–153. <https://doi.org/10.1016/j.imavis.2016.11.010>
21. Magnusson KEG, Jalden J, Gilbert PM, Blau HM. Global linking of cell tracks using the Viterbi algorithm. *IEEE Trans Med Imaging*. 2015; 34(4):911–929. <https://doi.org/10.1109/TMI.2014.2370951>
22. Dendorfer P, Rezatofighi H, Milan A, Shi J, Cremers D, Reid I, et al. CVPR19 Tracking and Detection Challenge: how crowded can it get?; 2019. Available from: <https://arxiv.org/abs/1906.04567>.
23. Pont-Tuset J, Perazzi F, Caelles S, Arbeláez P, Sorkine-Hornung A, Van Gool L. The 2017 DAVIS Challenge on video object segmentation; 2018. Available from: <https://arxiv.org/abs/1704.00675>.
24. Xu N, Yang L, Fan Y, Yue D, Liang Y, Yang J, et al. YouTube-VOS: a large-scale video object segmentation benchmark; 2018. Available from: <https://arxiv.org/abs/1809.03327>.
25. Sixta T, Cao J, Seebach J, Schnittler H, Flach B. Coupling cell detection and tracking by temporal feedback. *Mach Vis Appl*. 2020; 31(4):24. <https://doi.org/10.1007/s00138-020-01072-7>
26. Ljosa V, Sokolnicki KL, Carpenter AE. Annotated high-throughput microscopy image sets for validation. *Nat Methods*. 2012; 9:637. <https://doi.org/10.1038/nmeth.2083>
27. Ronneberger O, Fischer P, Brox T. U-Net: convolutional networks for biomedical image segmentation. In: Navab N, Hornegger J, Wells WM, Frangi AF, editors. *MICCAI 2015*. Springer International Publishing; 2015. p. 234–241.
28. Kuglin CD, Hines DC. The phase correlation image alignment method. In: *Proceedings of the IEEE International Conference on Cybernetics and Society*; 1975. p. 163–165.
29. Padfield D, Pittscher J, Roysam B. Coupled minimum-cost flow cell tracking for high-throughput quantitative analysis. *Med Image Anal*. 2011; 15(4):650–668. <https://doi.org/10.1016/j.media.2010.07.006>
30. Matula P, Maška M, Sorokin DV, Matula P, Ortiz-de Solórzano C, Kozubek M. Cell tracking accuracy measurement based on comparison of acyclic oriented graphs. *PLOS ONE*. 2015; 10(12):1–19. <https://doi.org/10.1371/journal.pone.0144959>
31. Kingma DP, Ba J. Adam: a method for stochastic optimization. In: *3rd International Conference on Learning Representations, San Diego*; 2015. Available from: <http://arxiv.org/abs/1412.6980>.
32. Svoboda D, Kozubek M, Stejskal S. Generation of digital phantoms of cell nuclei and simulation of image formation in 3d image cytometry. *Cytometry A*. 2009; 75A(6):494–509. <https://doi.org/10.1002/cyto.a.20714>

33. Stegmaier J. New methods to improve large-scale microscopy image analysis with prior knowledge and uncertainty. ME, KIT, Karlsruhe, GER. Karlsruhe, GER; 2017.
34. Stegmaier J, Arz J, Schott B, Otte JC, Kobitski A, Nienhaus GU, et al. Generating semi-synthetic validation benchmarks for embryomics. In: 2016 IEEE 13th International Symposium on Biomedical Imaging ISBI; 2016. p. 684–688.
35. Huang C, Ding H, Liu C. Segmentation of cell images based on improved deep learning approach. *IEEE Access*. 2020; 7:2169–3536. <https://doi.org/10.1109/ACCESS.2020.3001571>
36. Moshkov N, Mathe B, Kertesz-Farkas A, Hollandi R, Horvath P. Test-time augmentation for deep learning-based cell segmentation on microscopy images. *Sci Rep*. 2020; 10 (5068). <https://doi.org/10.1038/s41598-020-61808-3> PMID: 32193485