

# CoTracker: It is Better to Track Together


Nikita Karaev<sup>1,2</sup>, Ignacio Rocco<sup>1</sup>, Benjamin Graham<sup>1</sup>, Natalia Neverova<sup>1</sup>,  
Andrea Vedaldi<sup>1</sup>, and Christian Rupprecht<sup>2</sup>

<sup>1</sup> Meta AI

<sup>2</sup> Visual Geometry Group, University of Oxford

<https://co-tracker.github.io/>

nikita@robots.ox.ac.uk



**Fig. 1:** **CoTracker** is a new quasi-dense point tracker. It can track 70k points jointly on a single GPU, exploiting dependencies between different tracks to enhance performance. It has exceptional long-term tracking performance even in the presence of occlusions or when points leave the camera view (examples from DAVIS [43], top row displays object tracks, bottom row shows displacements on a dense regular grid. In frame 30, CoTracker successfully tracks the driver’s head even when it exists the camera view).

**Abstract.** We introduce CoTracker, a transformer-based model that tracks a large number of 2D points in long video sequences. Differently from most existing approaches that track points independently, CoTracker tracks them jointly, accounting for their dependencies. We show that joint tracking significantly improves tracking accuracy and robustness, and allows CoTracker to track occluded points and points outside of the camera view. We also introduce several innovations for this class of trackers, including using token proxies that significantly improve memory efficiency and allow CoTracker to track 70k points jointly and simultaneously at inference on a single GPU. CoTracker is an online algorithm that operates causally on short windows. However, it is trained utilizing unrolled windows as a recurrent network, maintaining tracks for long periods of time even when points are occluded or leave the field of view. Quantitatively, CoTracker substantially outperforms prior trackers on standard point-tracking benchmarks. Code and model weights are available at <https://co-tracker.github.io/>

**Keywords:** Point tracking · Motion estimation · Optical flow


## 1 Introduction

We are interested in estimating point correspondences in videos containing dynamic objects and a moving camera. Establishing point correspondences is an important problem in computer vision with many applications [1, 32, 55, 61]. There are two variants of this problem. In *optical flow*, the goal is to estimate the velocities of all points within a video frame. This estimation is performed jointly for all points, but the motion is only predicted at an infinitesimal distance. In *point tracking*, the goal is to estimate the motion of points over an extended period of time. For efficiency, trackers usually focus on a sparse set of points and treat them as statistically independent. This is the case even for recent techniques such as TAPIR [18] and PIPs++ [65], which employ modern architectures such as transformers and can track points even in the presence of occlusions. However, points often have strong statistical dependencies (*e.g.*, because they belong to the same object), which offers an opportunity for improvement.

We hypothesise that accounting for the dependency between tracked points can significantly improve tracking performance. To explore this idea, we introduce CoTracker, a new tracker that supports joint estimation of a very large number of tracks simultaneously, utilizing a transformer-based architecture and attention between tracks. As we show in ablations, tracking points jointly significantly improves the tracking accuracy, especially when points are occluded. In fact, we find that we can improve tracking quality further by adding *more* points to the tracker than required by the user. These additional *support points* expand the context of the tracker, porting to point tracking the idea of using context, common in visual object tracking [31, 36, 49].

In addition to joint tracking, CoTracker incorporates additional architectural design innovations. The network is a transformer operating in a sliding window fashion on a two-dimensional token representation, where dimensions are time and the set of tracked points [2]. Via suitable self-attention operators, the transformer considers each track as a whole for the duration of a window and can exchange information between tracks, thus exploiting their dependencies. However, attention can be expensive when there are many tracks. Hence, we introduce the concept of proxy tokens to point tracking that act similarly to registers [16, 29, 41] while also reducing the memory complexity. These tokens are processed as if they were a small number of additional tracks and allow the switch from expensive self-attention between tracks to efficient cross-attention between tracks and proxies. In this way, CoTracker can jointly track a *near-dense set of tracks* on a single GPU at inference.

CoTracker is designed as an online tracker that operates on relatively short windows of frames. Within a window, the tracks are initialized with queried points, and the network is tasked with progressively refining these initial estimations through iterative applications of the transformer. The windows partially overlap and communicate, similar to a recurrent network. Each subsequent overlapping window starts with refined predictions of the previous window and updates tracks for the new frames. We optimize the recurrent application of the network via unrolled training, porting this concept from recurrent networks and visual object



**Fig. 2: Non-joint tracking (top row) vs. joint tracking (bottom).** Background points are colored cyan while foreground points are magenta. Non-joint point tracking causes background points to follow object motion (cyan points in later frames), while tracking points together results in better tracks for foreground and background points.

tracking [21, 25] to point tracking. In this way, CoTracker achieves *excellent long-term tracking performance* as well; in particular, joint and recurrent tracking allows points to be tracked through long occlusions.

We train CoTracker on the synthetic TAP-Vid-Kubric [17] and evaluate it on TAP-Vid-{DAVIS, RGB-Stacking}, PointOdyssey [65] and DynamicReplica [34]. Our architecture works well for tracking single points and excels for groups of points, obtaining state-of-the-art tracking performance in several benchmarks compared to prior trackers. In summary, our contributions are the following: (i) we introduce the concept of *joint point tracking* by sharing information between tracked points through an attention mechanism; (ii) we propose to use *support points* to provide additional context for point tracking; (iii) we *reduce the model’s memory complexity* with proxy tokens; and (iv) we propose an *unrolled model training strategy*, which further improves tracking and occlusion accuracy.

## 2 Related work

*Optical flow.* Optical flow approximates dense instantaneous motion. Originally approached by studying the brightness constancy equation [6, 7, 26, 37], starting with FlowNet [19, 28] optical flow has been tackled using deep learning. More recently, DCFlow [60] has introduced the computation of a 4D cost volume, used in most follow-up works [52, 54, 57]. Notable is RAFT [54], which introduced incremental flow updates and inspired several follow-up works [33, 33, 59, 62]. CoTracker also uses 4D cost volumes and iterated updates but is applied to tracking.

Transformers [56] have also been applied to the optical flow problem [27, 47, 63]. Flowformer [27] drew inspiration from RAFT and proposed a transformer-based approach that tokenizes the 4D cost volume. GMFlow [63] replaced the update network with a softmax with self-attention for refinement. Perceiver IO [29] proposed a unified transformer for several tasks, including optical flow.

Optical flow can be applied to tracking by integrating predictions over time, but this approach accumulates errors in the form of drift [23], which motivates developing architectures like CoTracker that track points over longer horizons.

*Multi-frame optical flow.* Several authors [30, 38, 44, 46] have extended optical flow to multiple frames. Early methods used Kalman filtering [12, 20] to encourage temporal consistency. Modern multi-frame methods produce dense flow. RAFT [54] can be applied in a warm-start manner [50, 53] for multi-frame optical flow estimation. However, these methods are not designed for long-term tracking and do not consider points occluded for a long time. VideoFlow [46] extends optical flow to three and five consecutive frames by integrating forward and backward motion features during flow refinement. MFT [42] conducts optical flow estimation between distant frames and chooses the most reliable chain of optical flows. Our tracker scales to produce semi-dense tracks.

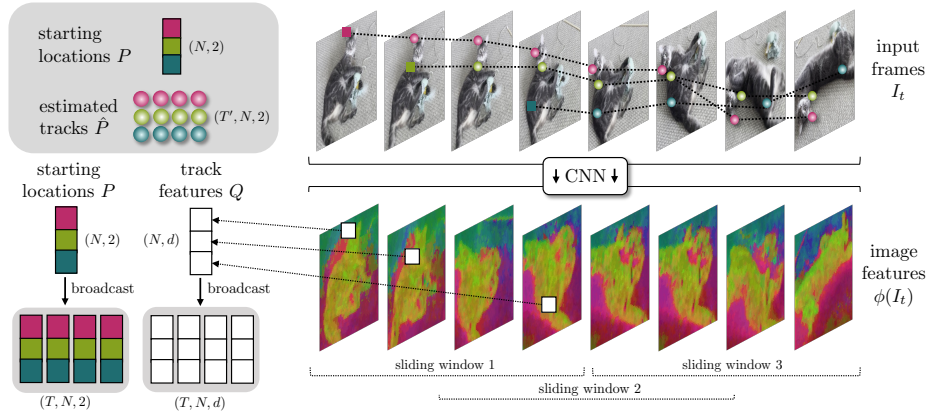
*Visual object tracking.* Before the emergence of deep learning, some authors proposed handcrafted joint trackers [5, 45], but few have considered doing so using deep networks like CoTracker. Our work is weakly related to *multiple object tracking* [10] where tracking through occlusions [64], with changes in appearance [39], and with temporal priors [48] have been extensively studied.

Some recent point trackers have been inspired by visual object tracking that predict the target states by extracting and combining features extracted from source and target image regions [3, 4, 14, 15, 35]. Transformers have also been applied in visual object tracking [11, 13]. However, our focus is the tracking of points, including background points, not objects.

*Tracking any point.* Particle Video [45] introduced the problem of Tracking Any Point in a video (TAP), and PIPs [23] proposed a better model which can track through occlusions more reliably by predicting tracks in a sliding window and restarting them from the last frame where the point was visible. However, PIP cannot track a point beyond the duration of a window. TAP-Vid [17] introduced a new benchmark and a simple baseline for TAP. While the original Particle Video does track points jointly, PIPs and TAP-Vid track points independently (in parallel). Inspired by TAP-Vid [17] and PIPs [23], TAPIR [18] further improved tracking with a two-stage, feed-forward point tracker that uses TAP-Vid-like matching and PIPs-like refinement.

PointOdyssey [65] addressed long-term tracking with PIPs++, a simplified version of PIPs, and introduced another benchmark for long-term tracking. However, PIPs++ still tracks points independently. OmniMotion [58] optimizes a volumetric representation for each video, refining correspondences in a canonical space. However, this approach requires test-time optimization, which, due to its cost, is not suitable for many practical applications, especially online ones.

Trackers and optical flow models are often trained using synthetic datasets [8, 19, 34, 40, 51, 65], as annotating real data can be challenging. Synthetic datasets provide accurate annotations, and training on them has demonstrated the ability to generalize to real-world data [40, 51].



**Fig. 3: CoTracker architecture.** We compute convolutional features  $\phi(I_t)$  for every frame and process them with sliding windows. To initialize track features  $Q$ , we bilinearly sample from  $\phi(I_t)$  with starting point locations  $P$ . Locations  $P$  also serve to initialize estimated tracks  $\hat{P}$ . See Fig. 4 for a visualization of one sliding window.

### 3 CoTracker

Our goal is to track 2D points throughout the duration of a video  $V = (I_t)_{t=1}^T$ , which is a sequence of  $T$  RGB frames  $I_t \in \mathbb{R}^{3 \times H \times W}$ . The goal of the *tracker* is to predict  $N$  point tracks  $P_t^i = (x_t^i, y_t^i) \in \mathbb{R}^2$ ,  $t = t^i, \dots, T$ ,  $i = 1, \dots, N$ , where  $t^i \in \{1, \dots, T\}$  is the time when the track starts. The tracker also predicts the *visibility flag*  $v_t^i \in \{0, 1\}$ , which tells if a point is visible or occluded in a given frame. To make the task definition unambiguous [17], we assume that each point is visible at the start of the track (*i.e.*,  $v_{t^i}^i = 1$ ). The tracker is thus given as input the video  $V$  and the starting locations and times  $(P_{t^i}^i, t^i)_{i=1}^N$  of  $N$  tracks, and outputs an estimate  $(\hat{P}_t^i = (\hat{x}_t^i, \hat{y}_t^i), \hat{v}_t^i)$  of the track locations and visibility for all valid times, *i.e.*,  $t \geq t^i$ .

#### 3.1 Transformer formulation

We implement the tracker as a transformer neural network (Figs. 3 and 4)  $\Psi : G \rightarrow O$ . The goal of this transformer is to improve an initial estimate of the tracks. Tracks are encoded as a grid  $G$  of input tokens  $G_t^i$ , one for each track  $i = 1, \dots, N$ , and time  $t = 1, \dots, T$ . The updated tracks are expressed by a corresponding grid  $O$  of output tokens  $O_t^i$ .

*Image features.* We extract dense  $d$ -dimensional appearance features  $\phi(I_t) \in \mathbb{R}^{d \times \frac{H}{k} \times \frac{W}{k}}$  from each video frame  $I_t$  using a convolutional neural network (trained end-to-end). We reduce the resolution by  $k = 4$  for efficiency. We also consider several scaled versions  $\phi_s(I_t) \in \mathbb{R}^{d \times \frac{H}{k2^{s-1}} \times \frac{W}{k2^{s-1}}}$ , of the features with strides  $s = 1, \dots, S$  and use  $S = 4$  scales. These downscaled features are obtained by applying average pooling to the base features.

*Track features.* The appearance of the tracks is captured by feature vectors  $Q_t^i \in \mathbb{R}^d$  (these are time-dependent to accommodate changes in the track appearance). They are initialized by broadcasting image features sampled at the starting locations and are updated by the neural network, as explained below.

*Spatial correlation features.* In order to facilitate matching tracks to images, we adopt correlation features  $C_t^i \in \mathbb{R}^S$  similar to RAFT [54]. Each  $C_t^i$  is obtained by comparing the track features  $Q_t^i$  to the image features  $\phi_s(I_t)$  around the current estimate  $\hat{P}_t^i$  of the track’s location. Specifically, the vector  $C_t^i$  is obtained by stacking the inner products  $[C_t^i]_{s\delta} = \langle Q_t^i, \phi_s(I_t)[\hat{P}_t^i/k_s + \delta] \rangle$ , where  $s = 1, \dots, S$  are the feature scales and  $\delta \in \mathbb{Z}^2, \|\delta\|_\infty \leq \Delta$  are offsets. The image features  $\phi_s(I_t)$  are sampled at non-integer locations by using bilinear interpolation and border padding. The dimension of  $C_t^i$  is  $(2\Delta + 1)^2 S = 196$  for our choice  $S = 4; \Delta = 3$ .

*Tokens.* The input tokens  $G(\hat{P}, \hat{v}, Q)$  code for position, visibility, appearance, and correlation of the tracks. They are given by the following concatenation of features with added positional encodings:

$$G_t^i = (\hat{P}_t^i - \hat{P}_1^i, \hat{v}_t^i, Q_t^i, C_t^i, \eta(\hat{P}_t^i - \hat{P}_1^i) + \eta'(\hat{P}_1^i) + \eta'(t)). \quad (1)$$

All the components except the last one have been introduced above. The last component is derived from the estimated position: it is the sinusoidal positional encoding  $\eta$  of the track location with respect to the initial location at time  $t = 1$ . We also add encodings  $\eta'$  of the start location  $P_1^i$  and for the time  $t$ , with appropriate dimensions. In fact, we found it beneficial to separately encode the position of points at the first frame and their relative displacement to this frame.

The output tokens  $O(\Delta\hat{P}, \Delta Q)$  contain updates for location and appearance, *i.e.*  $O_t^i = (\Delta\hat{P}_t^i, \Delta Q_t^i)$ .

*Iterated transformer application.* We apply the transformer  $M$  times in order to progressively improve the track estimates. Let  $m = 0, 1, \dots, M$  index the estimate, with  $m = 0$  denoting initialization. Each update computes

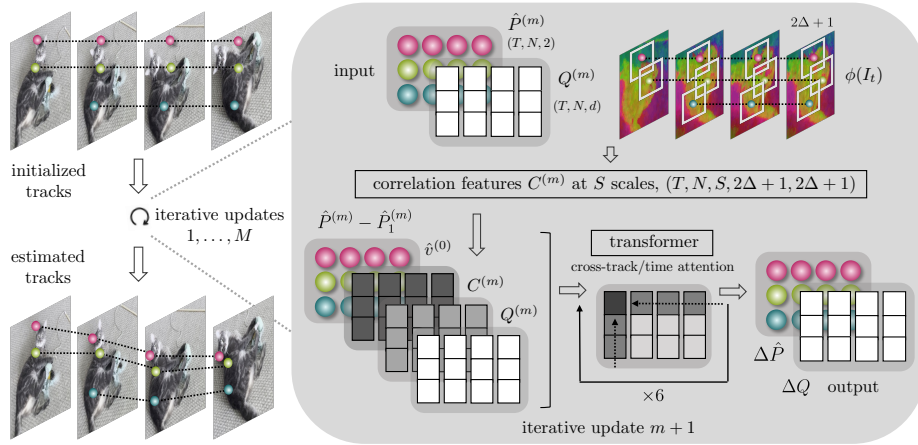
$$O(\Delta\hat{P}, \Delta Q) = \Psi(G(\hat{P}^{(m)}, \hat{v}^{(0)}, Q^{(m)}))$$

and sets  $\hat{P}^{(m+1)} = \hat{P}^{(m)} + \Delta\hat{P}$  and  $Q^{(m+1)} = Q^{(m)} + \Delta Q$ . The visibility mask  $\hat{v}$  is not updated iteratively, but only once after the last transformer application as  $\hat{v}^{(M)} = \sigma(WQ^{(M)})$ , where  $\sigma$  is the sigmoid activation function and  $W$  is a learned matrix of weights. We found that updating the visibility flag iteratively did not improve performance, likely due to the fact that predicting visibility requires predicting an accurate location first.

For  $m = 0$ , the position, visibility and appearance estimates  $\hat{P}^{(0)}, v^{(0)}$  and  $Q^{(0)}$  are initialized by broadcasting their query point’s initial location  $P_{t_i}^i$ , visibility  $v_{t_i}^i = 1$  (meaning visible) and appearance  $\phi(I_{t_i})[P_{t_i}^i/k]$  to all times  $t = 1, \dots, T$ .

### 3.2 Transformer architecture and proxy tokens

The transformer  $\Psi$  interleaves attention layers operating across the *time* and *track dimensions*, respectively. Factorising the attention [2] across time and tracks



**Fig. 4: CoTracker architecture.** Visualization of one sliding window with  $M$  updates. During one iteration, we update point tracks  $\hat{P}^{(m)}$  and track features  $Q^{(m)}$ .  $Q^{(0)}$  is initialized with the initially sampled features  $Q$  for all sliding windows,  $\hat{P}^{(0)}$  with the starting locations for the first window. For other windows,  $\hat{P}^{(0)}$  starts with predictions for frames processed in the preceding sliding window, and with the last predicted positions for the unseen frames. We compute visibility  $\hat{v}$  after the last update  $M$ .

makes the model computationally tractable: the complexity is reduced from  $O(N^2T^2)$  to  $O(N^2 + T^2)$ . However, for very large values of  $N$ , this cost is still prohibitive. We thus propose a new design, in which we introduce  $K$  proxy tracks, where  $K \ll N$  is a hyper-parameter. Proxy tracks are learned fixed tokens that we concatenate to the list of ‘regular’ tracks at the input to the transformer and discard at the output.

For time attention, ‘regular’ and proxy tracks are processed identically. For track attention, however, regular tracks cross-attend the proxies, but not each other, reducing the cost to  $O(NK + K^2 + T^2)$ . We discard the proxy tracks at the output of the transformer.

Proxies are often used to accelerate, *e.g.*, large graph neural networks; here, proxies are learnable query tokens. In computer vision, DETR [9] makes use of learnable queries in a transformer and methods like [29] use them to decode dense outputs. Our proxy tokens are different: They mirror and shadow the computation of the regular tracks, similarly to registers [16].

### 3.3 Windowed inference and unrolled training

An advantage of formulating tracking as the progressive refinement of a window of tracks is the ability to process arbitrarily long videos: It suffices to initialize the next window using partially overlapping tracks from the previous one.

Consider, in particular, a video  $V$  of length  $T' > T$  longer than the maximum window length  $T$  supported by the architecture. To track points throughout the

entire video  $V$ , we split the video in  $J = \lceil 2T'/T - 1 \rceil$  windows of length  $T$ , with an overlap of  $T/2$  frames.<sup>3</sup>

Let the superscript  $(m, j)$  denote the  $m$ -th application of the transformer to the  $j$ -th window. We thus have a  $M \times J$  grid of quantities  $(\hat{P}^{(m,j)}, \hat{v}^{(m,j)}, Q^{(m,j)})$ , spanning transformer iterations and windows. For  $m = 0$  and  $j = 1$ , these are initialized as for the single window case. Then, the transformer is applied  $M$  times to obtain the estimate  $(M, 1)$ . The latter is used to initialize estimate  $(0, 2)$  via broadcasting. Specifically, the first  $T/2$  components of  $\hat{P}^{(0,2)}$  are copies of the last  $T/2$  components of  $\hat{P}^{(M,1)}$ ; the last  $T/2$  components of  $\hat{P}^{(0,2)}$  are instead copies of the last time  $t = T/2 - 1$  from  $\hat{P}^{(M,1)}$ . The same update rule is used for  $\hat{v}^{(0,2)}$ , while  $Q^{(0,j)}$  is always initialized with the features  $Q$  of the track’s staring location. This process is repeated until estimate  $(M, J)$  is obtained.

The windowed transformer essentially operates as a recurrent network, so we train it in an unrolled fashion. Specifically, we optimize the track prediction error summed over iterated transformer applications and windows

$$\mathcal{L}_1(\hat{P}, P) = \sum_{j=1}^J \sum_{m=1}^M \gamma^{M-m} \|\hat{P}^{(m,j)} - P^{(j)}\|, \quad (2)$$


where  $\gamma = 0.8$  discounts early transformer updates. Here,  $P^{(j)}$  contains the ground-truth trajectories restricted to window  $j$  (trajectories which start in the middle of the window are padded backwards). The second loss is the cross entropy of the visibility flags  $\mathcal{L}_2(\hat{v}, v) = \sum_{j=1}^J \text{CE}(\hat{v}^{(M,j)}, v^{(j)})$ . While only a moderate number of windows are used in the loss during training due to the computational cost, at test time, we can unroll the windowed transformer applications arbitrarily, thus, in principle, handling any video length. In combination with joint tracking, unrolling allows the model to track points through occlusions of duration longer than one sliding window and to track points even outside of the camera view.

### 3.4 Support points

CoTracker can take advantage of tracking several points jointly. It is typical for applications to have several points to track, but in some cases, one might be interested in tracking a few or even a single point at inference time.

In these cases, we found it beneficial to track additional *support points* which are not explicitly requested by the user. Moreover, we have found that different configurations of such support points can lead to small differences in performance. We experiment with various configuration types, visualized in Fig. 5. With the “global” strategy, the support points form a regular grid across the whole image. With the “local” strategy, the grid of points is centred around the point we wish to track, thus allowing the model to focus on a neighbourhood of it. Note that these patterns are only considered at inference time and are used to improve the tracker’s accuracy for the target points by incorporating context.

<sup>3</sup> We assume that  $T$  is even. The last window is shorter if  $T/2$  does not divide  $T'$ .



**Fig. 5: Tracked and support points.** We visualize different configurations of points along with their tracks. Points represent the start of each track, which clearly illustrates grids of points in (b) and (c). (a) A single target point in a sequence from TAP-Vid-DAVIS. (b) A global grid of support points. (c) A local grid of support points. (d) Local and global support points. In Tab. 7 we evaluate CoTracker by combining different configurations of support points, as described in Sec. 3.4.


## 4 Experiments

We evaluate CoTracker on several standard real and synthetic tracking benchmarks. After discussing implementation details, datasets and benchmarks, we compare CoTracker to the state of the art (Sec. 4.1) and ablate our design, demonstrating the importance of joint tracking and other innovations (Sec. 4.2)

*Datasets and benchmarks.* We train the model on **TAP-Vid-Kubric** [17], as in [17, 18]. It consists of sequences of 24 frames showing 3D rigid objects falling to the ground under gravity and bouncing, generated using the Kubric engine [22]. Point tracks are selected randomly, primarily on objects and some on the background. Objects occlude each other and the background as they move.

For evaluation, we consider **TAP-Vid-DAVIS** [17], containing 30 real sequences of about 100 frames, and **TAP-Vid-RGB-Stacking** [17] with 50 synthetic sequences with objects being moved by robots of 200–300 frames each. Points are queried on random objects at random times and evaluation assesses both predictions of positions and visibility flags. This uses two evaluation protocols. In the “first” protocol, each point is queried only once in the video, at the first frame where it becomes visible. The tracker is expected to operate causally, predicting positions only for future frames. In the “strided” protocol, points are queried every five frames, and tracking is bidirectional. Given that most trackers (ours, PIPs, PIPs++) are causal, we run them twice, on the video and its reverse. We assess tracking accuracy using the TAP-Vid metrics [17]: Occlusion Accuracy (OA; accuracy of occlusion prediction treated as binary classification),  $\delta_{\text{avg}}^{\text{vis}}$  (fraction of visible points tracked within 1, 2, 4, 8 and 16 pixels, averaged over thresholds), and Average Jaccard (AJ, measuring jointly geometric and occlusion prediction accuracy). Following [17], the  $\delta_{\text{avg}}^{\text{vis}}$  and AJ tracking thresholds are applied after virtually resizing images to  $256 \times 256$  pixels. Note that TAP-Vid benchmarks evaluate only *visible* points.

We also evaluate on **PointOdyssey** [65], a more recent synthetic benchmark dataset for long-term tracking. It contains 100 sequences that are several thousand frames long, with objects and characters moving around the scene. Scenes are much more realistic than TAP-Vid-Kubric. We train and evaluate CoTracker on



**Fig. 6: Qualitative Results.** For PIPs++ (top), many points are incorrectly tracked and end up being ‘stuck’ on the front of the car. TAPIR (middle) works well for visible points but fails to handle occluded points. As soon as a point becomes occluded, it starts moving chaotically around the image. Our CoTracker (bottom) produces cleaner tracks. The tracks are also more ‘linear’ than those of PIPs++ or TAPIR, which is accurate as the primary motion is a homography (the observer does not translate).

PointOdyssey and report  $\delta_{\text{avg}}^{\text{vis}}$  and  $\delta_{\text{avg}}^{\text{occ}}$ ,  $\delta_{\text{avg}}$ . The last two are the same as  $\delta_{\text{avg}}^{\text{vis}}$ , but for occluded and all points, respectively; they can be computed because the dataset is synthetic so the ground-truth positions of invisible points are known. We also report their *Survival* rate, which is the average fraction of video frames until the tracker fails (detected when the tracking error exceeds 50 pixels).

**Dynamic Replica** [34] is a synthetic dataset for 3D reconstruction that contains long-term tracking annotations. It consists of 500 300-frame sequences of articulated models of people and animals. We evaluate models trained on Kubric and measure  $\delta_{\text{avg}}^{\text{vis}}$  and  $\delta_{\text{avg}}^{\text{occ}}$  on the “valid” split of this dataset.

*Implementation details.* Here we provide important implementation details and refer to the sup. mat. for others. We **train** the model for 50,000 iterations on 6,000 TAP-Vid-Kubric sequences of  $T' = 24$  frames using sliding window size  $T = 8$ , which takes 40 hours. We use 32 NVIDIA A100 80GB GPUs with a batch size of 1. Training tracks are sampled preferentially on objects. During training, we construct batches of  $N = 768$  tracks with visible points either in the first or middle frames of the sequence to train the tracker to handle both cases. In a similar manner, we also train a second version of CoTracker on the train split of PointOdyssey, using 128 tracks randomly sampled from sequences of length  $T' = 56$ .

A technical difficulty is that the number of tracks  $N$  varies from window to window, as new points can be added to the tracker at any time. While conceptually, this is handled by simply adding more tokens when needed, in practice, changing the number of tokens makes batching data for training difficult. Instead, we

Method	Train	DAVIS First			DAVIS Strided			RGB-S First		
		AJ $\uparrow$	$\delta_{\text{avg}}^{\text{vis}}$ $\uparrow$	OA $\uparrow$	AJ $\uparrow$	$\delta_{\text{avg}}^{\text{vis}}$ $\uparrow$	OA $\uparrow$	AJ $\uparrow$	$\delta_{\text{avg}}^{\text{vis}}$ $\uparrow$	OA $\uparrow$
TAP-Net [17]	K	33.0	48.6	78.8	38.4	53.1	82.3	—	—	—
OmniMotion [58]	—	52.8	66.9	<u>87.1</u>	51.7	67.5	85.3	<u>69.5</u>	<u>82.4</u>	<b>90.3</b>
PIPs [23]	FT++	42.2	64.8	77.7	52.4	70.0	83.6	—	59.1	—
MFT [42]	K++	47.3	66.8	77.8	56.1	70.8	86.9	—	—	—
PIPs++ [65]	PO	—	69.1	—	—	<u>73.7</u>	—	—	77.8	—
TAPIR [18]	K	<u>56.2</u>	<u>70.0</u>	86.5	<u>61.3</u>	73.6	<u>88.8</u>	60.3	74.1	85.5
CoTracker (Ours)	K	<b>62.2</b>	<b>75.7</b>	<b>89.3</b>	<b>65.9</b>	<b>79.4</b>	<b>89.9</b>	<b>71.6</b>	<b>83.3</b>	<u>89.6</u>

**Table 1: TAP-Vid benchmarks** We compare CoTracker to the best trackers available on TAP-Vid benchmarks utilizing both their “first” and “strided” evaluation protocols. During evaluation, TAPIR and OmniMotion have access to all video frames at once, while we process videos in an online manner using a causal sliding window. This puts online trackers like ours at a disadvantage on long videos with slow motion, such as in RGB-S. To address this, we decrease the frame rate for RGB-S evaluation by keeping every 5th frame. Training data: (K) Kubric, (K++) Kubric+more, (FT) FlyingThings++, (PO) Point Odyssey.

allocate enough tokens for all the tracked points, regardless of which window they are added in, and use masking to ignore tokens that are not yet used.

The video **resolution** can also affect tracking performance. For evaluation on TAP-Vid, we follow their protocol and downsample videos to  $256 \times 256$ , which ensures that no more information than expected is passed to the tracker. However, each tracker has a different native resolution at which it is trained ( $384 \times 512$  for PIPs and CoTracker;  $512 \times 896$  for PIPs++). We thus resize videos from  $256 \times 256$  to the native resolution of each tracker before running it to ensure fairness. For Dynamic Replica and PointOdyssey, we resize videos to the native tracker resolution, as there is no prescribed resolution.

#### 4.1 Comparisons to the State of the Art

We compare CoTracker to state-of-the-art trackers on TAP-Vid-DAVIS and TAP-Vid-RGB-Stacking in Tab. 1. For fairness, we use the same configuration of support points identified in the ablations (see Sec. 4.2) across all benchmarks. We also evaluate CoTracker on PointOdyssey in Tab. 2 and on Dynamic Replica in Tab. 3 by tracking jointly all the target points. CoTracker improves most metrics by substantial margins on the TAP-Vid benchmarks while generalizing well from synthetic Kubric to real DAVIS. CoTracker also improves tracking accuracy  $\delta_{\text{avg}}^{\text{vis}}$  in all cases. On PointOdyssey, despite using a window size of only 8 frames, CoTracker has a better survival rate than PIPs++, which uses a 128-frame sliding window. This shows the power of unrolled training, which trains the model to propagate information across a long series of sliding windows. The gap between CoTracker and other methods in Tab. 2 and Tab. 3 for  $\delta_{\text{avg}}^{\text{occ}}$  is higher than for

Method	PointOdyssey			
	$\delta_{\text{avg}} \uparrow$	$\delta_{\text{avg}}^{\text{vis}} \uparrow$	$\delta_{\text{avg}}^{\text{occ}} \uparrow$	Survival $\uparrow$
TAP-Net [17]	28.4	—	—	18.3
PIPs [24]	27.3	—	—	42.3
PIPs++ <sup>†</sup> [65]	29.0	32.4	18.8	47.0
CoTracker (Ours)	<b>30.2</b>	<b>32.7</b>	<b>24.2</b>	<b>55.2</b>

**Table 2: PointOdyssey.** All methods are trained on PointOdyssey. <sup>†</sup> results for [65] obtained using released code and model.

Method	Dynamic Replica		
	$\delta_{\text{avg}} \uparrow$	$\delta_{\text{avg}}^{\text{vis}} \uparrow$	$\delta_{\text{avg}}^{\text{occ}} \uparrow$
TAP-Net [17]	45.5	53.3	20.0
PIPs [24]	41.0	47.1	21.0
PIPs++ [65]	55.5	64.0	28.5
TAPIR [18]	56.8	66.1	27.2
CoTracker (Ours)	<b>61.6</b>	<b>68.9</b>	<b>37.6</b>

**Table 3: Dynamic Replica.** The gap between CoTracker and other methods in tracking occluded points’ accuracy  $\delta_{\text{avg}}^{\text{occ}}$  is higher than  $\delta_{\text{avg}}^{\text{vis}}$ .

Mode:	DAVIS First		Dynamic Replica			
	AJ $\uparrow$	$\delta_{\text{avg}} \uparrow$	OA	$\delta_{\text{avg}} \uparrow$	$\delta_{\text{avg}}^{\text{vis}} \uparrow$	$\delta_{\text{avg}}^{\text{occ}} \uparrow$
no joint	55.6	70.1	83.0	54.4	62.4	28.8
joint	<b>62.2</b>	<b>75.7</b>	<b>89.3</b>	<b>61.6</b>	<b>68.9</b>	<b>37.6</b>

**Table 4: Importance of joint tracking.** We track a single target point with or without additional support points on DAVIS, and all target points separately or jointly on Dynamic Replica.

Unrolled Training	DAVIS First		
	AJ $\uparrow$	$\delta_{\text{avg}} \uparrow$	OA $\uparrow$
$\times$	44.6	60.5	75.3
$\checkmark$	<b>62.2</b>	<b>75.7</b>	<b>89.3</b>

**Table 5: Unrolled training.** CoTracker is built for sliding window predictions. Using them during training is important.

$\delta_{\text{avg}}^{\text{vis}}$ , which shows that CoTracker excels in tracking occluded points thanks to jointly tracking groups of points.

## 4.2 Ablations

*Tracking together is better.* In Tab. 4, we demonstrate the importance of tracking points jointly, which is the main motivation of CoTracker. This is achieved by removing cross-track attention from the tracker, thus ignoring the dependencies between tracks entirely. For fairness, rather than simply removing the six cross-track attention layers, we maintain the same model size and replace them with twelve time-attention layers, resulting in a comparable model size. On Dynamic Replica, joint tracking improves the accuracy of occluded points  $\delta_{\text{avg}}^{\text{occ}}$   $28.8 \rightarrow 37.6$  (+30.6%) more than the accuracy of visible points  $\delta_{\text{avg}}^{\text{vis}}$   $62.4 \rightarrow 68.9$  (+10.4%), showing the effectiveness of joint tracking in understanding the scene motion, including currently invisible points.

*Importance of unrolled training.* We design CoTracker for tracking in a windowed manner, and in Tab. 5, we thus assess the effect of unrolling windows during training (Sec. 3.3). Switching off unrolled training decreases performance by 18 AJ points. Hence, unrolled training helps to track over long periods of time,  $> 10\times$  longer than the sequences used in training.

Num. proxy tokens	DAVIS First			Max. num. tracks	Time [s]
	AJ $\uparrow$	$\delta_{\text{avg}}$ $\uparrow$	OA $\uparrow$		
0	61.6	75.6	88.3	9.4k	207.3
32	60.2	74.4	88.5	69.2k	26.8
64	<b>62.2</b>	<b>75.7</b>	<b>89.3</b>	69.2k	27.9
128	60.9	74.8	88.4	69.2k	30.1

**Table 6:** The **proxy tokens** allow CoTracker to scale. We train models with different numbers of proxy tokens and show that they help to decrease both memory complexity and inference time. We report the maximum number of tracks that can fit on a 80 GB GPU.


Support:		DAVIS First		
global	local	AJ $\uparrow$	$\delta_{\text{avg}}$ $\uparrow$	OA $\uparrow$
$\times$	$\times$	55.6	70.1	83.0
$\checkmark$	$\times$	56.8	71.2	85.8
$\times$	$\checkmark$	60.4	75.4	87.3
$\checkmark$	$\checkmark$	<b>62.2</b>	<b>75.7</b>	<b>89.3</b>

**Table 7: Support points for TAP-Vid benchmarks.** We track single target points from DAVIS using additional support points for context, as described in Sec. 3.4. Adding points for context always helps, while additional global *and* local context works best.

*Effect of proxy tokens on scalability.* In Tab. 6, we assess the benefits of using proxy tokens during inference. For a fixed memory budget (80 GB), using proxy tokens instead of full self-attention allows us to track  $\times 7.4$  more points than without them; in fact, we can track a whole  $263 \times 263$  grid, which is quasi-dense for the input video resolution. Moreover, proxy tokens also reduce the time complexity of the model and make it  $\times 7$  faster at inference time for the maximum number of tracked points. The number of proxy tokens does not affect scalability but affects performance, with the best results obtained using 64 proxy tokens. In short, *proxy tokens allow to track almost one order of magnitude more tracks with higher accuracy than the naive self-attention.*

*Optimal support point configurations.* We compare the effect of choosing different configurations of support points from Fig. 5 in Tab. 7. In order to do so, we take a single benchmark point at a time and add additional support points (Sec. 3.4) to allow the model to perform joint tracking. This protocol, where a single benchmark point is tracked each time, is also important for fairness when compared to the state-of-the-art. This is because the benchmark points might be biased for objects, which is particularly evident in TAP-Vid-DAVIS (Fig. 5). Passing more than one such biased point at the same time to a joint tracker like ours could help it by revealing the outline of objects to the tracker. By considering one benchmark point at a time, we ensure that no such ground-truth information can leak into it. We use this scheme for the TAP-Vid benchmarks, which explicitly prohibits tracking more than one ground-truth point at a time.

Adding any configurations of support points helps, but the local configuration helps much more than the global, presumably because the model uses dependencies with other points located on the same object. The combination of global and local context works best, likely because the model can track both the camera motion and the object motion in this case.



**Fig. 7: Qualitative Results.** CoTracker can track points outside of the camera view. We visualize point displacements on a dense regular grid using an example from DAVIS [43]. In this example, the camera moves towards a group of people, while points in the background stay behind the camera view.

### 4.3 Limitations

Despite its high performance, CoTracker still occasionally makes tracking mistakes that a human would not make. Since CoTracker is trained on purely synthetic data, it sometimes does not generalize well to complex visual scenes with reflections and shadows. For example, CoTracker tends to track shadows together with the objects that cast them. Depending on the application this may be a desirable property (*e.g.* video editing) or it may not be (*e.g.* motion analysis).

## 5 Conclusions

We have presented CoTracker, a transformer-based point tracker that tracks several points jointly, accounting for their dependencies. CoTracker is state-of-the-art on the standard tracking benchmarks, often by a substantial margin, can track through occlusions and when points leave the field of view, even for hundreds of frames, and can track a very large number of points simultaneously. The transformer architecture is flexible and memory efficient. This allows for the integration of more functionalities in the future, such as 3D reconstruction.

*Acknowledgments* We want to thank Laurynas Karazija for evaluating model efficiency, Luke Melas-Kyriazi and Jianyuan Wang for their paper comments, Roman Shapovalov, Iurii Makarov, Shalini Maiti, and Adam W. Harley for the insightful discussions. Christian Rupprecht was supported by ERC-CoG UNION101001212 and VisualAI EP/T028572/1.

## References

1. Agarwal, S., Furukawa, Y., Snavely, N., Simon, I., Curless, B., Seitz, S.M., Szeliski, R.: Building rome in a day. *Communications of the ACM* **54**(10), 105–112 (2011)
2. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: *Proc. ICML* (July 2021)
3. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fully-convolutional siamese networks for object tracking. In: *Computer Vision–ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part II* 14. pp. 850–865. Springer (2016)
4. Bhat, G., Danelljan, M., Gool, L.V., Timofte, R.: Learning discriminative model prediction for tracking. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2019)
5. Birchfield, S.T., Pundlik, S.J.: Joint tracking of features and edges. In: *Proc. CVPR* (2008)
6. Black, M.J., Anandan, P.: A framework for the robust estimation of optical flow. In: *Proc. ICCV* (1993)
7. Bruhn, A., Weickert, J., Schnörr, C.: Lucas/kanade meets horn/schunck: Combining local and global optic flow methods. *IJCV* **61** (2005)
8. Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: *Proc. ECCV* (2012)
9. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: *Proc. ECCV* (2020)
10. Chen, F., Wang, X., Zhao, Y., Lv, S., Niu, X.: Visual object tracking: A survey. *CVIU* **222** (2022)
11. Chen, X., Yan, B., Zhu, J., Wang, D., Yang, X., Lu, H.: Transformer tracking. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 8126–8135 (June 2021)
12. Chin, T.M., Karl, W.C., Willsky, A.S.: Probabilistic and sequential computation of optical flow using temporal coherence. *IEEE Trans. on Image Processing* **3**(6) (1994)
13. Cui, Y., Jiang, C., Wang, L., Wu, G.: Mixformer: End-to-end tracking with iterative mixed attention. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 13608–13618 (2022)
14. Danelljan, M., Bhat, G., Khan, F.S., Felsberg, M.: Atom: Accurate tracking by overlap maximization. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4660–4669 (2019)
15. Danelljan, M., Bhat, G., Shahbaz Khan, F., Felsberg, M.: Eco: Efficient convolution operators for tracking. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 6638–6646 (2017)
16. Darcet, T., Oquab, M., Mairal, J., Bojanowski, P.: Vision transformers need registers. *arXiv preprint arXiv:2309.16588* (2023)
17. Doersch, C., Gupta, A., Markeeva, L., Recasens, A., Smaira, L., Aytar, Y., Carreira, J., Zisserman, A., Yang, Y.: Tap-vid: A benchmark for tracking any point in a video. *arXiv* (2022)
18. Doersch, C., Yang, Y., Vecerik, M., Gokay, D., Gupta, A., Aytar, Y., Carreira, J., Zisserman, A.: Tapir: Tracking any point with per-frame initialization and temporal refinement (2023)
19. Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., Brox, T.: Flownet: Learning optical flow with convolutional networks. In: *Proc. ICCV* (2015)

20. Elad, M., Feuer, A.: Recursive optical flow estimation—adaptive filtering approach. *J. of Visual Communication and Image Representation* **9**(2) (1998)
21. Girshick, R., Iandola, F., Darrell, T., Malik, J.: Deformable part models are convolutional neural networks. In: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. pp. 437–446 (2015)
22. Greff, K., Belletti, F., Beyer, L., Doersch, C., Du, Y., Duckworth, D., Fleet, D.J., Gnanapragasam, D., Golemo, F., Herrmann, C., et al.: Kubric: A scalable dataset generator. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3749–3761 (2022)
23. Harley, A.W., Fang, Z., Fragkiadaki, K.: Particle video revisited: Tracking through occlusions using point trajectories. In: *Proc. ECCV* (2022)
24. Harley, A.W., Fang, Z., Fragkiadaki, K.: Particle videos revisited: Tracking through occlusions using point trajectories. In: *Proc. ECCV* (2022)
25. Held, D., Thrun, S., Savarese, S.: Learning to track at 100 fps with deep regression networks. In: *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*. pp. 749–765. Springer (2016)
26. Horn, B.K., Schunck, B.G.: Determining optical flow. *Artificial intelligence* **17**(1-3) (1981)
27. Huang, Z., Shi, X., Zhang, C., Wang, Q., Cheung, K.C., Qin, H., Dai, J., Li, H.: Flowformer: A transformer architecture for optical flow. In: *Proc. ECCV* (2022)
28. Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: FlowNet 2.0: Evolution of optical flow estimation with deep networks. In: *Proc. CVPR* (2017)
29. Jaegle, A., Borgeaud, S., Alayrac, J., Doersch, C., Ionescu, C., Ding, D., Koppula, S., Zoran, D., Brock, A., Shelhamer, E., Hénaff, O.J., Botvinick, M.M., Zisserman, A., Vinyals, O., Carreira, J.: Perceiver IO: A general architecture for structured inputs & outputs. In: *Proc. ICLR* (2022)
30. Janai, J., Guney, F., Ranjan, A., Black, M., Geiger, A.: Unsupervised learning of multi-frame optical flow with occlusions. In: *Proc. ECCV* (2018)
31. Jia, X., Lu, H., Yang, M.H.: Visual tracking via adaptive structural local sparse appearance model. In: *2012 IEEE Conference on computer vision and pattern recognition*. pp. 1822–1829. IEEE (2012)
32. Jiang, H., Sun, D., Jampani, V., Yang, M.H., Learned-Miller, E., Kautz, J.: Super slomo: High quality estimation of multiple intermediate frames for video interpolation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 9000–9008 (2018)
33. Jiang, S., Lu, Y., Li, H., Hartley, R.: Learning optical flow from a few matches. In: *Proc. CVPR* (2021)
34. Karaev, N., Rocco, I., Graham, B., Neverova, N., Vedaldi, A., Ruppel, C.: Dynamicstereo: Consistent dynamic depth from stereo videos. In: *Proc. CVPR* (2023)
35. Li, F., Tian, C., Zuo, W., Zhang, L., Yang, M.H.: Learning spatial-temporal regularized correlation filters for visual tracking. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4904–4913 (2018)
36. Li, Y., Zhu, J., Hoi, S.C.: Reliable patch trackers: Robust visual tracking by exploiting reliable patches. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 353–361 (2015)
37. Lucas, B.D., Kanade, T.: An iterative image registration technique with an application to stereo vision. In: *Proc. IJCAI*. vol. 2 (1981)
38. Luo, J., Wan, Z., Li, B., Dai, Y., et al.: Continuous parametric optical flow. In: *Thirty-seventh Conference on Neural Information Processing Systems* (2023)

39. Matthews, L., Ishikawa, T., Baker, S.: The template update problem. *PAMI* **26**(6) (2004)
40. Mayer, N., Ilg, E., Hausser, P., Fischer, P., Cremers, D., Dosovitskiy, A., Brox, T.: A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In: *Proc. CVPR* (2016)
41. Nagrani, A., Yang, S., Arnab, A., Jansen, A., Schmid, C., Sun, C.: Attention bottlenecks for multimodal fusion. *Advances in Neural Information Processing Systems* **34**, 14200–14213 (2021)
42. Neoral, M., Šerých, J., Matas, J.: Mft: Long-term tracking of every pixel (2023)
43. Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Van Gool, L.: The 2017 davis challenge on video object segmentation. *arXiv* (2017)
44. Ren, Z., Gallo, O., Sun, D., Yang, M.H., Sudderth, E.B., Kautz, J.: A fusion approach for multi-frame optical flow estimation. In: *Proc. WACV* (2019)
45. Sand, P., Teller, S.: Particle video: Long-range motion estimation using point trajectories. *IJCV* **80** (2008)
46. Shi, X., Huang, Z., Bian, W., Li, D., Zhang, M., Cheung, K.C., See, S., Qin, H., Dai, J., Li, H.: Videoflow: Exploiting temporal cues for multi-frame optical flow estimation. *arXiv* (2023)
47. Shi, X., Huang, Z., Li, D., Zhang, M., Cheung, K.C., See, S., Qin, H., Dai, J., Li, H.: Flowformer++: Masked cost volume autoencoding for pretraining optical flow estimation. *arXiv* (2023)
48. Sidenbladh, H., Black, M.J., Fleet, D.J.: Stochastic tracking of 3d human figures using 2d image motion. In: *Proc. ECCV* (2000)
49. Song, Y., Ma, C., Gong, L., Zhang, J., Lau, R.W., Yang, M.H.: Crest: Convolutional residual learning for visual tracking. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2555–2564 (2017)
50. Sui, X., Li, S., Geng, X., Wu, Y., Xu, X., Liu, Y., Goh, R., Zhu, H.: Craft: Cross-attentional flow transformer for robust optical flow. In: *Proc. CVPR* (2022)
51. Sun, D., Vlasic, D., Herrmann, C., Jampani, V., Krainin, M., Chang, H., Zabih, R., Freeman, W.T., Liu, C.: Autoflow: Learning a better training set for optical flow. In: *Proc. CVPR* (2021)
52. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: *Proc. CVPR* (2018)
53. Sun, S., Chen, Y., Zhu, Y., Guo, G., Li, G.: Skflow: Learning optical flow with super kernels. *arXiv* (2022)
54. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: *Proc. ECCV* (2020)
55. Teed, Z., Deng, J.: Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in neural information processing systems* **34**, 16558–16569 (2021)
56. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: *Proc. NeurIPS* (2017)
57. Wang, J., Zhong, Y., Dai, Y., Zhang, K., Ji, P., Li, H.: Displacement-invariant matching cost learning for accurate optical flow estimation. *Proc. NeurIPS* **33** (2020)
58. Wang, Q., Chang, Y.Y., Cai, R., Li, Z., Hariharan, B., Holynski, A., Snavely, N.: Tracking everything everywhere all at once (2023)
59. Xu, H., Yang, J., Cai, J., Zhang, J., Tong, X.: High-resolution optical flow from 1d attention and correlation. In: *Proc. CVPR* (2021)
60. Xu, J., Ranftl, R., Koltun, V.: Accurate optical flow via direct cost volume processing. In: *Proc. CVPR* (July 2017)

61. Yang, G., Vo, M., Neverova, N., Ramanan, D., Vedaldi, A., Joo, H.: Banmo: Building animatable 3d neural models from many casual videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2863–2873 (2022)
62. Zhang, F., Woodford, O.J., Prisacariu, V.A., Torr, P.H.: Separable flow: Learning motion cost volumes for optical flow estimation. In: Proc. CVPR (2021)
63. Zhao, S., Zhao, L., Zhang, Z., Zhou, E., Metaxas, D.: Global matching with overlapping attention for optical flow estimation. In: Proc. CVPR (2022)
64. Zhao, T., Nevatia, R.: Tracking multiple humans in crowded environment. In: Proc. CVPR. vol. 2 (2004)
65. Zheng, Y., Harley, A.W., Shen, B., Wetzstein, G., Guibas, L.J.: Pointodyssey: A large-scale synthetic dataset for long-term point tracking. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 19855–19865 (2023)