



TRACKASTRA: Transformer-based cell tracking for live-cell microscopy

Benjamin Gallusser¹  and Martin Weigert^{1,2,3} 

¹ Institute of Bioengineering, School of Life Sciences, EPFL, Switzerland

² Center for Scalable Data Analytics and AI (ScaDS.AI), Dresden/Leipzig, Germany

³ Faculty of Computer Science, TU Dresden, Germany

`benjamin.gallusser@epfl.ch, marweigert@gmail.com`

Abstract. Cell tracking is a ubiquitous image analysis task in live-cell microscopy. Unlike multiple object tracking (MOT) for natural images, cell tracking typically involves hundreds of similar-looking objects that can divide in each frame, making it a particularly challenging problem. Current state-of-the-art approaches follow the tracking-by-detection paradigm, *i.e.* first all cells are detected per frame and successively linked in a second step to form biologically consistent cell tracks. Linking is commonly solved via discrete optimization methods, which require manual tuning of hyperparameters for each dataset and are therefore cumbersome to use in practice. Here we propose TRACKASTRA, a general purpose cell tracking approach that uses a simple transformer architecture to directly learn pairwise associations of cells within a temporal window from annotated data. Importantly, unlike existing transformer-based MOT pipelines, our learning architecture also accounts for dividing objects such as cells and allows for accurate tracking even with simple greedy linking, thus making strides towards removing the requirement for a complex linking step. The proposed architecture operates on the full spatio-temporal context of detections within a time window by avoiding the computational burden of processing dense images. We show that our tracking approach performs on par with or better than highly tuned state-of-the-art cell tracking algorithms for various biological datasets, such as bacteria, cell cultures and fluorescent particles. We provide code at <https://github.com/weigertlab/trackastra>.

1 Introduction

The accurate tracking of cells in microscopy videos is an important step in many biological experiments, for example when studying the temporal dynamics of bacteria colonies, eukaryotic cell cultures, or whole-embryo development [19, 31, 55]. This task not only consists of identifying the trajectories of individual cells, but as well requires to correctly assign mother-daughter relationships of dividing cells over multiple generations, which is crucial for inference of cell lineages. Typical live-cell videos can contain hundreds or thousands of visually similar cells that continuously divide and that exhibit various local and global movement patterns in between divisions, which contributes to the complexity

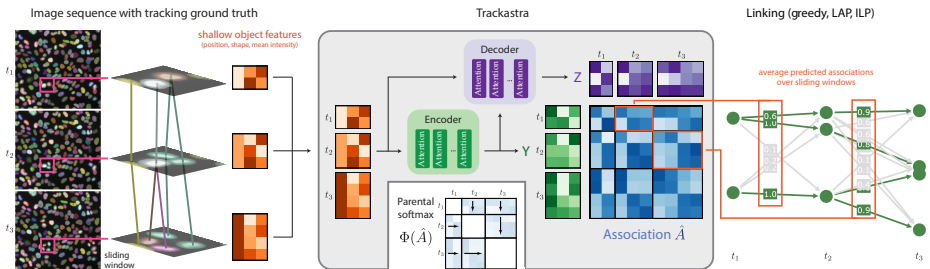


Fig. 1: Overview of TRACKASTRA. Given frame-by-frame object detections in a live-cell video, object features are extracted from a small temporal window and passed as tokens into an encoder-decoder transformer, to predict pairwise associations \hat{A} . We apply a *parental softmax* normalization on \hat{A} to guide the learning directly towards biologically plausible associations. Finally, we build a candidate graph by averaging the predictions \hat{A} over a sliding window, and obtain a tracking solution by pruning the graph with either a greedy algorithm or discrete optimization.

of the task (*cf.* Fig. 2 for examples). While cell tracking has been performed manually during initial experiments of embryo development [48], many specialized semi-automated and fully automated algorithms have been developed over the last two decades [16]. Most current state-of-the-art cell tracking methods follow the *tracking-by-detection* approach, where first cell instances are detected in each frame of a recorded video, which then in a second step are linked across frames to form biologically valid tracks. In recent years, several robust deep-learning-based methods have emerged for the previously tedious and parameter-sensitive first cell detection step [15, 38, 40, 44, 45]. However, the second linking step is still commonly solved with specialized discrete optimization methods such as integer linear programming (ILP) that allow to enforce the biological constraints on a graph data structure [2, 17, 20, 39]. The costs for such optimization methods have to be carefully chosen, and can range from simple features such as Euclidean distance or intersection-over-union between objects [19, 50] to explicit motion models [52]. More recently, annotated tracking ground truth videos have been used to learn cell flow models with convolutional neural networks for better cost prediction [12–14, 17, 27, 31, 47], or to directly learn object associations in a local context via graph neural networks [1, 42] or via iterative semantic segmentation [37]. Furthermore in [2, 51], multiple segmentation hypotheses are considered during the discrete linking optimization problem, and a feasible unique solution is enforced with adequate additional constraints.

In this paper, we explore the potential of directly learning association costs between segmented objects using a straightforward transformer architecture. This approach models the all-to-all interactions between objects within a short temporal window, while relying solely on shallow per-object features like position and basic shape features. Importantly, we aim to learn the correct associations even for dividing objects, in order to reduce the reliance on computationally

expensive combinatorial optimizers and to potentially allow for a greedy linking post-processing only based on the learned associations.

The transformer neural network architecture [54] was originally introduced in the context of language translation, a sequence-to-sequence task, where the input is a set of tokens embedded in a continuous vector space. These are then fed through (self/cross) attention layers that enable the model to reason across all tokens at once. This particular inductive bias has produced state-of-the-art results for various predictive tasks using different elementary tokens, for example image patches as tokens for image classification [7], object detection [3] and segmentation [24], amino acids as tokens for protein structure prediction [21], and keypoints as tokens for image feature matching [26] or tasks on point clouds [28]. Notably, multiple object tracking (MOT) in the domain of natural images has recently been successfully approached with transformers in [4, 36, 49, 59]. Since the temporal succession of cell detections in a video can be seen as a sequence of tokens, we here investigate whether a purely transformer-based model can be used to learn the correct associations between cell detections in a video, especially in the presence of dividing cells.

Our contributions are as follows: *i)* We provide TRACKASTRA (*T*racking-by-*A*ssociation with *T*ransformers), which is, to the best of our knowledge, the first plain transformer-based cell tracking method that accounts for dividing objects. Our method is simple, formulating tracking as a direct association prediction task followed by greedy linking, thereby reducing the reliance on the computationally expensive optimizer step whose hyperparameter tuning is a common practical challenge in cell tracking experiments. *ii)* We show that very simple object features such as position or basic shape features are enough for the subsequent model to learn the correct associations, which removes the need for dedicated visual feature extractors such as pretrained CNNs. *iii)* To ensure the biological consistency of the learned associations, we introduce a blockwise *parental softmax* normalization method for the association matrix, which allows for one-to-many associations, but not many-to-one. *iv)* We evaluate our approach experimentally on three datasets from different modalities and demonstrate that TRACKASTRA outperforms even baselines that are highly tuned for specific model organisms.

1.1 Related work

Cell tracking in microscopy videos is closely related to the extensively studied *multiple object tracking* (MOT) problem in computer vision [29], for which in recent years several transformer-based approaches have been proposed [4, 36, 49, 59]. In contrast to MOT [25], however, the cell tracking problem exhibits two crucial differences that motivate our work: First, raw image frames are typically grayscale and contain thousands of objects with very similar visual appearance that are hard to distinguish locally. Second, objects are allowed to split (but generally not to fuse), which introduces a more complex combinatorial solution space and precludes computationally efficient tracking approaches commonly used in MOT, such as network flows [41].

MOT with transformers In [6, 36, 49, 57] MOT is done autoregressively end-to-end with a transformer, *i.e.* the attention is computed between all candidate object detection tokens in frame t (the keys) and the aggregated token representations of already linked tracks until frame $t - 1$ (the queries). Alternatively, the attention matrix can be computed with bi-directional temporal context between single detections only, considering all detections within a local sliding window [35, 59]. Closest to our work is the tracking-by-detection approach *Global Tracking Transformers* [59], which, however, cannot account for cell divisions and heavily relies on appearance features of each detection that are extracted with a convolutional neural network (CNN) from the raw video frames.

Learning based association prediction for cell tracking In microscopy image analysis, transformer-based tracking has been used in the context of single particle tracking [58] where most detections look similar like in our problem setting, however without supporting dividing objects. In [37], the association step involves densely predicting for each object its mask in the subsequent frame, which is then used to generate appropriate linking costs. For dividing objects, explicitly learning associations between objects has been explored with Graph Neural Networks in [1, 42], where object features extracted by a CNN interact via the given structure of a hypothesis graph. This constrains detection interactions to a certain locality, in contrast to direct all-to-all interactions in a transformer.

2 Method

Our proposed TRACKASTRA method operates on raw image sequences and corresponding detections (or segmentation masks) and uses an encoder-decoder transformer to directly predict the pairwise association matrix A between all detections in a local window of consecutive time frames. Specifically, we construct a token for each object and timepoint within the local window and use this sequence of tokens as input to the transformer. The predicted associations \hat{A} are then used as costs in a candidate track graph that is pruned either greedily or via discrete optimization to obtain the final cell tracks. An overview of the full pipeline is shown in Fig. 1. The following sections describe the dataset and training target construction, the transformer architecture, the loss function, the inference and final link assignment, and implementation details.

2.1 Dataset and association matrix construction

Let $I_1, I_2, \dots, I_T \in \mathbb{R}^{w \times h}$ be an image sequence that is grouped into overlapping windows $S_1, \dots, S_{T-s+1} \in \mathbb{R}^{s \times w \times h}$ of size s . Each window S_n contains a set of detections $\{d_i\}$ that each correspond to a time point $t_i \in \mathbb{N}$, a center point $p_i \in \mathbb{R}^2$, a segmentation mask $m_i \in \{0, 1\}^{w \times h}$, and other potential object features $z_i \in \mathbb{R}^z$ such as basic shape descriptors or mean image intensity of the instance. The goal of the model is to predict an association probability matrix $\hat{A} = (\hat{a}_{ij})$ between all d_i in the window S_n . To construct the target association matrix

$A = (a_{ij})$ the set of detections $\{d_i\}$ is matched to the set of ground truth objects $V = \{v_k\}$ and their ground truth associations. Each ground truth object again corresponds to a time point t_k , center point p_k , and a segmentation mask m_k and the tracking associations can be described as a directed tree $G = (V, E)$. An edge e_{kl} , $k, l \in V$ exists only if $t_k + 1 = t_l$ and the objects v_k and v_l represent the same cell at different time points, or if v_k is the mother cell of v_l . As a simple matching criterion between detections d_i and ground truth objects v_k we use

$$M_{ik} = \max \left(\text{IoU}(m_i, m_k), 1 - \frac{\|p_i - p_k\|_2}{\delta_{max}} \right) > 0.5 \quad , \quad (1)$$

where δ_{max} is a distance threshold and IoU denotes the intersection-over-union. The final matching is then obtained by solving a minimum cost bipartite matching problem based on the costs M_{ik} between $\{d_i\}$ and $\{v_k\}$. Finally, for all matched pairs of detections (d_i, v_{k_i}) and (d_j, v_{k_j}) , we set $a_{ij} = 1$ if v_{k_i} and v_{k_j} are part of the same sub-lineage, *i.e.* iff $v_{k_i} \in \text{descendants}(v_{k_j})$ or $v_{k_i} \in \text{ancestors}(v_{k_j})$, otherwise we set $a_{ij} = 0$. Note that this way, also associations across non-adjacent timepoints as well as appearing and disappearing objects are supported.

2.2 Transformer architecture

The input tokens $x_i \in \mathbb{R}^d$ are constructed by using learned Fourier spatial positional encodings Θ for the detection positions p_i , concatenating them with the object features z_i , and projecting them onto the token dimensionality d :

$$x_i = W_{inp}[\Theta(p_i), z_i] \quad . \quad (2)$$

where z_i are the low-dimensional feature vector containing shallow texture and morphological features (such as mask area or mean intensity) and W_{inp} is a linear projection layer mapping the concatenated tensor to \mathbb{R}^d . The model consists of an encoder-decoder transformer architecture of $2L$ multi-head attention layers with 4 heads each (*cf.* Fig. 1):

$$\mathcal{A}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d}} + M \right) V \quad , \quad (3)$$

where Q are the projected attention *queries*, K the projected *keys*, V the projected *values*, and M is a mask disabling attention for all token pairs whose distance is larger than a user defined threshold d_{max} , *i.e.* $M_{ij} = 0$ if $\|p_i - p_j\|_2 \leq d_{max}$ and $M_{ij} = -\infty$ otherwise. In every attention layer, we additionally use rotary positional embeddings (RoPE [46]) to directly and efficiently inject relative spatial and temporal information. The encoder f transforms the input tokens using L self-attention layers $\mathcal{A}_f^\ell(X, X, X)$ to obtain representations $Y = f(X)$. The decoder g uses L cross-attention layers $\mathcal{A}_g^\ell(X, Y, Y)$ to obtain a second set of representations $Z = g(X, Y)$. In between attention layers we use a simple two-layer MLP with GeLU activation, layer normalization and add residual connections following [54]. Finally, we apply two-layer MLPs to Y and Z and compute the logits \hat{A} of the final association matrix as the outer product $\hat{A} = \text{MLP}_Y(Y) \cdot \text{MLP}_Z(Z)^T$.

2.3 Parental softmax

Given the predicted association logits \hat{A} , a simple approach to extract association probabilities $\tilde{A} \in (0, 1)$ would be to apply a sigmoid to each entry of \hat{A} , *i.e.* $\tilde{A} = \sigma(\hat{A})$. However, this approach does not enforce the combinatorial constraints of cell tracking, *i.e.* the uniqueness of each object’s parent while allowing for more than one child, as well as appearance and disappearance of objects. To remedy this, we propose a logit normalization that we call *parental softmax* and which ensures that the block-wise sum of all entries in the vector of possible parent associations for each d_i is at most one (*cf.* Fig. 1). Concretely, we define the parental softmax $\Phi(\hat{A})$ as

$$\tilde{A} = \Phi(\hat{A})_{ij} = \frac{\exp(\hat{A}_{ij})}{1 + \sum_{i' \in \mathcal{P}_j} \exp(\hat{A}_{i'j})} \quad , \quad (4)$$

where $\mathcal{P}_j = \{d_{i'} | t_{i'} = t_j - 1, \forall i' \in D\}$ denotes all detections in the frame before detection d_j . Note that adding a constant to the denominator (*quiet softmax*) allows for detections to not be assigned to any parent detection, accommodating for appearing and disappearing objects. We then define the loss to be minimized during training as

$$\mathcal{L}(A, \hat{A}, W) = \mathcal{L}_{BCE}(A, \Phi(\hat{A}), W) + \lambda \mathcal{L}_{BCE}(A, \sigma(\hat{A}), W) \quad , \quad (5)$$

where \mathcal{L}_{BCE} is the usual element-wise binary cross-entropy loss, $\lambda \in \mathbb{R}$ is a small fixed hyperparameter (we use $\lambda = 10^{-2}$ throughout), and W is a weighting factor for each matrix element. The elementwise weighting terms W are set to

$$w_{ij} = \begin{cases} 0 & t_j - t_i > \Delta t & \text{temporal cutoff} \\ & \vee t_j - t_i < 1 & \text{only forward links} \\ 1 + \lambda_{\text{div}} & \deg^+(v_{k_i}) = 2 & \text{dividing cells} \\ 1 + \lambda_{\text{cont}} & \deg^+(v_{k_i}) = 1 & \text{continuing tracks} \\ 1 & \text{otherwise} \end{cases} \quad , \quad (6)$$

where $\deg^+(v)$ is the out-degree of vertex v in G . We choose $\Delta t = 2$, $\lambda_{\text{div}} = 10$ and $\lambda_{\text{cont}} = 1$ as fixed hyperparameters. This choice effectively up-weights the loss for cell divisions and continuing tracks, and removes the loss for associations that are not used during the linking step.

2.4 Inference and linking

Inference is done with a sliding window of size s as in training. To obtain global scalar association scores $0 \leq \bar{a}_{i'j'} \leq 1$ from $\hat{A}^{(1)}, \dots, \hat{A}^{(T-s+1)}$, where i' and j' are global detection indices in a video I_1, I_2, \dots, I_T , we take the mean over the $s - 1$ windows that include this association

$$\bar{a}_{i'j'} = \frac{1}{s-1} \sum_{\{S_n | i', j' \in S_n\}} \tilde{a}_{i'j'}^{(S_n)} \quad . \quad (7)$$

Next, we build a candidate graph $G_C = (V, E)$ with a maximum admissible Euclidean distance $dist_{max}$ between detections in adjacent time frames. For this, we use associations $\bar{a}_{i'j'}$ with $t'_{j'} - t'_i = 1$, *i.e.* the upper blockwise diagonal of \bar{A} . To generate a first association candidate graph we directly discard small associations with $\bar{a}_{i'j'} < \alpha$ with $\alpha = 0.05$. This candidate graph is then pruned to a solution graph $G_S = (V_S, E_S)$ with $V_S \subseteq V, E_S \subseteq E$ with one of the following linking algorithms:

Greedy We iteratively add edges and their incident nodes to G_S , ordered by descending edge probability, if the edge probability $\theta \geq 0.5$ and if the edge does not violate the biological constraints (*i.e.* at most two children, and at most one parent per vertex)

$$\deg^+(v) \leq 2 \quad \forall v \in V_S \quad , \quad \deg^-(v) \leq 1 \quad \forall v \in V_S . \quad (8)$$

Linear assignment problem (LAP) We use the established two-step LAP as described by Jaqaman *et al.* [19], implemented in [9]. In the first step, linear chains are formed, which are connected to full cell lineages in a second step. We set a maximum linking distance adapted to the respective dataset, and use the default values for all other hyperparameters.

Integer linear program (ILP) We solve a global ILP with all detections as graph vertices and associations $\{\bar{a}_{i'j'}\}$ with $t_{j'} - t_{i'} = 1$ as edges. The formulation enforces the biological constraints in Eq. (8), as described in [31]. We set the parameters of the ILP, *i.e.* the linear weights of different classes of costs, to values that balance the likelihoods of appearance, disappearance and divisions of cells.

2.5 Implementation details

TRACKASTRA can be trained on a single GPU (*e.g.* Nvidia RTX 4090) for prototypical 2D cell tracking datasets. The transformer is implemented in *PyTorch*. We set window size $s = 6$, embedding dimension $d = 256$, number of encoder and decoder attention layers $L = 6$, the maximum number of tokens per window $|D| = 2048$, and batch size 8. As shallow object features z_i we use the mean intensity, the object area and the inertia tensor of the object region [18]. We apply the following data augmentations jointly to all frames in a window: flips, shifts, rotations, shear, scaling, intensity shifting and scaling, and temporal subsampling. The augmentations are applied directly to the object features. TRACKASTRA association prediction scales to videos with thousands of objects, *e.g.*, inference runs at ~ 1 FPS on a single GPU for 2k objects per frame. The ILP is implemented in *Motile* [10].

3 Experiments

In the following we present tracking experiments on three datasets with varying characteristics, preceded by the employed evaluation measures.

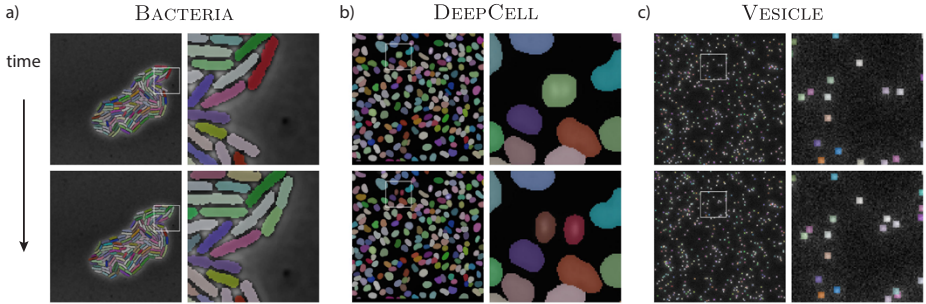


Fig. 2: Cell tracking datasets evaluated in the experiments section. a) BACTERIA dataset from [55] that shows dense colonies of growing and dividing bacteria. b) DEEPCELL dataset of moving and dividing cells with labeled nuclei (DynamicNuclearNet from [42]). c) VESICLE dataset from the ISBI particle tracking challenge [5] that shows synthetically generated images of fluorescently labeled particles.

3.1 Metrics

We report absolute errors per video for multiple elementary error types: False positive (FP) edges, false negative (FN) edges, as well as false positive and false negative divisions. To obtain an aggregated error measure, we further report the Acyclic Oriented Graphs Matching (AOGM) measure [32], which accounts for the number of operations to transform a predicted graph into a reference ground truth (GT) graph. To start, predicted detections are matched to reference detections if they cover more than half of a reference detection’s area. After that, the following operations are performed in the given order to transform the predicted graph into the reference graph:

1. *NS*: (*node split*) Split predicted node that false merged multiple reference nodes, and delete the incident edges of the original node (does not occur when linking GT detections).
2. *FN*: Add FN node.
3. *FP*: Delete FP node and its incident edges (does not occur when linking GT detections).
4. *ED*: Delete FP edge.
5. *EA*: Add FN edge.
6. *EC*: Change semantics of predicted edge (linear chain link *vs.* division link).

AOGM is then defined as the weighted sum of all operations (following [53])

$$\text{AOGM} = 5 \cdot |NS| + 10 \cdot |FN| + 1 \cdot |FP| + 1 \cdot |ED| + 1.5 \cdot |EA| + 1 \cdot |EC|, \quad (9)$$

which represents the quality of a cell tracking solution for a given video in terms of errors to tolerate, in contrast to the usual dominant fraction of simple correctly assigned links. Note that cell divisions are counted implicitly in AOGM. We also report TRA, which is the normalized version of AOGM commonly used in cell

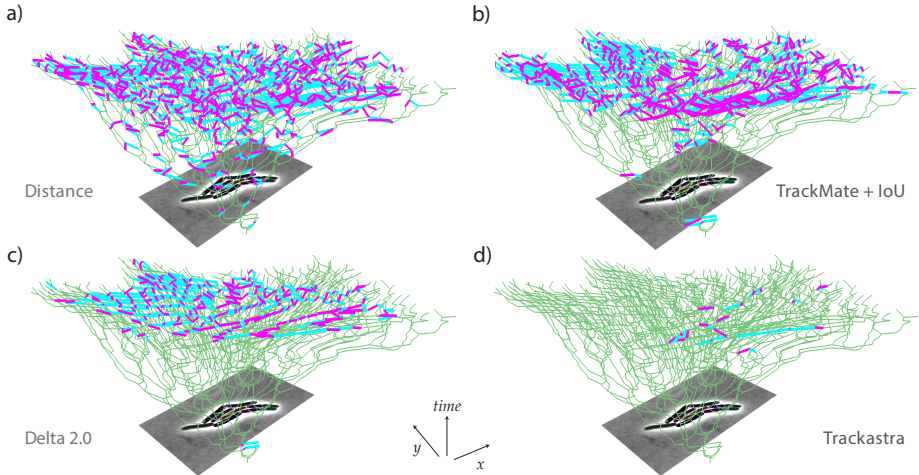


Fig. 3: Error trees on a challenging BACTERIA test video. Time on the vertical axis, edges colored as true positive (green), false positive (magenta) and false negative (cyan).

tracking [53]

$$\text{TRA} = 1 - \frac{\min(\text{AOGM}, \text{AOGM}_0)}{\text{AOGM}_0}, \quad 0 \leq \text{TRA} \leq 1, \quad (10)$$

where AOGM_0 is the AOGM of an empty predicted graph. Finally, when using predicted segmentations, we also report AOGM^+ , which removes the constant penalty incurred by missing detections.

3.2 Bacteria colony tracking (BACTERIA)

Here we use one of the largest public bacteria tracking datasets [55]. This dataset (denoted BACTERIA) consists of 39 videos of six different types of bacteria colonies, containing roughly 100k cells with full segmentation and tracking annotations and showing roughly 9k divisions. This dataset is challenging, as objects are densely packed and the colonies grow and divide, leading to large displacements between frames (*cf.* Fig. 2a). We split it into 31 training, two validation and six test videos, and compare our approach on ground truth detections against various baselines, including DELTA 2.0 [37] that was explicitly created for bacterial colony tracking (Tab. 1). As expected, simple Euclidean-distance-based tracking achieves poor scores, even if a powerful integer linear programming (ILP) linker is used, since the frame rate in this dataset is low and the objects move notably from frame to frame (*cf.* Fig. 2a, Fig. 3a). A standard procedure for bacteria tracking using the accordingly configured and widely used cell tracking tool TrackMate [50] also leads to poor results, since the movements are often too large even for a generously configured intersection-over-union-based tracker (*cf.* Fig. 3b). The state-of-the-art deep-learning-based bacteria tracking algorithm DELTA 2.0 [37],

Table 1: Tracking results on BACTERIA (using ground truth detections). The test set consists of six videos with a mean of 2912 edges and 236 divisions. \downarrow -metrics report mean absolute errors per video. We show results for three runs per model. The last two rows show results of the general model (*cf.* Sec. 3.5).

Method	Linking	TRA \uparrow	AOGM \downarrow	FP edges \downarrow	FN edges \downarrow	FP divs \downarrow	FN divs \downarrow
Distance	greedy	0.944	1511	447	451	298	194
Distance	ILP	0.962	1065	303	303	165	193
TrackMate (overlap) [50]	LAP	0.957	872	256	292	77	50
DELTA 2.0 [37]	greedy	0.996	118	40	43	12	4
TRACKASTRA (points only)	ILP	0.995	136	39	39	20	29
TRACKASTRA	greedy	0.999	36	10	12	5	3
TRACKASTRA	ILP	0.999	23	7	8	2	3
TRACKASTRA-general	greedy	0.999	29	8	9	4	3
TRACKASTRA-general	ILP	0.999	19	6	7	2	3

trained on the same dataset as TRACKASTRA, reduces the total AOGM notably, as shown in Fig. 3c. Surprisingly, training TRACKASTRA using only center point coordinates as features (*i.e.* without any appearance information) performs only slightly worse than DELTA 2.0, which in turn does make use of both the grayscale images and the segmentation masks. This highlights the strong cues that object motion alone can provide for biological datasets. When training TRACKASTRA with shallow object features, already using greedy linking based on the predicted associations reduces the total errors (AOGM) by $\sim 70\%$ compared to the state of the art (Tab. 1, Fig. 3d, Vid. S1). When adding an ILP linker, we are further reducing the total error to an almost perfect tracking result (AOGM = 23 *vs.* 118 for DELTA 2.0).

3.3 Nuclei tracking (DEEPCELL)

Next, we evaluate TRACKASTRA on DynamicNuclearNet [42], the largest publicly available dataset for cell nuclei tracking. This dataset (here denoted DEEPCELL) contains 130 videos of fluorescently labeled nuclei from five different cell types, containing roughly 600k cells with full segmentation and tracking annotations and showing roughly 2k divisions. Due to the comparatively lower rate of cell divisions, we choose to reduce the impact of divisions in the loss reweighting W by setting $\lambda_{\text{div}} = 2$ for this dataset. We adhere to the training-validation-test split as defined in [42]. We train a TRACKASTRA model based on ground truth segmentations, as well as a model based on the segmentations from the Caliban pipeline [42] (*cf.* Tab. 2). We compare our approach to two recent methods that learn linking costs explicitly [1, 42], and take the scores reported in [42].

Using ground truth segmentations TRACKASTRA generally outperforms Caliban [42], the state-of-the-art model that is specifically tuned for DEEPCELL (*cf.* Tab. 2). Specifically the total number of errors (AOGM) is more than halved when using an ILP (AOGM = 7.9 *vs.* 18.1), and shows already a notable reduction when using a simple greedy linker (AOGM = 11.9 *vs.* 18.1, *cf.* Vid. S2). However,

Table 2: Tracking results on DEEPCELL. The test set consists of 12 videos with a mean of 4018 edges and 15 divisions. Scores for Baxter, Caliban and CellTrackerGNN as evaluated in [42]. Metrics represent mean per test video. AA: Association accuracy [13]. AOGM+ removes the constant penalty due to missing detections (498.6).

Costs	Linking	Ground truth objects				Caliban segmentations [42]				
		AOGM ↓	TRA ↑	Div F1 ↑	AA ↑	AOGM+ ↓	AOGM ↓	TRA ↑	Div F1 ↑	AA ↑
Baxter [30]	greedy	-	0.997	0.72	1.00	-	-	0.987	0.60	0.98
CellTrackerGNN [1]	greedy	128.3	0.999	0.18	0.93	204.3	702.9	0.988	0.13	0.89
Caliban [42]	LAP	18.1	1.000	0.97	0.99	75.4	574.0	0.991	0.92	0.97
TRACKASTRA	greedy	11.9	1.000	0.90	1.00	93.2	591.8	0.991	0.71	0.96
TRACKASTRA	ILP	7.9	1.000	0.94	1.00	66.8	565.4	0.991	0.79	0.98
TRACKASTRA-general	greedy	7.4	1.000	0.96	1.00	67.8	566.4	0.991	0.59	0.98
TRACKASTRA-general	ILP	5.8	1.000	0.94	1.00	66.6	565.2	0.991	0.65	0.98

Table 3: Out-of-domain results on HELA. Specialized models trained only on DEEPCELL (using ground truth detections), TRACKASTRA-general trained on a diverse dataset. The test set contains two videos with a mean of 16835 edges and 151 divisions.

Method	Linking	TRA ↑	AOGM ↓	FP edges ↓	FN edges ↓	FP divs ↓	FN divs ↓
Caliban [42]	LAP	0.994	931	36	275	61	151
TRACKASTRA	greedy	0.999	254	52	94	24	41
TRACKASTRA	ILP	0.999	190	72	48	35	17
TRACKASTRA-general	greedy	0.999	145	48	25	22	4
TRACKASTRA-general	ilp	1.000	96	47	24	16	4

Caliban slightly outperforms TRACKASTRA in terms of Division F1, which might be due to explicit modeling of division events in Caliban. We additionally show in Tab. 2 results of CellTrackerGNN [1] as reported in [42], which are worse than both Caliban and TRACKASTRA. Note that the available CellTrackerGNN models have been trained on datasets from the Cell Tracking Challenge [34], which might explain their worse performance on DEEPCELL.

On predicted segmentations (using the provided masks from [42]), our best in-domain model again outperforms Caliban in terms of total errors after removing the constant penalty for missing detections ($\text{AOGM}+ = 66.8$ *vs.* 75.4, *cf.* Tab. 2). However, TRACKASTRA is again not able to reach the Division F1 of Caliban. As before, CellTrackerGNN models have been trained on datasets from the Cell Tracking Challenge, and are therefore not performing well on DEEPCELL.

To probe the out-of-domain capabilities of TRACKASTRA, we take a model trained on DEEPCELL and apply it to a similar dataset of fluorescently tagged nuclei of HeLa cells from the Cell Tracking Challenge (denoted HELA), refer to Tab. 3. Interestingly, TRACKASTRA significantly outperforms Caliban ($\text{AOGM} = 190$ *vs.* 931), potentially because of its use of only shallow input features (such as positions and basic shape descriptors), which likely helps the model prevent overfitting. Notably, the general model (TRACKASTRA-general, *cf.* Sec. 3.5) trained on a diverse dataset performs substantially better than the specialized models.

Table 4: Tracking results on the VESICLE dataset from the ISBI particle tracking challenge (using GT detections). For a definition of the metrics see [5].

Density	Method	Vesicles (low)			Vesicles (mid)			Vesicles (high)		
		$\alpha \uparrow$	$\beta \uparrow$	$JSC_\theta \uparrow$	$\alpha \uparrow$	$\beta \uparrow$	$JSC_\theta \uparrow$	$\alpha \uparrow$	$\beta \uparrow$	$JSC_\theta \uparrow$
low	LAP [8, 19]	0.953	0.947	0.979	0.753	0.703	0.704	0.568	0.490	0.515
	KF [22]	0.937	0.924	0.959	0.673	0.609	0.787	0.477	0.389	0.643
	MoTT [58]	0.926	0.891	0.925	0.800	0.733	0.874	0.652	0.544	0.748
	TRACKASTRA	0.945	0.936	0.972	0.798	0.757	0.895	0.672	0.602	0.806

3.4 ISBI particle tracking challenge

Here we assess how well TRACKASTRA performs for single molecule tracking, a domain closely related to cell tracking where, however, objects do not divide. Specifically, we use the challenging VESICLE dataset from the ISBI particle tracking challenge [5] (*cf.* Fig. 2c). This gives us the chance to compare TRACKASTRA against a state-of-the-art transformer-based tracking approach, MoTT [58], as well as classical methods such as the Kalman filter (KF) [22] and the linear assignment problem (LAP) [8, 19] (*cf.* Tab. 4). Interestingly, even in this context TRACKASTRA outperforms MoTT for the case of high vesicle density when using ground truth detections in terms of the association accuracy (α, β) as well as Jaccard similarity coefficient (JSC_θ), which is a measure of the overlap between the predicted and ground truth tracks (Tab. 4). This highlights the general applicability of our approach, as no adjustments to the model or the training procedure were required.

3.5 Multi-domain general TRACKASTRA model

We train a TRACKASTRA model on a diverse dataset, including all data used in the experiments described before (Secs. 3.2 to 3.4), as well as additional data from the Cell Tracking Challenge [34] and from [11, 23, 43, 56]. We evaluate this model (denoted TRACKASTRA-general) on the BACTERIA and DEEPCELL test sets. Notably, on both BACTERIA and DEEPCELL, the general model outperforms the specialized single-domain TRACKASTRA models (*cf.* Tabs. 1 and 2). Furthermore, when applying the general model to the out-of-domain HELA dataset (which was not part of any training set), it outperforms the specialized TRACKASTRA model trained on DEEPCELL, demonstrating the importance of a large, diverse training set for out-of-domain tracking performance (*cf.* Tab. 3).

Finally, we also submit a single TRACKASTRA-general model to the cell linking benchmark of the Cell Tracking Challenge (7th edition, 2024). Our model performs best overall according to structural as well as biological metrics across 13 diverse datasets, as well as best individually for roughly half of the datasets [33].

3.6 Ablations

Parental softmax: Here we explore whether the proposed parental softmax improves performance for dividing objects. Specifically, we compare in Tab. 5

Table 5: Ablation for parental softmax on BACTERIA (using GT detections). Test set as in Tab. 1. We report results for two runs per model.

Linking	Parental softmax	AOGM ↓	FP edges ↓	FN edges ↓	FP divs ↓	FN divs ↓
greedy	✗	28.4	7.6	9.5	4.6	3.4
greedy	✓	23.0	5.6	8.4	3.8	1.8
ILP	✗	18.8	5.4	6.2	2.4	3.1
ILP	✓	14.8	4.2	5.2	1.8	1.8

results on BACTERIA with and without using parental softmax, *i.e.* when using only sigmoid normalization by setting $\lambda = 1$ and removing the first parental loss component in Eq. (5). We find that the parental softmax reduces the number of errors (AOGM) by $\sim 20\%$ for both a greedy and an ILP linker on BACTERIA, demonstrating the positive impact of the parental softmax on the tracking performance for dividing objects.

Transformer size: Here we vary the number L of attention layers in both the encoder and decoder, as well as the embedding dimension d , using slightly smaller TRACKASTRA models for computational efficiency (*cf.* Fig. 4a,b). As expected, using no attention layers ($L = 0$), *i.e.* directly predicting the association matrix with the two projection heads leads to poor performance, confirming the importance of attention across all detections. Furthermore, using more than $L = 6$ layers does not lead to a notable improvement.

Window size: We run experiments with different window sizes s , *i.e.* different temporal context available to the model (*cf.* Fig. 4c). When using only $s = 2$ we find that the performance substantially deteriorates, whereas window sizes $s \in (3, 6)$ all lead to good results, demonstrating that a relatively small temporal context is already enough for decent tracking results.

Other components: Finally, we ablate basic components of TRACKASTRA in Tab. 6. Removing positional encodings as well as the basic shape descriptors from TRACKASTRA increases errors notably, while the decrease in performance when replacing the tailored ILP optimizer with a simple greedy linker is less pronounced.

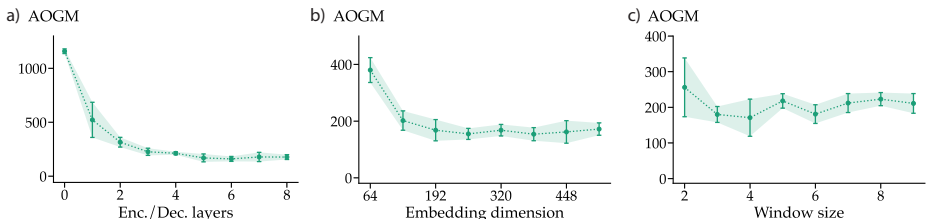
**Fig. 4:** Ablations on BACTERIA (using ground truth detections) using only center points as features, and a LAP linker. Lower is better. We show results for three runs per model.

Table 6: Ablations on BACTERIA (using GT detections). Test set as in Tab. 1. We show results for three runs per model.

Rel. pos. encoding	Object features	ILP	TRA \uparrow	AOGM \downarrow	FP edges \downarrow	FN edges \downarrow	FP divs \downarrow	FN divs \downarrow
		\checkmark	0.989	314	72	88	35	44
			0.992	235	72	73	28	42
\checkmark			0.994	164	42	47	24	30
\checkmark		\checkmark	0.995	136	39	39	20	29
\checkmark	\checkmark		0.999	36	10	12	5	3
\checkmark	\checkmark	\checkmark	0.999	23	7	8	2	3

4 Discussion

We presented TRACKASTRA, a robust method to track dividing objects such as cells using a powerful and scalable transformer architecture. TRACKASTRA uses only object positions and shallow object features as input, making it readily applicable to new scenarios whenever a domain specific detection method is already available. We demonstrate that TRACKASTRA performs well across different imaging modalities and biological model systems that share the particularity of dividing, yet non-fusing objects. Additionally, we show that a single model trained on a large cross-modality dataset is able to generalize well to datasets from various different domains, which is crucial for practitioners. Finally, TRACKASTRA has the potential for tracking other types of dividing objects, such as icebergs in satellite imagery. An important limitation of the presented approach is that it currently does not correct faulty detection inputs, but we anticipate that training a detection model and TRACKASTRA end-to-end will address this issue. Furthermore, we currently do not make use of the pairwise association predictions from non-adjacent time frames, which we expect to be beneficial for circumventing faulty detections, for example for track gap closing in case of missing detections. Moreover, while the presented results are limited to 2D datasets, TRACKASTRA is expected to scale well to 3D datasets since the architecture does not require to process dense images. Overall our work demonstrates the potential that pure transformer-based architectures can hold for the field of cell tracking.

Acknowledgements

We thank Arlo Sheridan, Talmo Pereira, Uwe Schmidt and Albert Dominguez for helpful discussions, Morgan Schwartz for releasing the code for the Caliban benchmarks, Simon van Vliet and Johannes Seiffarth for creating large-scale annotated datasets, and the EPFL School of Life Sciences ELISIR program and CARIGEST SA for their generous funding. Additionally we would like to acknowledge the Janelia Trackathon 2023 and the resulting metrics library *traccuracy*.

References

1. Ben-Haim, T., Riklin Raviv, T.: Graph neural network for cell tracking in microscopy videos. In: ECCV, pp. 610–626 (2022)
2. Bragantini, J., Lange, M., Royer, L.: Large-Scale Multi-Hypotheses Cell Tracking Using Ultrametric Contours Maps. arXiv:2308.04526 (2023)
3. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-End Object Detection with Transformers. In: ECCV, pp. 213–229 (2020)
4. Chen, X., Yan, B., Zhu, J., Wang, D., Yang, X., Lu, H.: Transformer tracking. In: CVPR. pp. 8126–8135 (2021)
5. Chenouard, N., Smal, I., De Chaumont, F., Maška, M., Sbalzarini, I.F., Gong, Y., Cardinale, J., Carthel, C., Coraluppi, S., Winter, M., et al.: Objective comparison of particle tracking methods. *Nature Methods* **11**(3), 281–289 (2014)
6. Chu, P., Wang, J., You, Q., Ling, H., Liu, Z.: Transmot: Spatial-temporal graph transformer for multiple object tracking. In: WACV. pp. 4870–4880 (2023)
7. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021)
8. Ershov, D., Phan, M.S., Pylvänäinen, J.W., Rigaud, S.U., Le Blanc, L., Charles-Orszag, A., Conway, J.R., Laine, R.F., Roy, N.H., Bonazzi, D., et al.: TrackMate 7: integrating state-of-the-art segmentation algorithms into tracking pipelines. *Nature Methods* **19**(7), 829–832 (2022)
9. Fukai, Y.T., Kawaguchi, K.: LapTrack: linear assignment particle tracking with tunable metrics. *Bioinformatics* **39**(1) (2023)
10. Funke, J., Lambert, T., Malin-Mayor, C., Gallusser, B., Jug, F., Pascual Ramos, A.C.: Motile: Multi-object tracker using integer linear equations (2023), <https://github.com/funkelab/motile>
11. Funke, J., Mais, L., Champion, A., Dye, N., Kainmueller, D.: A benchmark for epithelial cell tracking. In: ECCV Workshops (2018)
12. Hayashida, J., Bise, R.: Cell tracking with deep learning for cell detection and motion estimation in low-frame-rate. In: MICCAI. pp. 397–405 (2019)
13. Hayashida, J., Nishimura, K., Bise, R.: MPM: Joint representation of motion and position map for cell tracking. In: CVPR. pp. 3823–3832 (2020)
14. Hayashida, J., Nishimura, K., Bise, R.: Consistent Cell Tracking in Multi-frames with Spatio-Temporal Context by Object-Level Warping Loss. In: WACV. pp. 1759–1768 (2022)
15. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: ICCV. pp. 2961–2969 (2017)
16. Hirsch, P., Epstein, L., Guignard, L.: Mathematical and bioinformatic tools for cell tracking. In: Cell Movement in Health and Disease, pp. 341–361 (2022)
17. Hirsch, P., Malin-Mayor, C., Santella, A., Preibisch, S., Kainmueller, D., Funke, J.: Tracking by Weakly-Supervised Learning and Graph Optimization for Whole-Embryo *C. elegans* lineages. In: MICCAI. pp. 25–35 (2022)
18. Jähne, B.: Spatio-temporal image processing: theory and scientific applications. chap. 8: Tensor Methods (1993)
19. Jaqaman, K., Loerke, D., Mettlen, M., Kuwata, H., Grinstein, S., Schmid, S.L., Danuser, G.: Robust single-particle tracking in live-cell time-lapse sequences. *Nature Methods* **5**(8), 695–702 (2008)
20. Jug, F., Levinkov, E., Blasse, C., Myers, E.W., Andres, B.: Moral lineage tracing. In: CVPR. pp. 5926–5935 (2016)

21. Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., et al.: Highly accurate protein structure prediction with alphafold. *Nature* **596**(7873), 583–589 (2021)
22. Kalman, R.E., et al.: A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* **82**(1), 35–45 (1960)
23. Ker, D.F.E., Eom, S., Sanami, S., Bise, R., Pascale, C., Yin, Z., Huh, S.i., Osuna-Highley, E., Junkers, S.N., Helfrich, C.J., et al.: Phase contrast time-lapse microscopy datasets with automated and manual cell tracking annotations. *Scientific Data* **5** (2018)
24. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment Anything. In: ICCV. pp. 4015–4026 (2023)
25. Leal-Taixé, L., Milan, A., Reid, I., Roth, S., Schindler, K.: Motchallenge 2015: Towards a benchmark for multi-target tracking. arXiv:1504.01942 (2015)
26. Lindenberger, P., Sarlin, P.E., Pollefeys, M.: Lightglue: Local feature matching at light speed. In: ICCV. pp. 17627–17638 (2023)
27. Löffler, K., Mikut, R.: Embedtrack—simultaneous cell segmentation and tracking through learning offsets and clustering bandwidths. *IEEE Access* **10**, 77147–77157 (2022)
28. Lu, D., Xie, Q., Wei, M., Gao, K., Xu, L., Li, J.: Transformers in 3d point clouds: A survey. arXiv:2205.07417 (2022)
29. Luo, W., Xing, J., Milan, A., Zhang, X., Liu, W., Kim, T.K.: Multiple object tracking: A literature review. *Artificial Intelligence* **293**, 103448 (2021)
30. Magnusson, K.E., Jaldén, J., Gilbert, P.M., Blau, H.M.: Global linking of cell tracks using the viterbi algorithm. *IEEE Trans. on Medical Imaging* **34**(4), 911–929 (2014)
31. Malin-Mayor, C., Hirsch, P., Guignard, L., McDole, K., Wan, Y., Lemon, W.C., Kainmueller, D., Keller, P.J., Preibisch, S., Funke, J.: Automated reconstruction of whole-embryo cell lineages by learning from sparse annotations. *Nature Biotechnology* **41**(1), 44–49 (2023)
32. Matula, P., Maška, M., Sorokin, D.V., Matula, P., Ortiz-de Solórzano, C., Kozubek, M.: Cell Tracking Accuracy Measurement Based on Comparison of Acyclic Oriented Graphs. *PLOS ONE* (2015)
33. Maška, M., Cunha, A., Meijering, E., Muñoz-Barrutia, A., Riklin Raviv, T., Stegmaier, J., Uhlmann, V., Ortiz de Solórzano, C., Kozubek, M.: Cell tracking challenge - cell linking benchmark. <https://celltrackingchallenge.net/latest-clb-results/>, accessed: 2024-07-15
34. Maška, M., Ulman, V., Delgado-Rodriguez, P., Gómez-de Mariscal, E., Nečasová, T., Guerrero Peña, F.A., Ren, T.I., Meyerowitz, E.M., Scherr, T., Löffler, K., Mikut, R., et al.: The Cell Tracking Challenge: 10 years of objective benchmarking. *Nature Methods* **20**, 1010–1020 (2023)
35. Meinhardt, T., Feiszli, M., Fan, Y., Leal-Taixe, L., Ranjan, R.: Novis: A case for end-to-end near-online video instance segmentation. arXiv:2308.15266 (2023)
36. Meinhardt, T., Kirillov, A., Leal-Taixe, L., Feichtenhofer, C.: Trackformer: Multi-object tracking with transformers. In: CVPR. pp. 8844–8854 (2022)
37. O'Connor, O.M., Alnahhas, R.N., Lugagne, J.B., Dunlop, M.J.: Delta 2.0: A deep learning pipeline for quantifying single-cell spatial and temporal dynamics. *PLOS Computational Biology* **18**(1) (2022)
38. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: CVPR. pp. 779–788 (2016)
39. Schiegg, M., Hanslovsky, P., Kausler, B.X., Hufnagel, L., Hamprecht, F.A.: Conservation tracking. In: ICCV. pp. 2928–2935 (2013)

40. Schmidt, U., Weigert, M., Broaddus, C., Myers, G.: Cell detection with star-convex polygons. In: MICCAI. pp. 265–273 (2018)
41. Schuster, S., Vernaza, P., Choi, W., Chandraker, M.: Deep network flow for multi-object tracking. In: CVPR. pp. 6951–6960 (2017)
42. Schwartz, M.S., Moen, E., Miller, G., Dougherty, T., Borba, E., Ding, R., Graf, W., Pao, E., Valen, D.V.: Caliban: Accurate cell tracking and lineage construction in live-cell imaging experiments with deep learning. *bioRxiv* (2023)
43. Seiffarth, J., Blöbaum, L., Löffler, K., Scherr, T., Grünberger, A., Scharr, H., Mikut, R., Nöh, K.: Data for - tracking one in a million: Performance of automated tracking on a large-scale microbial data set (2022), <https://doi.org/10.5281/zenodo.7260137>
44. Soelistyo, C.J., Ulicna, K., Lowe, A.R.: Machine learning enhanced cell tracking. *Frontiers in Bioinformatics* **3** (2023)
45. Stringer, C., Wang, T., Michaelos, M., Pachitariu, M.: Cellpose: a generalist algorithm for cellular segmentation. *Nature Methods* **18**(1), 100–106 (2021)
46. Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., Liu, Y.: Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing* **568**, 127063 (2024)
47. Sugawara, K., Çevrim, Ç., Averof, M.: Tracking cell lineages in 3d by incremental deep learning. *eLife* (2022)
48. Sulston, J.E., Schierenberg, E., White, J.G., Thomson, J.N.: The embryonic cell lineage of the nematode *Caenorhabditis elegans*. *Dev. Biology* **100**(1), 64–119 (1983)
49. Sun, P., Cao, J., Jiang, Y., Zhang, R., Xie, E., Yuan, Z., Wang, C., Luo, P.: Transtrack: Multiple-object tracking with transformer. *arXiv: 2012.15460* (2020)
50. Tinevez, J.Y., Perry, N., Schindelin, J., Hoopes, G.M., Reynolds, G.D., Laplantine, E., Bednarek, S.Y., Shorte, S.L., Eliceiri, K.W.: Trackmate: An open and extensible platform for single-particle tracking. *Nature Methods* **115**, 80–90 (2017)
51. Türetken, E., Wang, X., Becker, C.J., Haubold, C., Fua, P.: Network flow integer programming to track elliptical cells in time-lapse sequences. *IEEE Transactions on Medical Imaging* **36**(4), 942–951 (2016)
52. Ulicna, K., Vallardi, G., Charras, G., Lowe, A.R.: Automated deep lineage tree analysis using a bayesian single cell tracking approach. *Frontiers in Computer Science* **3**, 734559 (2021)
53. Ulman, V., Maška, M., Magnusson, K.E., Ronneberger, O., Haubold, C., Harder, N., Matula, P., Matula, P., Svoboda, D., Radojevic, M., et al.: An objective comparison of cell-tracking algorithms. *Nature Methods* **14**(12), 1141–1152 (2017)
54. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *NeurIPS* (2017)
55. van Vliet, S., Winkler, A.R., Spriewald, S., Stecher, B., Ackermann, M., et al.: Spatially correlated gene expression in bacterial groups: the role of lineage history, spatial gradients, and cell-cell interactions. *Cell Systems* **6**(4), 496–507 (2018)
56. Zargari, A., Lodewijk, G.A., Mashhadi, N., Cook, N., Neudorf, C.W., Araghbidikashani, K., Hays, R., Kozuki, S., Rubio, S., Hrabeta-Robinson, E., et al.: Deepsea is an efficient deep-learning model for single-cell segmentation and tracking in time-lapse microscopy. *Cell Reports Methods* **3**(6) (2023)
57. Zeng, F., Dong, B., Zhang, Y., Wang, T., Zhang, X., Wei, Y.: MOTR: End-to-end multiple-object tracking with transformer. In: ECCV. pp. 659–675 (2022)
58. Zhang, Y., Yang, G.: A motion transformer for single particle tracking in fluorescence microscopy images. In: MICCAI. pp. 503–513 (2023)
59. Zhou, X., Yin, T., Koltun, V., Krähenbühl, P.: Global tracking transformers. In: CVPR. pp. 8771–8780 (2022)