

MFT: Long-Term Tracking of Every Pixel

Michal Neoral, Jonáš Šerých, Jiří Matas

CMP Visual Recognition Group, Faculty of Electrical Engineering, Czech Technical University in Prague

{neoramic, serycjon, matas}@fel.cvut.cz

Abstract

We propose *MFT – Multi-Flow dense Tracker* – a novel method for dense, pixel-level, long-term tracking. The approach exploits optical flows estimated not only between consecutive frames, but also for pairs of frames at logarithmically spaced intervals. It selects the most reliable sequence of flows on the basis of estimates of its geometric accuracy and the probability of occlusion, both provided by a pre-trained CNN. We show that MFT achieves competitive performance on the TAP-Vid benchmark, outperforming baselines by a significant margin, and tracking densely orders of magnitude faster than the state-of-the-art point-tracking methods. The method is insensitive to medium-length occlusions and it is robustified by estimating flow with respect to the reference frame, which reduces drift.

1. Introduction

Reliable dense optical flow has a significant enabling potential for diverse computer vision applications, including structure-from-motion, video editing, and augmented reality. Despite the widespread use of optical flow between consecutive frames for motion estimation in videos, generating consistent and dense long-range motion trajectories has been under-explored and remains a challenging task.

A simple baseline method for obtaining point-to-point correspondences in a video, e.g. for augmented reality, concatenates interpolated optical flow to form trajectories of a pixel, i.e. the set of projections of the pre-image of the pixel, for all frames in a sequence. However, such approach suffers from several problems: error accumulation leading to drift, sensitivity to occlusion and non-robustness, since a single poorly estimated optical flow damages the long-term correspondences for future frames. This results in trajectories that quickly diverge and become inconsistent, particularly in complex scenes involving large motions, repetitive patterns and illumination changes. Additionally, concatenated optical flow between consecutive frames cannot recover trajectories after occlusions. Few optical flow approaches estimate occluded regions or uncertainty of esti-




Figure 1. **Overview of the MFT.** MFT tracks a query point (black square) by chaining optical flows. Each chain consists of a previously computed chain from frame 0 up to frame $(t - \Delta)$ (dashed arrow, white dot), and an optical flow vector computed between frames $(t - \Delta)$ and t (solid arrow). MFT forms multiple candidate chains with varying Δ . The best candidate (black dot) is selected according to uncertainty and occlusion scores. This is done in parallel, independently for each pixel in the reference frame.

ated optical flow.

Another baseline approach — matching every frame with the reference — is neither prone to drift nor occlusions, but has other weaknesses. As the pose and illumination conditions change in the sequence, the matching problem becomes progressively more difficult. In the datasets used for evaluation in this paper, match-to-reference performs worse than consecutive frame optical flow concatenation.

Addressing both weaknesses, we propose a novel method for dense long-term pixel-level tracking. It is based on calculating flow not only for consecutive frames, but also for pairs of frames with logarithmically spaced time differences (see Fig. 1). We show that when equipped with suitable estimates of accuracy and of being occluded, a simple strategy for selecting the most reliable concatenation of the set of flows leads to dense and accurate long-term flow trajectories. It is insensitive to medium-length occlusions and, helped by estimating the flow with respect to more distant frames, its drift is reduced.

The idea to obtain long-term correspondences by calculating a set of optical flows, rather than just flow between



Figure 2. **MFT – Multi-Flow Tracker application: video editing.** A WOW! logo, inserted in frame 0 of sequences from selected standard datasets [43,53], propagated by MFT. Frames at 0%, 50%, and 100% of the sequence shown. Full sequences in the supplementary.

consecutive images, appeared for the first time in [12]. This led to a sequence of papers on the topic [9, 10, 13]. The performance of these early, pre-conv-net methods is difficult to assess. They were mainly qualitatively, i.e. visually, tested on a few videos that are not available.

The paper introduces the following contributions: A point-tracking method that is (i) capable of tracking all pixels in a video based on CNN optical flow estimation, (ii) conceptually simple and can be trained and evaluated on a single customer grade GPU. We show (iii) a simple yet effective strategy for selection of long-term optical flow chain candidates, and (iv) how to select the most reliable candidate on the basis of spatial accuracy and occlusion probability obtained by small CNNs trained on synthetic data. We publish the results and the method code ¹.

Experimentally the method outperforms baselines by a large margin and provide a good speed/performance balance, running orders of magnitude faster than the state-of-the-art for video point tracking [16,32] when used for dense point tracking. Fig. 2 shows an application of the proposed method for video editing.

2. Related Work

Object Tracking. Historically, object tracking algorithms [3, 14, 31] estimated the location of an object spec-

ified in the first frame by a bounding box output in every frame of the video sequence. More recently [33, 42], the focus of tracking methods shifted to segmentation of the object or regions specified in the initial frame. Nevertheless, algorithms that are model-free, i.e. are able to track any object specified in the first frame, do not provide point-to-point, dense correspondences.

Structure-from-motion (SfM) and SLAM are two related techniques that can be used for tracking points. Although some methods can estimate the position of points densely [18], they are limited to static scenes. Non-rigid SfM techniques exist but are limited to a closed set of object categories since they require parametric 2D or 3D models [1, 35]. N-NRSfM [46] is template-free, but prior to building a 3D model, it requires accurate 2D long-term tracks of points (chicken-egg problem). Some approaches [55, 56] utilize differentiable rendering or NeRF [40] to create deformable 3D models for tracking points on surfaces. However, their exhaustive computation makes them impractical for real-world usage.

Optical flow estimation is a well-studied problem in computer vision that aims to estimate dense pixel-level displacements between consecutive frames [22]. Modern methods employ deep learning techniques [17, 24, 28, 47, 49] trained on synthetic data. State-of-the-art optical flow methods, such as RAFT [49] and FlowFormer [24], estimate opti-

¹<https://github.com/serycjon/MFT>

cal flow from 4D correlation cost volume of features for all pixel-pairs. While these methods achieve high accuracy for dense estimation of flow between pairs of consecutive frames, estimating accurate flow between distant frames remains a problem, especially for large displacements or large object deformation.

Li *et al.* [34] combines feature matching and optical flow restricted with a deformable mesh. NeuralMarker [23] is trained to find correspondences between the template image and its distorted version inserted into random background image. These approaches allow recovery from occluded regions. However, they are inapplicable for dense tracking in dynamic scenes with non-rigid objects.

To track points over multiple consecutive frames, some methods [4, 37, 48, 50] have proposed to concatenate estimated optical flow. However, they cannot recover from partial occlusions. Standard OF benchmarks [6, 39] do not evaluate occlusion predictions and consequently most OF methods do not detect occlusions at all. Moreover, concatenating optical flow results in error accumulation over time and induce drift in the tracked points. Although some optical flow methods have been proposed to estimate the flow from more than two frames [41, 44], they still operate in a frame-by-frame manner and do not handle partial occlusions well. Therefore, achieving long-term, pixel-wise tracking with optical flow remains a challenging problem in computer vision.

Multi-step-flow (MSF) algorithms [11–13] address the limitations of concatenation-based approaches for long-term dense point tracking. These algorithms construct long-term dense point tracks by merging optical flow estimates computed over varying time steps. This enables handling of temporarily occluded points by skipping them until they reappear. However, they rely on the brightness constancy assumption, which leads to failure over distant frames. The MSF approach has been updated in subsequent works [9, 10] by introducing the multi-step integration and statistical selection (MISS) approach. MISS generates a large number of motion path candidates by randomly selecting reference frames and weighting them based on estimated quality. The optimal candidate path is then determined through global spatial-smoothness optimization. However, these methods are computationally intensive and limited to tracking a small patch of a single object.

In comparison, our proposed MFT picks the best path based on occlusion and uncertainty estimated from correlation cost volume for individual optical flows. Although some optical flow methods estimate occlusions [26, 29, 36, 41, 57, 58] or uncertainty of estimated optical flow [27, 52, 54], state-of-the-art optical flow methods [24, 49] do not provide such estimates. We are the first to employ estimation of occlusion and optical flow uncertainty for the dense and robust long-term tracking of points.

Feature matching identifies corresponding points or regions between images or frames in a sequence. Typically, feature matching is carried out sparsely on estimated key-points. While some dense estimation methods have been developed in the past [4], they have not been able to match the performance of their sparse counterparts until recent advancements, such as the COTR approach [30]. Note that feature matching is still performed only between pairs of frames and estimation of point positions may only be provided independently for target frames.

Point tracking aims to track a set of physical points in a video as introduced in TAP-Vid [15]. A baseline method TAP-Net [15] computes cost volume (similar to RAFT [49]) for a single query point independently for each frame of the sequence. A two-branch network then estimates the position and visibility of the query point in the targeted frame. PIPs [20] focuses on tracking points through occlusions by processing the video in fixed-sized temporal windows. It does not re-detect the target after longer occlusions. PIPs use test-time linking of estimated trajectories since it is limited to tracking in eight consecutive frames only. Particle Video [45] prunes tracked points on occlusion and creates new tracks on disocclusion, however these are not linked together. TAPIR [16] combines the per-frame point localization from TAP-Net [15] with a temporal processing inspired by PIPs [20], but uses a time-wise convolution instead of fixed size frame batches. CoTracker [32] processes query points with a sliding-window transformer that enables multiple tracks to influence each other. This works best when a single query point is tracked at a time, supported by an auxiliary grid of queries. Compared to our proposed approach, these methods do not track densely, but instead focus on tracking individual query points. OmniMotion [51] tracks densely. It pre-processes the video by computing optical flow between all pairs of frames. It represents the whole video with a quasi-3D volume, a NeRF [40]-like network and a set of $2D \leftrightarrow$ quasi-3D bijections. The representation is globally optimized to obtain consistent motion estimates.

3. Method

The proposed method for long-term tracking of every pixel in a template is based on combining optical flow fields computed over different time spans, hence we call it Multi-Flow Tracker, or MFT in short. Given a sequence of $H \times W$ -sized video frames I_0, I_1, \dots, I_N and a list of positions on the reference (template) frame $\mathbf{p}_{i,0} = (x_i, y_i), i \in \{1, \dots, HW\}$ the method predicts the corresponding positions $p_{i,t}$ in all the other frames $t \in \{1, \dots, N\}$, together with an occlusion flag $o_{i,t}$. At time t , the MFT outputs are formed by combining the MFT result from a previous time $t - \Delta$, with the flow from $t - \Delta$ to the current frame t (see Fig. 1). Note that this is not combining only two flows, but




Figure 3. **Optical flow and occlusions.** OF methods are typically [24, 49] trained to ignore occlusions and to predict the ground-truth flow (red) even when occluded in the second frame. Continuing tracking after an occlusion would result in the target drifting to the occluding object. Example from Sintel [6].

appending single flow to a previously computed, arbitrarily long chain of flows. MFT constructs a set of candidate results with varying Δ , then the best candidate is chosen independently for each template position. To rank the candidates, MFT computes and propagates an occlusion map and an uncertainty map in addition to the optical flow fields. Detecting occlusions is necessary to prevent drift to occluding objects as shown in Fig. 3. The position uncertainty serves to pick the most accurate of the candidates. We now describe how the occlusion and uncertainty maps are formed, followed by a detailed description of the proposed MFT.

3.1. Occlusion and Uncertainty

Current optical flow methods typically compute the flow from a cost-volume inner representation and image features [24, 47, 49]. Given a pair of input images, I_a and I_b , the cost-volume encodes similarity between each position in I_a and (possibly a subset of) positions in I_b . We propose to re-use the cost-volume as an input to two small CNNs for occlusion and uncertainty estimation. In both cases we use two convolutional layers with kernel size 3. The first layer has 128 output channels and ReLU activation. Both networks take the same input as the flow estimation head and each outputs a $H \times W$ map.

Occlusion: Similar to [26, 41, 57], we formulate the occlusion prediction as a binary classification. The network should output 1 for any point in I_a that is not visible in I_b and 0 otherwise. We train it on datasets with occlusion ground-truth labels (Sintel [6], FlyingThings [38], and Kubric [19]) using standard cross-entropy loss. The trained CNN achieves 0.96 accuracy on Sintel validation set.

Uncertainty: We train the uncertainty CNN with the uncertainty loss function from [5, 21]

$$\mathcal{L}_u = \frac{1}{2\sigma^2} l_H(\|\vec{x} - \vec{x}^*\|_2) + \frac{1}{2} \log(\sigma^2) \quad (1)$$

where x is the predicted flow, x^* the ground truth flow, σ^2 the predicted uncertainty and l_H is the Huber loss func-

tion [25]. The uncertainty CNN predicts $\alpha = \log(\sigma^2)$ to improve numerical stability during training. We output σ^2 during inference.

We sum the occlusion loss and \mathcal{L}_u weighted by $\frac{1}{5}$. Note that we only train the occlusion and uncertainty networks, keeping the pre-trained optical flow fixed.

3.2. MFT – Multi-Flow Tracker

The MFT tracker is initialized with the first frame of a video. It then outputs a triplet $\bar{\mathbf{FOU}}_{0 \rightarrow t} = (\bar{\mathbf{F}}_{0 \rightarrow t}, \bar{\mathbf{O}}_{0 \rightarrow t}, \bar{\mathbf{U}}_{0 \rightarrow t})$ at each consequent frame I_t . The $\bar{\mathbf{F}}_{0 \rightarrow t}$ is a $H \times W \times 2$ map of position differences between frame number 0 and t , in the classical optical flow format. The $\bar{\mathbf{O}}_{0 \rightarrow t}$ and $\bar{\mathbf{U}}_{0 \rightarrow t}$ are $H \times W$ maps with the current occlusions and uncertainties respectively. On the initialization frame, all three maps contain zeros only (no motion, no occlusion, no uncertainty), on the first frame after initialization, the triplet is directly the output of the optical flow network and the proposed occlusion and uncertainty CNNs. On all the following frames, the results are not the direct outputs of the network, but instead they are formed by chaining two $(\mathbf{F}, \mathbf{O}, \mathbf{U})$ triplets together.

The MFT is parameterized by D , a set of time deltas. We set $D = \{\infty, 1, 2, 4, 8, 16, 32\}$ (logarithmically spaced) by default. For every $\Delta \in D$, we create a result candidate that is formed by chaining two parts – a previously computed result $\bar{\mathbf{FOU}}_{0 \rightarrow (t-\Delta)}$ and a network output $\mathbf{FOU}_{(t-\Delta) \rightarrow t}$ as shown in Fig. 4. To keep the notation simple, we write $(t - \Delta)$, but in fact we compute $\max(0, t - \Delta)$ to avoid invalid negative frame numbers.

To do the chaining, we first define a new map $\bar{\mathbf{P}}_{(t-\Delta)}$ storing the point positions in time $(t - \Delta)$. For each position $\mathbf{p} = (x, y)$ in the initial frame, the position in time $(t - \Delta)$ is calculated as

$$\bar{\mathbf{P}}_{(t-\Delta)}[\mathbf{p}] = \mathbf{p} + \bar{\mathbf{F}}_{0 \rightarrow (t-\Delta)}[\mathbf{p}], \quad (2)$$

where $\mathbf{A}[\mathbf{b}]$ means the value in a map \mathbf{A} at integer spatial coordinates \mathbf{b} . To form the candidate $\mathbf{F}_{0 \rightarrow t}^\Delta$, we add the optical flow $\mathbf{F}_{(t-\Delta) \rightarrow t}$, sampled at the appropriate position to the motion between frames 0 and $(t - \Delta)$.

$$\mathbf{F}_{0 \rightarrow t}^\Delta[\mathbf{p}] = \bar{\mathbf{F}}_{0 \rightarrow (t-\Delta)}[\mathbf{p}] + \mathbf{F}_{(t-\Delta) \rightarrow t}[\bar{\mathbf{P}}_{(t-\Delta)}[\mathbf{p}]]_s \quad (3)$$

where $\mathbf{A}[\mathbf{b}]_s$ means the value in a map \mathbf{A} sampled at possibly non-integer spatial coordinates \mathbf{b} with bilinear interpolation. When chaining two occlusion scores, we take their maximum.

$$\mathbf{O}_{0 \rightarrow t}^\Delta[\mathbf{p}] = \max(\bar{\mathbf{O}}_{0 \rightarrow (t-\Delta)}[\mathbf{p}]; \mathbf{O}_{(t-\Delta) \rightarrow t}[\bar{\mathbf{P}}_{(t-\Delta)}[\mathbf{p}]]_s) \quad (4)$$

Since we threshold the occlusion scores in the end to get a binary decision, this corresponds to an “or” operation – the chain is declared occluded whenever at least one of its parts is occluded.




Figure 4. **Schematic explanation of the MFT tracking procedure.** At the current frame, time t (right), the tracker creates a set of result candidates, each formed by a different chain of optical flows. In this example, the first candidate ③ is formed by chaining the result ① previously computed in time $(t - \Delta_1)$ with flow ② estimated between frames $(t - \Delta_1)$ and t . We use bilinear interpolation (red) to sample the flow field, since the positions in $(t - \Delta_1)$ usually do not align with the pixel grid. The flow ③ into the current frame t is constructed by summing the two flow vectors. We repeat this procedure for Δ_2 , again summing the result ④ for frame $(t - \Delta_2)$ with ⑤ – the bilinearly sampled flow field from $(t - \Delta_2)$ to t . When chaining the flows, we also chain their occlusion and uncertainty maps. Finally, we select the candidate (③), or (⑥) with the lowest uncertainty score among the ones not occluded, or mark the result occluded when all candidates predict occlusion. Current point position shown in blue, grid-aligned flow vectors in black, interpolated flow vectors in red.

The uncertainties are chained by addition, as they represent the variance of the sum of flows, assuming independence of individual uncertainties.

$$\mathbf{U}_{0 \rightarrow t}^\Delta[\mathbf{p}] = \bar{\mathbf{U}}_{0 \rightarrow (t-\Delta)}[\mathbf{p}] + \mathbf{U}_{(t-\Delta) \rightarrow t}[\bar{\mathbf{P}}_{(t-\Delta)}[\mathbf{p}]]_s \quad (5)$$

We repeat the chaining procedure for each $\Delta \in D$ to obtain up to $|D|$ different result candidates. Finally, we select the best Δ , Δ^* according to candidate uncertainty and occlusion maps. In particular, we pick the Δ that has the lowest uncertainty score among the unoccluded candidates. When all the candidates are occluded (occlusion score larger than a threshold θ_o), all candidates are equally good and the first one is selected.

$$\Delta^*[\mathbf{p}] = \arg \min_{\Delta \in D} \mathbf{U}_{0 \rightarrow t}^\Delta[\mathbf{p}] + \infty \cdot \llbracket \mathbf{O}_{0 \rightarrow t}^\Delta[\mathbf{p}] > \theta_o \rrbracket, \quad (6)$$

where $\llbracket x \rrbracket$ is the Iverson bracket (equal to 1 when condition x holds, 0 otherwise). Notice that we select the Δ^* independently for each position. For example with $D = \{\infty, 1\}$, the flows are computed either directly between the template and the current frame ($\Delta = \infty$), or from the previous to the current frame ($\Delta = 1$) as in the traditional OF setup. For some parts of the image, it is better to use $\Delta = \infty$, because having a direct link to the template does not introduce drift. On the other hand, on some parts of the image the appearance might have significantly changed over the longer time span, making the direct flow not reliable at the current frame. In such case a long chain of $\Delta = 1$ flows might be preferred. Note that MFT usually switches back and forth between the used Δ s during the tracking. A single template query point might be tracked using a chain of $\Delta = 1$ flows for some time, then it might switch to the direct $\Delta = \infty$ flow for some frames (possibly undoing any accumulated drift), then back to $\Delta = 1$ and so on.

The final result at frame t is formed by selecting the result from the candidate corresponding to Δ^* in each pixel, e.g., for the flow output $\bar{\mathbf{F}}_{0 \rightarrow t}$ we have

$$\bar{\mathbf{F}}_{0 \rightarrow t}[\mathbf{p}] = \mathbf{F}_{0 \rightarrow t}^{\Delta^*}[\mathbf{p}] \quad (7)$$

Finally, MFT memorizes and outputs the resulting triplet $\bar{\mathbf{FOU}}_{0 \rightarrow t}$ and discard memorized results that will no longer be needed (more than $\max(D \setminus \{\infty\})$ frames old). Given query positions $\mathbf{p}_{i,0}$ on the template frame 0, we compute their current positions and occlusion flags by bilinear interpolation of the $\bar{\mathbf{FOU}}$ result.

$$\mathbf{p}_{i,t} = \mathbf{p}_{i,0} + \bar{\mathbf{F}}_{0 \rightarrow t}[\mathbf{p}_{i,0}]_s \quad (8)$$

$$o_{i,t} = \bar{\mathbf{O}}_{0 \rightarrow t}[\mathbf{p}_{i,0}]_s \quad (9)$$

3.3. Implementation Details

For the optical flow, we use the official RAFT [49] implementation with author-provided weights. Both the occlusion and the uncertainty CNNs operate on the same inputs as the RAFT flow regression CNN, i.e. samples from the RAFT cost-volume, context features, and Conv-GRU outputs. We train on Sintel [6], FlyingThings [38], and Kubric [19]. We sample training images with equal probability from each dataset. Because the Kubric images are smaller than the RAFT training pipeline expects, we randomly upscale them with scale ranging between 3.2 \times and 4.6 \times . We train the occlusion and the uncertainty network for 50k iterations with the original RAFT training hyperparameters, which takes around 10 hours on a single GPU.

The MFT tracker is implemented in PyTorch and all the operations are performed on GPU. Note that the optical flows and the occlusion and uncertainty maps can be pre-computed offline. When the $\Delta = \infty$ is not included in D , the number of pre-computed flow fields needed to be stored in order to be able track forward or backward from any frame in a video is less than $N^2|D|$. Pre-computing flows for $\Delta = \infty$ (direct from template) and all possible template frames is not practical, as the number of stored flow fields grows quadratically with the number of frames N . With the flows for other Δ s pre-computed, MFT needs to compute just one OF per frame during inference, so the tracking speed stays reasonably fast.

On a GeForce RTX 2080 Ti GPU (i7-8700K CPU @ 3.70GHz), the chaining of the flow, occlusion and uncertainty maps takes approximately 1.3ms for each Δ candidate with videos of 512×512 resolution. On average, the preparation of all the result candidates takes 8ms. The per-pixel selection of the best one adds additional 0.6ms. Computing a single RAFT flow, including the extra occlusion and uncertainty outputs, takes 60ms. Altogether, the full MFT runs at 2.3FPS. With pre-computed flows MFT runs at over 100FPS, making it suitable for interactive applications in, e.g., film post-production. We set $\theta_o = 0.02$ empirically.

4. Experiments

Since there is no benchmark for dense long-term point tracking, we evaluate the MFT on the recently introduced TAP-Vid DAVIS and TAP-Vid Kinetics datasets [15] for sparse point tracking. The datasets consists of 30 videos from DAVIS 2017 [43] and 1189 videos from Kinetics-700 [7, 8] respectively, rescaled to 256×256 resolution, semi-automatically annotated with positions and occlusion flags of ≈ 20 selected points.

Evaluation protocol: The TAP-Vid benchmark uses two evaluation modes: “first” and “strided”. In the “first” mode, the tracker is initialized on the first frame where the currently evaluated ground-truth tracked point becomes visible, and is only evaluated on the following frames. In the “strided” mode, the tracker is initialized on frames 0, 5, 10, ... if the currently evaluated tracked point is visible in the given frame. The tracker is then evaluated on both the following and the preceding frames, we thus run our MFT method two times, forward and backward in time, starting on the initialization frame.

Evaluation metrics: The TAP-Vid benchmark uses three metrics. The occlusion prediction quality is measured by occlusion classification accuracy (OA). The accuracy of the predicted positions, $\langle \delta_{avg}^x \rangle$, is measured by fraction of visible points with position error under a threshold, averaged over thresholds 1, 2, 4, 8, 16. Both occlusion and position accuracy are captured by Average Jaccard (AJ), see [15] for more details.

4.1. Flow Delta Ablation

In Table 1, we show the impact of using different sets D of Δ s. We evaluate two baselines – (1) basic chaining of consecutive optical flows ($\Delta = 1$), and (2), computing the optical flow directly between the template and the current frame ($\Delta = \infty$). The first one performs better in all metrics, as the OF is computed on pairs of consecutive images, which it was trained to do, and the test sequences are not long enough to induce significant drift by error accumulation. Note that the performance in the strided evaluation mode is better, because the sequences are on average two times shorter and contain less occlusions.

	flow delta set D	DAVIS - first			DAVIS - strided		
		AJ	$\langle \delta_{avg}^x \rangle$	OA	AJ	$\langle \delta_{avg}^x \rangle$	OA
(1)	{1}	38.3	54.5	69.3	48.9	61.8	80.8
(2)	{ ∞ }	38.3	50.8	65.5	47.9	58.0	76.3
(3)	{ $\infty, 1$ }	46.4	63.7	76.7	55.0	68.1	85.8
(4)	{ $\infty, 1, 2, 4, 8, 16, 32$ }	47.3	66.8	77.8	56.1	70.8	86.9
(5)	{1, 2, 4, 8, 16, 32}	47.4	66.2	77.3	55.7	70.2	86.5

Table 1. **TAP-Vid Davis benchmark – evaluation of MFT on variants based on different sets D** of time differences Δ used in optical flow; ∞ indicates OF between the template and the current frame. Performance measured by occlusion accuracy (OA), position accuracy ($\langle \delta_{avg}^x \rangle$), and combined measure AJ. For definition of $\langle \delta_{avg}^x \rangle$ and AJ, see text. Bold best, underline second.

	Resolution H×W	DAVIS - first			DAVIS - strided		
		AJ	$\langle \delta_{avg}^x \rangle$	OA	AJ	$\langle \delta_{avg}^x \rangle$	OA
(1)	256×256	33.0	47.7	70.2	41.4	54.6	83.6
(2)	256×256→512×512	47.3	66.8	77.8	56.1	70.8	86.9
(3)	256×256→256×ratio	40.5	58.5	76.9	49.2	63.8	86.4
(4)	256×256→480×ratio	49.2	69.2	77.9	58.8	73.9	87.7
(5)	orig res.→480×ratio	52.3	71.9	79.5	61.9	76.1	88.8
(6)	orig res.→720×ratio	54.0	74.0	79.1	64.3	78.7	88.1

Table 2. **TAP-Vid Davis benchmark – evaluation of MFT for different image resolutions.** Performance measured as in Table 1.

Combining the basic chaining with the direct OF, line (3) in Table 1, the performance increases in all metrics, showing the effectivity of the proposed candidate selection mechanism. Row (4) is the full MFT method which achieves the overall best results. The final experiment (5) works without the direct flow. This means that we can pre-compute all the optical flows needed to track from any frame in any time direction, and store them in storage space proportional to the number of frames $2N|D|$. Note that attempting to do that with $\infty \in D$ would result in storage requirements proportional to N^2 . The last version achieves second best overall performance. Visual performance of the baselines and full MFT is shown in Fig. 5. All results in Table 1 were obtained on $2\times$ upscaled images as discussed in the next section which is equivalent to adding one upsampling layer to the RAFT feature pyramid.

4.2. Input Resolution Ablation

The official TAP-Vid benchmark is evaluated on videos rescaled to 256×256 resolution, which is small compared with the RAFT training set. Because of this, we upscale the 256×256 videos to 512×512 resolution. In all the experiments, the output positions are scaled back to the 256×256 resolution for evaluation. Rows (1) and (2) in Table 2 show that this upscaling improves the performance by a large margin on all three metrics. This shows that RAFT is sensitive to input sizes, note that no information was added to the images when upscaling.

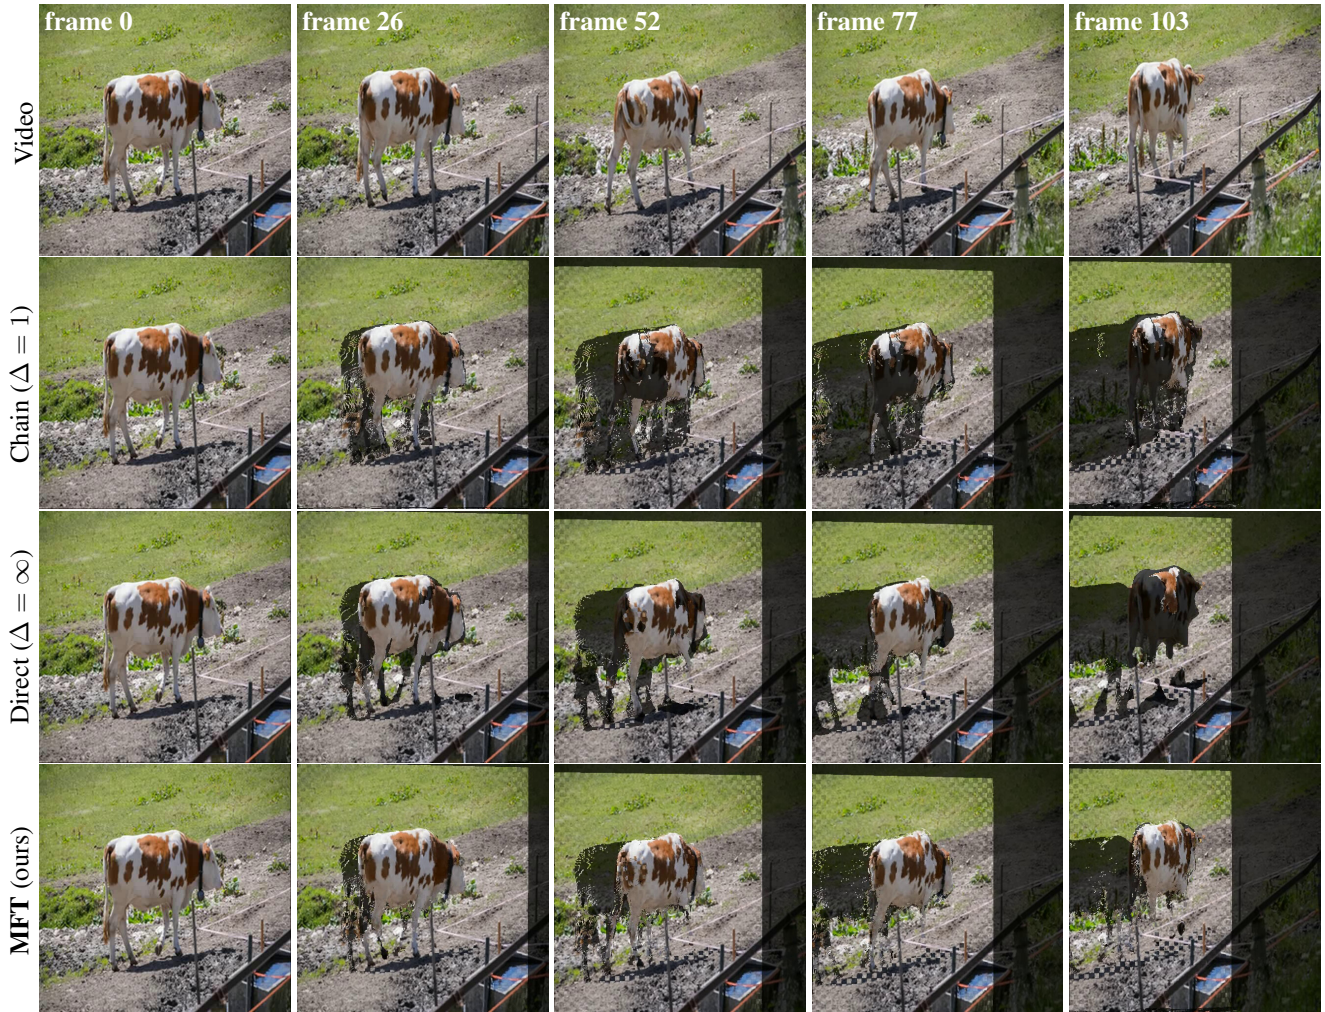


Figure 5. Result visualizations sampled at 25%, 50%, 75% and 100% of the input video (*top*) length. We take the first frame of a video and set its transparency with a checkerboard pattern. We then warp the resulting image using the outputs of each method and overlay the result on the current frame. The checkerboard pattern is visible when the tracking results are incorrect, or when the illumination changed between the template and the current frame. Pixels without a correspondence on the template frame are darkened. *Row 2*: simple flow chaining $\Delta = 1$. A short occlusion by the tail makes the tracker lose track in the back half of the cow. *Row 3*: direct flow $\Delta = \infty$. The tracker survives the occlusion but loses track when the cow rotates away from the camera. *Bottom*: the proposed MFT handles both the short occlusion and the appearance change, tracking well on background and most of the cow’s body. All trackers fail on the legs which are too thin for the RAFT optical flow. Best viewed zoomed-in on a screen.

The aspect ratio of the original videos is changed during the scaling from full DAVIS resolution to the 256×256 . This makes the video contents appear distorted and changes the motion statistics. Consequently we perform several experiments with varying video resolutions but keeping the original aspect ratio. In the first two, (rows (3), (4) in Table 2), we upsample the 256×256 videos. This way we stick as close to the TAP-Vid protocol as possible, only requiring the original video aspect ratio as an extra input. In (3), we keep the image height unchanged and only upscale the width such that the aspect ratio is not changed wrt the full resolution videos. All the metrics improve compared

to the no scaling variant (1). Also, when we upscale the images to larger size (4), the performance increases.

In the last two rows (5), (6), we skip the TAP-Vid downscaling to 256×256 and instead downscale to the target resolution directly from the full-resolution DAVIS videos. This preserves high-frequency details more than doing the downscale-upscale cycle. Thanks to this, row (5) is better than (4), although the input resolution is the same in both. Even larger resolution (6) again improves the $\langle \delta_{avg}^x \rangle$ and the AJ metric for the cost of small (below one percent point) decrease in occlusion accuracy.

Because we downscale directly from the full resolution,

without the 256×256 intermediate step, the results of (5) and (6) are not directly comparable with the original TAP-Vid benchmark table, but are closer to a real-world scenario.

4.3. Comparison With the State-of-the-Art

On the TAP-Vid benchmark, the proposed MFT tracker performs third best, after the state-of-the-art sparse point-tracking methods [16, 32], out-performing the other dense point-tracker OmniMotion [51]. MFT runs at over 2FPS, which is orders of magnitude faster than the alternative methods evaluated densely, tracking every pixel and not just selected few. The speed/performance balance makes MFT favorable for dense point-tracking. Additionally, the optical flows can be pre-computed (only $2N \log N$ flows needed for a video of length N with logarithmically spaced flow delta set D) resulting in tracking at over 100FPS from any frame in the video, both forward and backward. This makes MFT a good candidate for interactive applications such as video editing. The complete results, including the inference speeds, are shown in Table 3. Both MFT and OmniMotion [51] can be seen as post-processing of a set of RAFT optical flows. The MFT strategy performs better than the complex model and global optimization in OmniMotion.

One MFT weakness we have observed are spurious re-detections. MFT sometimes matches out-of-view parts of the template to visually similar parts of the current frame. Single such incorrect re-detection can “restart” a flow chain, affecting the performance for the rest of the video. A typical example is tracking of a point on a road surface. When the camera moves such that the original point moves far out of view, the tracklet sometimes suddenly jumps to a newly uncovered patch of the road. Both the appearance of the incorrectly matched point and its image context is often very similar to the template frame, *e.g.*, a relatively texture-less black road some distance below a car wheel.

BADJA evaluation. In addition to TAP-Vid DAVIS, we evaluate the MFT on BADJA [2] benchmark with videos of animals annotated with 2D positions of selected joints. The benchmark measures the percentage of points with position error under a permissive threshold $0.2\sqrt{A}$, where A is the area of the animal segmentation mask. Thanks to this, the MFT performs well even though the ground truth points (joints) are located under the surface, and thus, MFT cannot track them directly. In Table 4, we evaluate against the BADJA results of PIPs [20] and their RAFT baseline. In terms of median of the per-sequence results, MFT performs the best. The mean score is affected by a single failure sequence, *dog-a*, on which the dog turns shortly after the first frame, making most of the tracklets occluded. The assumption that a joint can be approximately tracked by tracking a nearby point on the surface becomes invalid in such case.

Method	FPS	DAVIS - first			DAVIS - strided			Kinetics - first		
		AJ	$<\delta_{avg}^x$	OA	AJ	$<\delta_{avg}^x$	OA	AJ	$<\delta_{avg}^x$	OA
TAP-Net [15]	0.11†	33.0	48.6	78.8	38.4	53.1	82.3	38.4	54.4	80.6
PIPs [20]	2e-4†	-	-	-	42.0	59.4	82.1	31.7	53.7	72.9
OmniMotion [51]	2e-3†	-	-	-	51.7	67.5	85.3	-	-	-
MFT (ours)	2.32	47.3	66.8	77.8	56.1	70.8	86.9	39.6	60.4	72.7
TAPIR [16]	0.04†	56.2	70.0	86.5	61.3	72.3	87.6	49.6	64.2	85.0
CoTracker [32]	0.04	60.6	75.4	89.3	64.8	79.1	88.7	48.7	64.3	86.5

Table 3. **Evaluation on TAP-Vid benchmark. MFT performs well while being orders of magnitude faster than other methods when evaluated densely.** Performance measured as in Table 1. Results for other methods are from [15, 16, 32, 51]. FPS: speed of dense (every pixel) tracking on 512×512 video in Frames Per Second. Speeds marked with † were extrapolated from timing info in [16, 51], details in supplementary.

	a	b	c	d	e	f	g	Avg.	Med.
RAFT	64.6	65.6	69.5	13.8	39.1	37.1	29.3	45.6	39.1
PIPs	76.3	81.6	83.2	34.2	44.0	57.4	59.5	62.3	59.5
MFT	81.8	82.0	75.7	6.9	47.9	55.8	62.7	59.0	62.7

Table 4. **BADJA [2] benchmark – evaluation of MFT against PIPs [20].** Performance measured by the PCK-T measure, *i.e.*, the percentage of points with error under a threshold. Bold best. Results for PIPs and RAFT from [20]. The labeled individual sequences include (a) bear, (b) camel, (c) cows, (d) dogs-a, (e) dog, (f) horse-h, and (g) horse-l.

5. Conclusions

We proposed MFT – a novel method for long-term tracking of every pixel on the template frame. Its novelties include an introduction of two small CNNs estimating occlusion and flow uncertainty maps that are highly effective in selecting accurate flow chains that exploit flow computed both between consecutive and non-consecutive frames. MFT performs well on the point-tracking benchmark TAP-Vid [15], and enables tracking all template pixels densely much faster (2.4 FPS vs 0.04 FPS) than the state-of-the-art point-trackers. With pre-computed flows, MFT tracks densely at over 100FPS, enabling real-time interactivity for applications such as video editing. Flow fields needed for MFT can be pre-computed offline, with storage requirements growing log-linearly with the video sequence length. We also evaluated MFT on BADJA dataset, showing competitive performance on animal joint tracking. We also show that accuracy of the popular RAFT optical flow increases significantly with input image resolution, even when upscaling low-resolution images which does not provide additional details. We will publish the MFT code and weights. **Acknowledgements.** This work was supported by Toyota Motor Europe, by the Grant Agency of the Czech Technical University in Prague, grant No.SGS23/173/OHK3/3T/13, and by the Research Center for Informatics project CZ.02.1.01/0.0/0.0/16_019/0000765 funded by OP VVV.

References

- [1] Benjamin Biggs, Oliver Boyne, James Charles, Andrew Fitzgibbon, and Roberto Cipolla. Who left the dogs out? 3d animal reconstruction with expectation maximization in the loop. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 195–211. Springer, 2020. [2](#)
- [2] Benjamin Biggs, Thomas Roddick, Andrew Fitzgibbon, and Roberto Cipolla. Creatures great and SMAL: Recovering the shape and motion of animals from video. In *ACCV*, 2018. [8](#)
- [3] David S Bolme, J Ross Beveridge, Bruce A Draper, and Yui Man Lui. Visual object tracking using adaptive correlation filters. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 2544–2550. IEEE, 2010. [2](#)
- [4] Thomas Brox and Jitendra Malik. Object segmentation by long term analysis of point trajectories. In *European conference on computer vision*, pages 282–295. Springer, 2010. [3](#)
- [5] David Brüggemann, Christos Sakaridis, Prune Truong, and Luc Van Gool. Refign: Align and refine for adaptation of semantic segmentation to adverse conditions. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3174–3184, 2023. [4](#)
- [6] Daniel J Butler, Jonas Wulff, Garrett B Stanley, and Michael J Black. A naturalistic open source movie for optical flow evaluation. In *Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part VI 12*, pages 611–625. Springer, 2012. [3](#), [4](#), [5](#)
- [7] Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019. [6](#)
- [8] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. [6](#)
- [9] Pierre-Henri Conze, Philippe Robert, Tomas Crivelli, and Luce Morin. Dense long-term motion estimation via statistical multi-step flow. In *2014 International Conference on Computer Vision Theory and Applications (VISAPP)*, volume 3, pages 545–554. IEEE, 2014. [2](#), [3](#)
- [10] Pierre-Henri Conze, Philippe Robert, Tomas Crivelli, and Luce Morin. Multi-reference combinatorial strategy towards longer long-term dense motion estimation. *Computer Vision and Image Understanding*, 150:66–80, 2016. [2](#), [3](#)
- [11] Tomás Crivelli, Pierre-Henri Conze, Philippe Robert, Matthieu Fradet, and Patrick Pérez. Multi-step flow fusion: towards accurate and dense correspondences in long video shots. In *British Machine Vision Conference*, 2012. [3](#)
- [12] Tomas Crivelli, Pierre-Henri Conze, Philippe Robert, and Patrick Pérez. From optical flow to dense long term correspondences. In *2012 19th IEEE International Conference on Image Processing*, pages 61–64. IEEE, 2012. [2](#), [3](#)
- [13] Tomas Crivelli, Matthieu Fradet, Pierre-Henri Conze, Philippe Robert, and Patrick Pérez. Robust optical flow integration. *IEEE Transactions on Image Processing*, 24(1):484–498, 2014. [2](#), [3](#)
- [14] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Atom: Accurate tracking by overlap maximization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4660–4669, 2019. [2](#)
- [15] Carl Doersch, Ankush Gupta, Larisa Markeeva, Adria Recasens Contente, Lucas Smaira, Yusuf Aytar, Joao Carreira, Andrew Zisserman, and Yi Yang. TAP-Vid: A benchmark for tracking any point in a video. *Advances in Neural Information Processing Systems*, 2022. [3](#), [6](#), [8](#)
- [16] Carl Doersch, Yi Yang, Mel Vecerik, Dilara Gokay, Ankush Gupta, Yusuf Aytar, Joao Carreira, and Andrew Zisserman. TAPIR: Tracking any point with per-frame initialization and temporal refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10061–10072, October 2023. [2](#), [3](#), [8](#)
- [17] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Häusser, Caner Hazırbaş, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. Flownet: Learning optical flow with convolutional networks. In *ICCV*, pages 2758–2766, Dec. 2015. [2](#)
- [18] Mariia Gladkova, Nikita Korobov, Nikolaus Demmel, Aljoša Ošep, Laura Leal-Taixé, and Daniel Cremers. Directtracker: 3d multi-object tracking using direct image alignment and photometric bundle adjustment. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3777–3784. IEEE, 2022. [2](#)
- [19] Klaus Greff, Francois Belletti, Lucas Beyer, Carl Doersch, Yilun Du, Daniel Duckworth, David J Fleet, Dan Gnanaprasam, Florian Golemo, Charles Herrmann, et al. Kubric: A scalable dataset generator. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3749–3761, 2022. [4](#), [5](#)
- [20] Adam W Harley, Zhaoyuan Fang, and Katerina Fragkiadaki. Particle video revisited: Tracking through occlusions using point trajectories. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*, pages 59–75. Springer, 2022. [3](#), [8](#)
- [21] Yihui He, Chenchen Zhu, Jianren Wang, Marios Savvides, and Xiangyu Zhang. Bounding box regression with uncertainty for accurate object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2888–2897, 2019. [4](#)
- [22] Berthold KP Horn and Brian G Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. [2](#)
- [23] Zhaoyang Huang, Xiaokun Pan, Weihong Pan, Weikang Bian, Yan Xu, Ka Chun Cheung, Guofeng Zhang, and Hongsheng Li. Neuralmarker: A framework for learning general marker correspondence. *ACM Transactions on Graphics (TOG)*, 41(6):1–10, 2022. [3](#)
- [24] Zhaoyang Huang, Xiaoyu Shi, Chao Zhang, Qiang Wang, Ka Chun Cheung, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Flowformer: A transformer architecture for optical flow. *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVII*, pages 668–685, 2022. [2](#), [3](#), [4](#)

- [25] Peter J Huber. Robust estimation of a location parameter. *Breakthroughs in statistics: Methodology and distribution*, pages 492–518, 1992. 4
- [26] Junhwa Hur and Stefan Roth. Iterative residual refinement for joint optical flow and occlusion estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5754–5763, 2019. 3, 4
- [27] Eddy Ilg, Ozgun Cicek, Silvio Galesso, Aaron Klein, Osama Makansi, Frank Hutter, and Thomas Brox. Uncertainty estimates and multi-hypotheses networks for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 652–667, 2018. 3
- [28] Eddy Ilg, Nikolaus Mayer, Tomoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [29] Eddy Ilg, Tomoy Saikia, Margret Keuper, and Thomas Brox. Occlusions, motion and depth boundaries with a generic network for disparity, optical flow or scene flow estimation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 614–630, 2018. 3
- [30] Wei Jiang, Eduard Trulls, Jan Hosang, Andrea Tagliasacchi, and Kwang Moo Yi. Cotr: Correspondence transformer for matching across images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6207–6217, 2021. 3
- [31] Zdenek Kalal, Krystian Mikolajczyk, and Jiri Matas. Tracking-learning-detection. *IEEE transactions on pattern analysis and machine intelligence*, 34(7):1409–1422, 2011. 2
- [32] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. CoTracker: It is better to track together. *arXiv preprint arXiv:2307.07635*, 2023. 2, 3, 8
- [33] Matej Kristan, Aleš Leonardis, Jiří Matas, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kämäräinen, Martin Danelljan, Luka Čehovin Zajc, Alan Lukežič, Ondrej Drbohlav, et al. The eighth visual object tracking vot2020 challenge results. In *European Conference on Computer Vision*, pages 547–601. Springer, 2020. 2
- [34] Wenbin Li, Darren Cosker, and Matthew Brown. Drift robust non-rigid optical flow enhancement for long sequences. *Journal of Intelligent & Fuzzy Systems*, 31(5):2583–2595, 2016. 3
- [35] Xueting Li, Sifei Liu, Shalini De Mello, Kihwan Kim, Xiaolong Wang, Ming-Hsuan Yang, and Jan Kautz. Online adaptation for consistent mesh reconstruction in the wild. *Advances in Neural Information Processing Systems*, 33:15009–15019, 2020. 2
- [36] Shuaicheng Liu, Kunming Luo, Nianjin Ye, Chuan Wang, Jue Wang, and Bing Zeng. Oiflow: Occlusion-inpainting optical flow estimation by unsupervised learning. *IEEE Transactions on Image Processing*, 30:6420–6433, 2021. 3
- [37] Yu Liu, Jianbing Shen, Wenguan Wang, Hanqiu Sun, and Ling Shao. Better dense trajectories by motion in videos. *IEEE transactions on cybernetics*, 49(1):159–170, 2017. 3
- [38] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4040–4048, 2016. 4, 5
- [39] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3061–3070, 2015. 3
- [40] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2, 3
- [41] Michal Neoral, Jan Šochman, and Jiří Matas. Continual occlusion and optical flow estimation. In C.V. Jawahar, Hongdong Li, Greg Mori, and Konrad Schindler, editors, *Computer Vision – ACCV 2018*, pages 159–174, Cham, 2019. Springer International Publishing. 3, 4
- [42] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Computer Vision and Pattern Recognition*, 2016. 2
- [43] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675v2*, 2017. 2, 6
- [44] Zhile Ren, Orazio Gallo, Deqing Sun, Ming-Hsuan Yang, Erik B Sudderth, and Jan Kautz. A fusion approach for multi-frame optical flow estimation. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2077–2086. IEEE, 2019. 3
- [45] Peter Sand and Seth Teller. Particle video: Long-range motion estimation using point trajectories. *International Journal of Computer Vision*, 80:72–91, 2008. 3
- [46] Vikramjit Sidhu, Edgar Tretschk, Vladislav Golyanik, Antonio Agudo, and Christian Theobalt. Neural dense non-rigid structure from motion with latent space constraints. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*, pages 204–222. Springer, 2020. 2
- [47] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Janf Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. *CVPR*, 2018. 2, 4
- [48] Narayanan Sundaram, Thomas Brox, and Kurt Keutzer. Dense point trajectories by gpu-accelerated large displacement optical flow. In *ECCV*, 2010. 3
- [49] Zachary Teed and Jia Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *European Conference on Computer Vision*, pages 402–419. Springer, 2020. 2, 3, 4, 5
- [50] Heng Wang, Alexander Kläser, Cordelia Schmid, and Cheng-Lin Liu. Dense trajectories and motion boundary descriptors for action recognition. *International journal of computer vision*, 103:60–79, 2013. 3
- [51] Qianqian Wang, Yen-Yu Chang, Ruojin Cai, Zhengqi Li, Bharath Hariharan, Aleksander Holynski, and Noah

- Snavely. Tracking everything everywhere all at once. *arXiv:2306.05422*, 2023. 3, 8
- [52] Anne S Wannenwetsch, Margret Keuper, and Stefan Roth. Probflow: Joint optical flow and uncertainty estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 1173–1182, 2017. 3
- [53] Ning Xu, Linjie Yang, Yuchen Fan, Dingcheng Yue, Yuchen Liang, Jianchao Yang, and Thomas Huang. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327*, 2018. 2
- [54] Gengshan Yang and Deva Ramanan. Volumetric correspondence networks for optical flow. In *Advances in Neural Information Processing Systems*, pages 793–803, 2019. 3
- [55] Gengshan Yang, Deqing Sun, Varun Jampani, Daniel Vlasic, Forrester Cole, Huiwen Chang, Deva Ramanan, William T Freeman, and Ce Liu. Lasr: Learning articulated shape reconstruction from a monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15980–15989, 2021. 2
- [56] Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. Banmo: Building animatable 3d neural models from many casual videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2863–2873, 2022. 2
- [57] Congxuan Zhang, Cheng Feng, Zhen Chen, Weiming Hu, and Ming Li. Parallel multiscale context-based edge-preserving optical flow estimation with occlusion detection. *Signal Processing: Image Communication*, 101:116560, 2022. 3, 4
- [58] Shengyu Zhao, Yilun Sheng, Yue Dong, Eric I Chang, Yan Xu, et al. Maskflownet: Asymmetric feature matching with learnable occlusion mask. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6278–6287, 2020. 3