

# Tracking Everything Everywhere All at Once

Qianqian Wang<sup>1,2</sup> Yen-Yu Chang<sup>1</sup> Ruojin Cai<sup>1</sup> Zhengqi Li<sup>2</sup>  
Bharath Hariharan<sup>1</sup> Aleksander Holynski<sup>2,3</sup> Noah Snavely<sup>1,2</sup>

<sup>1</sup>Cornell University <sup>2</sup>Google Research <sup>3</sup>UC Berkeley



Figure 1: We present a new method for estimating full-length motion trajectories for every pixel in every frame of a video, as illustrated by the motion paths shown above. For clarity, we only show sparse trajectories for foreground objects, though our method computes motion for *all* pixels. Our method yields accurate, coherent long-range motion even for fast-moving objects, and robustly tracks through occlusions as shown in the *dog* and *swing* examples. For context, in the second row we depict the moving object at different moments in time.

## Abstract

We present a new test-time optimization method for estimating dense and long-range motion from a video sequence. Prior optical flow or particle video tracking algorithms typically operate within limited temporal windows, struggling to track through occlusions and maintain global consistency of estimated motion trajectories. We propose a complete and globally consistent motion representation, dubbed OmniMotion, that allows for accurate, full-length motion estimation of every pixel in a video. OmniMotion represents a video using a quasi-3D canonical volume and performs pixel-wise tracking via bijections between local and canonical space. This representation allows us to ensure global consistency, track through occlusions, and model any combination of camera and object motion. Extensive evaluations on the TAP-Vid benchmark and real-world footage show that our approach outperforms prior state-of-the-art methods by a large margin both quantitatively and qualitatively. See our project page for more results: [omnimotion.github.io](https://omnimotion.github.io).

## 1. Introduction

Motion estimation methods have traditionally followed one of two dominant approaches: sparse feature tracking and dense optical flow [53]. While each type of method has proven effective for their respective applications, neither representation fully models the motion of a video: pairwise optical flow fails to capture motion trajectories over long temporal windows, and sparse tracking does not model the motion of all pixels.

A number of approaches have sought to close this gap, i.e., to estimate both *dense* and *long-range* pixel trajectories in a video. These range from methods that simply chain together two-frame optical flow fields, to more recent approaches that directly predict per-pixel trajectories across multiple frames [22]. Still, these methods all use limited context when estimating motion, disregarding information that is either temporally or spatially distant. This locality can result in accumulated errors over long trajectories and spatio-temporal inconsistencies in the motion estimates. Even when prior methods do consider long-range context [53], they op-

erate in the 2D domain, resulting in a loss of tracking during occlusion events. All in all, producing both *dense* and *long-range* trajectories remains an open problem in the field, with three key challenges: (1) maintaining accurate tracks across long sequences, (2) tracking points through occlusions, and (3) maintaining coherence in space and time.

In this work, we propose a holistic approach to video motion estimation that uses *all* the information in a video to jointly estimate full-length motion trajectories for every pixel. Our method, which we dub *OmniMotion*, uses a quasi-3D representation in which a canonical 3D volume is mapped to per-frame local volumes through a set of local-canonical bijections. These bijections serve as a flexible relaxation of dynamic multi-view geometry, modeling a combination of camera and scene motion. Our representation guarantees cycle consistency, and can track all pixels, even while occluded (“*Everything, Everywhere*”). We optimize our representation per video to jointly solve for the motion of the entire video “*All at Once*”. Once optimized, our representation can be queried at any continuous coordinate in the video to receive a motion trajectory spanning the entire video.

In summary, we propose an approach that: 1) produces globally consistent full-length motion trajectories for all points in an entire video, 2) can track points through occlusions, and 3) can tackle in-the-wild videos with any combination of camera and scene motion. We demonstrate these strengths quantitatively on the TAP video tracking benchmark [14], where we achieve state-of-the-art performance, outperforming all prior methods by a large margin.

## 2. Related Work

**Sparse feature tracking.** Tracking features [4, 41] across images is essential for a wide range of applications such as Structure from Motion (SfM) [1, 54, 57] and SLAM [16]. While sparse feature tracking [12, 42, 55, 65] can establish long-range correspondence, this correspondence is limited to a set of distinctive interest points, and often restricted to rigid scenes. Hence, below we focus on work that can produce dense pixel motion for general videos.

**Optical flow.** Optical flow has traditionally been formulated as an optimization problem [6, 7, 23, 73]. However, recent advances have enabled direct prediction of optical flow using neural networks with improved quality and efficiency [19, 24, 25, 59]. One leading method, RAFT [64], estimates flow through iterative updates of a flow field based on 4D correlation volumes. While optical flow methods allow for precise motion estimation between consecutive frames, they are not suited to long-range motion estimation: chaining pairwise optical flow into longer trajectories results in drift and fails to handle occlusions, while directly computing optical flow between distant frames (i.e., larger displacements) often results in temporal inconsistency [8, 73].

Multi-frame flow estimation methods [26, 28, 50, 68] can address some limitations of two-frame optical flow, but still struggle to handle long-range motion.

**Feature matching.** While optical flow methods are typically intended to operate on consecutive frames, other techniques can estimate dense correspondences between distant pairs of video frames [40]. Several methods learn such correspondences in a self- or weakly-supervised manner [5, 9, 12, 36, 51, 69, 71, 76] using cues like cycle consistency [27, 72, 81], while others [17, 29, 60, 66, 67] use stronger supervision signals such as geometric correspondences generated from 3D reconstruction pipelines [38, 54]. However, pairwise matching approaches typically do not incorporate temporal context, which can lead to inconsistent tracking over long videos and poor occlusion handling. In contrast, our method produces smooth trajectories through occlusions.

**Pixel-level long-range tracking.** A notable recent approach, PIPs [22], estimates multi-frame trajectories through occlusions by leveraging context within a small temporal window (8 frames). However, to produce motion for videos longer than this temporal window, PIPs still must chain correspondences, a process that (1) is prone to drift and (2) will lose track of points that remain occluded beyond the 8-frame window. Concurrent to our work, several works develop learning-based methods for predicting long-range pixel-level tracks in a feedforward manner. MFT [45] learns to select the most reliable sequences of optical flows to perform long-range tracking. TAPIR [15] tracks points by employing a matching stage inspired by TAP-Net [14] and a refinement stage inspired by PIPs [22]. CoTracker [30] proposes a flexible and powerful tracking algorithm with a transformer architecture to track points throughout a video. Our contribution is complementary to these works: the output of any of these methods can be used as the input supervision when optimizing a global motion representation.

**Video-based motion optimization.** Most conceptually related to our approach are classical methods that optimize motion globally over an entire video [2, 11, 35, 52, 53, 58, 61, 70]. Particle video, for instance, produces a set of semi-dense long-range trajectories (called *particles*) from initial optical flow fields [53]. However, it does not track through occlusions; an occluded entity will be treated as a different particle when it re-appears. Rubinstein et al. [52] further proposed a combinatorial assignment approach that can track through occlusion and generate longer trajectories. However, this method only produces semi-dense tracks for videos with simple motion, whereas our method estimates long-range motion for all pixels in general videos. Also related is ParticleSfM [80], which optimizes long-range correspondences from pairwise optical flows. Unlike our approach, ParticleSfM focuses on camera pose estimation within an SfM framework, where only correspondences from static regions

are optimized, and dynamic objects are treated as outliers.

**Neural video representations.** While our method shares similarities with recent methods that model videos using coordinate-based multi-layer perceptrons (MLPs) [43,56,63], prior work has primarily focused on problems such as novel view synthesis [37, 39, 46, 47, 75] and video decomposition [31, 79]. In contrast, our work targets the challenge of dense, long-range motion estimation. Though some methods for dynamic novel view synthesis produce 2D motion as a by-product, these systems require known camera poses and the resulting motion is often erroneous [20]. Some dynamic reconstruction methods [74, 77, 78] can also produce 2D motion, but these are often object-centric with a focus on articulated objects. Alternatively, video decomposition-based representations such as Layered Neural Atlases [31] and Deformable Sprites [79] solve for a mapping between each frame and a global texture atlas. Frame-to-frame correspondence can be derived by inverting this mapping, but this process is expensive and unreliable. Furthermore, these methods are limited to representing videos using a limited number of layers with fixed ordering, restricting their ability to model complex, real-world videos.

### 3. Overview

We propose a test-time optimization approach for estimating dense and long-range motion from a video sequence. Our method takes as input a collection of frames and pairwise noisy motion estimates (e.g., optical flow fields), and uses these to solve for a complete, globally consistent motion representation for the entire video. Once optimized, our representation can be queried with any pixel in any frame to produce a smooth, accurate motion trajectory across the full video. Our method identifies when points are occluded, and even tracks points through occlusions. In the following sections, we describe our underlying representation, dubbed *OmniMotion* (Sec. 4), then describe our optimization process (Sec. 5) for recovering this representation from a video.

### 4. OmniMotion representation

As discussed in Sec. 1, classical motion representations, such as pairwise optical flow, lose track of objects when they are occluded, and can produce inconsistencies when correspondences are composed over multiple frames. To obtain accurate, consistent tracks even through occlusion, we therefore need a *global* motion representation, i.e., a data structure that encodes the trajectories of all points in a scene. One such global representation is a decomposition of a scene into a set of discrete, depth-separated layers [31, 79]. However, most real-world scenes cannot be represented as a set of fixed, ordered layers: e.g., consider the simple case of an object rotating in 3D. At the other extreme is full 3D reconstruction that disentangles 3D scene geometry, camera

pose and scene motion. This, however, is an extremely ill-posed problem. Thus, we ask: can we accurately track real-world motion without explicit dynamic 3D reconstruction?

We answer this question using our proposed representation, *OmniMotion* (illustrated in Fig. 2). *OmniMotion* represents the scene in a video as a canonical 3D volume that is mapped to local volumes for each frame through local-canonical bijections. The local-canonical bijections are parametrized as neural networks and capture both camera and scene motion without disentangling the two. As such, the video can be considered as a rendering of the resulting local volumes from a fixed, static camera.

Because *OmniMotion* does not explicitly disentangle camera and scene motion, the resulting representation is not a physically accurate 3D scene reconstruction. Instead we call it a *quasi-3D* representation. This relaxation of dynamic multi-view geometry allows us to sidestep ambiguities that make dynamic 3D reconstruction challenging. Yet we retain properties needed for consistent and accurate long-term tracking through occlusion: first, by establishing bijections between each local frame and a canonical frame, *OmniMotion* guarantees globally cycle-consistent 3D mappings across all local frames, which emulates the one-to-one correspondences between real-world, metric 3D reference frames. Second, *OmniMotion* retains information about all scene points that are projected onto each pixel, along with their relative depth ordering, enabling points to be tracked even when they are temporarily occluded from view.

In the following sections, we describe our quasi-3D canonical volume and 3D bijections, and then describe how they can be used to compute motion between any two frames.

#### 4.1. Canonical 3D volume

We represent a video’s content using a canonical volume  $G$  that acts as a three-dimensional atlas of the observed scene. As in NeRF [43], we define a coordinate-based network  $F_{\theta}$  over  $G$  that maps each canonical 3D coordinate  $\mathbf{u} \in G$  to a density  $\sigma$  and color  $\mathbf{c}$ . The density stored in  $G$  is key, as it tells us where the surfaces are in canonical space. Together with the 3D bijections, this allows us to track surfaces over multiple frames as well as reason about occlusion relationships. The color stored in  $G$  allows us to compute a photometric loss during optimization.

#### 4.2. 3D bijections

We define a continuous bijective mapping  $\mathcal{T}_i$  that maps 3D points  $\mathbf{x}_i$  from each local coordinate frame  $L_i$  to the canonical 3D coordinate frame as  $\mathbf{u} = \mathcal{T}_i(\mathbf{x}_i)$ , where  $i$  is the frame index. Note that the canonical coordinate  $\mathbf{u}$  is time-independent and can be viewed as a globally consistent “index” for a particular scene point or 3D trajectory across time. By composing these bijective mappings and their inverses, we can map a 3D point from one local 3D coordinate



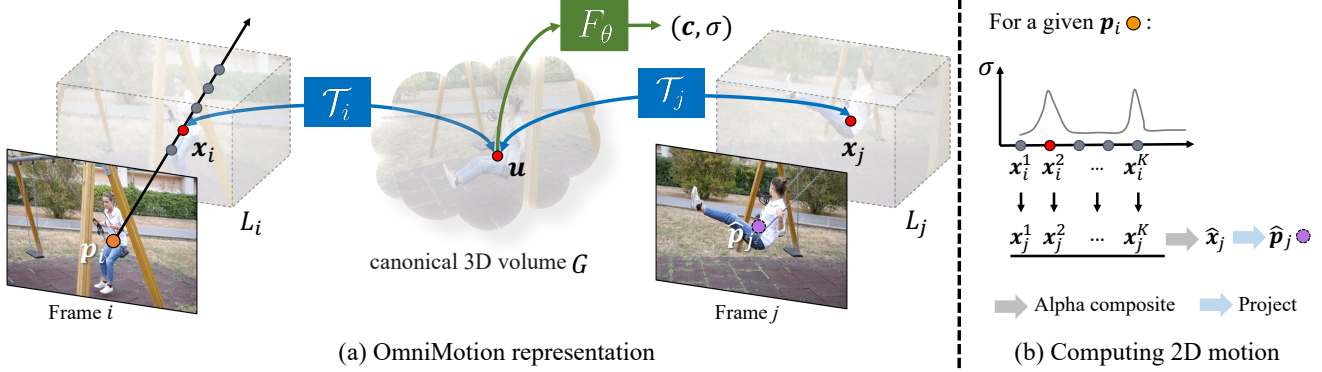


Figure 2: *Method overview.* (a) Our OmniMotion representation is comprised of a canonical 3D volume  $G$  and a set of bijections  $\mathcal{T}_i$  that map between each frame’s local volume  $L_i$  and the canonical volume  $G$ . Any local 3D location  $\mathbf{x}_i$  in frame  $i$  can be mapped to its corresponding canonical location  $\mathbf{u}$  through  $\mathcal{T}_i$ , and then mapped back to another frame  $j$  as  $\mathbf{x}_j$  through the inverse mapping  $\mathcal{T}_j^{-1}$ . Each location  $\mathbf{u}$  in  $G$  is associated with a color  $\mathbf{c}$  and density  $\sigma$ , computed using a coordinate-based MLP  $F_\theta$ . (b) To compute the corresponding 2D location for a given query point  $\mathbf{p}_i$  mapped from frame  $i$  to  $j$ , we shoot a ray into  $L_i$  and sample a set of points  $\{\mathbf{x}_i^k\}_{k=1}^K$ , which are then mapped first to the canonical space to obtain their densities, and then to frame  $j$  to compute their corresponding local 3D locations  $\{\mathbf{x}_j^k\}_{k=1}^K$ . These points  $\{\mathbf{x}_j^k\}_{k=1}^K$  are then alpha-composited and projected to obtain the 2D corresponding location  $\hat{\mathbf{p}}_j$ .

frame ( $L_i$ ) to another ( $L_j$ ):

$$\mathbf{x}_j = \mathcal{T}_j^{-1} \circ \mathcal{T}_i(\mathbf{x}_i). \quad (1)$$

Bijjective mappings ensure that the resulting correspondences between 3D points in individual frames are all cycle consistent, as they arise from the same canonical point.

To allow for expressive maps that can capture real-world motion, we parameterize these bijections as invertible neural networks (INNs). Inspired by recent work on homeomorphic shape modeling [34, 48], we use Real-NVP [13] due to its simple formulation and analytic invertibility. Real-NVP builds bijective mappings by composing simple bijective transformations called affine coupling layers. An affine coupling layer splits the input into two parts; the first part is left unchanged, but is used to parametrize an affine transformation that is applied to the second part.

We modify this architecture to also condition on a per-frame latent code  $\psi_i$  [34, 48]. Then all invertible mappings  $\mathcal{T}_i$  are parameterized by the same invertible network  $M_\theta$ , but with different latent codes:  $\mathcal{T}_i(\cdot) = M_\theta(\cdot; \psi_i)$ .

### 4.3. Computing frame-to-frame motion

Given this representation, we now describe how we can compute 2D motion for any query pixel  $\mathbf{p}_i$  in frame  $i$ . Intuitively, we “lift” the query pixel to 3D by sampling points on a ray, “map” these 3D points to a target frame  $j$  using bijections  $\mathcal{T}_i$  and  $\mathcal{T}_j$ , “render” these mapped 3D points from the different samples through alpha compositing, and finally “project” back to 2D to obtain a putative correspondence.

Specifically, since we assume that camera motion is subsumed by the local-canonical bijections  $\mathcal{T}_i$ , we simply use a

fixed, orthographic camera. The ray at  $\mathbf{p}_i$  can then be defined as  $\mathbf{r}_i(z) = \mathbf{o}_i + z\mathbf{d}$ , where  $\mathbf{o}_i = [\mathbf{p}_i, 0]$  and  $\mathbf{d} = [0, 0, 1]$ . We sample  $K$  samples on the ray  $\{\mathbf{x}_i^k\}$ , which are equivalent to appending a set of depth values  $\{z_i^k\}_{k=1}^K$  to  $\mathbf{p}_i$ . Despite not being a physical camera ray, it captures the notion of multiple surfaces at each pixel and suffices to handle occlusion.

Next we obtain densities and colors for these samples by mapping them to the canonical space and then querying the density network  $F_\theta$ . Taking the  $k$ -th sample  $\mathbf{x}_i^k$  as an example, its density and color can be written as  $(\sigma_k, \mathbf{c}_k) = F_\theta(M_\theta(\mathbf{x}_i^k; \psi_i))$ . We can also map each sample along the ray to a corresponding 3D location  $\mathbf{x}_j^k$  in frame  $j$  (Eq. 1).

We can now aggregate the correspondences  $\mathbf{x}_j^k$  from all samples to produce a single correspondence  $\hat{\mathbf{x}}_j$ . This aggregation is similar to how the colors of sample points are aggregated in NeRF: we use alpha compositing, with the alpha value for the  $k$ -th sample as  $\alpha_k = 1 - \exp(-\sigma_k)$ . We then compute  $\hat{\mathbf{x}}_j$  as:

$$\hat{\mathbf{x}}_j = \sum_{k=1}^K T_k \alpha_k \mathbf{x}_j^k, \text{ where } T_k = \prod_{l=1}^{k-1} (1 - \alpha_l) \quad (2)$$

A similar process is used to composite  $\mathbf{c}^k$  to get the image-space color  $\hat{\mathbf{C}}_i$  for  $\mathbf{p}_i$ .  $\hat{\mathbf{x}}_j$  is then projected using our stationary orthographic camera model to yield the predicted 2D corresponding location  $\hat{\mathbf{p}}_j$  for the query location  $\mathbf{p}_i$ .

## 5. Optimization

Our optimization process takes as input a video sequence and a collection of noisy correspondence predictions (from an existing method) as guidance, and generates a complete, globally consistent motion estimate for the entire video.

### 5.1. Collecting input motion data

For most of our experiments, we use RAFT [64] to compute input pairwise correspondence. We also experimented with another dense correspondence method, TAP-Net [14], and demonstrate in our evaluation that our approach consistently works well given different types of input correspondence. Taking RAFT as an example, we begin by exhaustively computing all pairwise optical flows. Since optical flow methods can produce significant errors under large displacements, we apply cycle consistency and appearance consistency checks to filter out spurious correspondences. We also optionally augment the flows through chaining, when deemed reliable. Additional details about our flow collection process are provided in the supplemental material. Despite the filtering, the (now incomplete) flow fields remain noisy and inconsistent. We now introduce our optimization method that consolidates these noisy, incomplete pairwise motion into complete and accurate long-range motion.

### 5.2. Loss functions

Our primary loss function is a flow loss. We minimize the mean absolute error (MAE) between the predicted flow  $\hat{\mathbf{f}}_{i \rightarrow j} = \hat{\mathbf{p}}_j - \mathbf{p}_i$  from our optimized representation and the supervising input flow  $\mathbf{f}_{i \rightarrow j}$  derived from running optical flow:

$$\mathcal{L}_{\text{flo}} = \sum_{\mathbf{f}_{i \rightarrow j} \in \Omega_f} \|\hat{\mathbf{f}}_{i \rightarrow j} - \mathbf{f}_{i \rightarrow j}\|_1 \quad (3)$$

where  $\Omega_f$  is the set of all the filtered pairwise flows. In addition, we minimize a photometric loss defined as the mean squared error (MSE) between the predicted color  $\hat{C}_i$  and the observed color  $C_i$  in the source video frame:

$$\mathcal{L}_{\text{pho}} = \sum_{(i, \mathbf{p}) \in \Omega_p} \|\hat{C}_i(\mathbf{p}) - C_i(\mathbf{p})\|_2^2 \quad (4)$$

where  $\Omega_p$  is the set of all pixel locations over all frames. Last, to ensure temporal smoothness of the 3D motion estimated by  $M_\theta$ , we apply a regularization term that penalizes large accelerations. Given a sampled 3D location  $\mathbf{x}_i$  in frame  $i$ , we map it to frame  $i - 1$  and frame  $i + 1$  using Eq. 1, yielding 3D points  $\mathbf{x}_{i-1}$  and  $\mathbf{x}_{i+1}$  respectively, and then minimize 3D acceleration as in [37]:

$$\mathcal{L}_{\text{reg}} = \sum_{(i, \mathbf{x}) \in \Omega_x} \|\mathbf{x}_{i+1} + \mathbf{x}_{i-1} - 2\mathbf{x}_i\|_1 \quad (5)$$

where  $\Omega_x$  is the union of local 3D spaces for all frames. Our final combined loss can be written as:

$$\mathcal{L} = \mathcal{L}_{\text{flo}} + \lambda_{\text{pho}} \mathcal{L}_{\text{pho}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}} \quad (6)$$

where weights  $\lambda$  control the relative importance of each term.

The intuition behind this optimization is to leverage the bijections to a single canonical volume  $G$ , photo consistency, and the natural spatiotemporal smoothness provided

by the coordinate-based networks  $M_\theta$  and  $F_\theta$  to reconcile inconsistent pairwise flow and fill in missing content in the correspondence graphs.

### 5.3. Balancing supervision via hard mining

The exhaustive pairwise flow input maximizes the useful motion information available to the optimization stage. However, this approach, especially when coupled with the flow-filtering process, can result in an unbalanced collection of motion samples in dynamic regions. Rigid background regions typically have many reliable pairwise correspondences, while fast-moving and deforming foreground objects often have many fewer reliable correspondences after filtering, especially between distant frames. This imbalance can lead the network to focus entirely on dominant (simple) background motions, and ignore the challenging moving objects that represent a small portion of the supervisory signal.

To address this issue, we propose a simple strategy for mining hard examples during training. Specifically, we periodically cache flow predictions and compute error maps by calculating the Euclidean distance between the predicted and input flows. During optimization, we guide sampling such that regions with high errors are sampled more frequently. We compute these error maps on consecutive frames, where we assume our supervisory optical flow is most reliable. Please refer to the supplement for more details.

### 5.4. Implementation details

**Network.** Our mapping network  $M_\theta$  consists of six affine coupling layers. We apply positional encoding [43, 63] with 4 frequencies to each layer’s input coordinates before computing the scale and translation. We use a single 2-layer MLP with 256 channels implemented as a GaborNet [18] to compute the latent code  $\psi_i$  for each frame  $i$ . The input to this MLP is the time  $t_i$ . The dimensionality of latent code  $\psi_i$  is 128. The canonical representation  $F_\theta$  is also implemented as a GaborNet, but with 3 layers of 512 channels.

**Representation.** We normalize all pixel coordinates  $\mathbf{p}_i$  to the range  $[-1, 1]$ , and set the near and far depth range to 0 and 2, defining a local 3D space for each frame as  $[-1, 1]^2 \times [0, 2]$ . While our method can place content at arbitrary locations in the canonical space  $G$ , we initialize mapped canonical locations given by  $M_\theta$  to be roughly within a unit sphere to ensure well-conditioned input to the density/color network  $F_\theta$ . To improve numerical stability during training, we apply the contraction operation in mip-NeRF 360 [3] to canonical 3D coordinates  $\mathbf{u}$  before passing them to  $F_\theta$ .

**Training.** We train our representation on each video sequence with Adam [32] for 200K iterations. In each training batch, we sample 256 correspondences from 8 pairs of images, for a total of 1024 correspondences. We sample

$K = 32$  points on each ray using stratified sampling. More training details can be found in the supplemental material.

## 6. Evaluation

### 6.1. Benchmarks

We evaluate our method on the TAP-Vid benchmark [14], which is designed to evaluate the performance of point tracking across long video clips. TAP-Vid consists of both real-world videos with accurate human annotations of point tracks and synthetic videos with perfect ground-truth point tracks. Each point track is annotated through the entire video, and is labeled as occluded when not visible.

**Datasets.** We evaluate on the following datasets from TAP-Vid: 1) **DAVIS**, a real dataset of 30 videos from the DAVIS 2017 validation set [49], with clips ranging from 34-104 frames and an average of 21.7 point track annotations per video. 2) **Kinetics**, a real dataset of 1,189 videos each with 250 frames from the Kinetics-700-2020 validation set [10] with an average of 26.3 point track annotations per video. To make evaluation tractable for test-time optimization approaches like ours, we randomly sample a subset of 100 videos and evaluate all methods on this subset. 3) **RGB-Stacking** [33], a synthetic dataset of 50 videos each with 250 frames and 30 tracks. We exclude the synthetic Kubric dataset [21] as it is primarily intended for training. For quantitative evaluation, we adhere to the TAP benchmark protocol and evaluate all methods at  $256 \times 256$  resolution, but all qualitative results are run at a higher resolution (480p).

**Evaluation metrics.** Following the TAP-Vid benchmark, we report both the position and occlusion accuracy of predicted tracks. We also introduce a new metric measuring temporal coherence. Our evaluation metrics include:

- $\delta_{\text{avg}}^x$  measures the average position accuracy of visible points across 5 thresholds: 1, 2, 4, 8, and 16 pixels. The accuracy  $< \delta^x$  at each threshold  $\delta^x$  is the fraction of points that are within  $\delta^x$  pixels of their ground truth position.
- **Average Jaccard (AJ)** evaluates both occlusion and position accuracy on the same thresholds as  $< \delta_{\text{avg}}^x$ . It categorizes predicted point locations as true positives, false positives, and false negatives, and is defined as the ratio of true positives to all points. True positives are points within the threshold of visible ground truth points. False positives are points that are predicted as visible, but where the ground truth is occluded or beyond the threshold. False negatives are ground truth visible points that are predicted as occluded or are beyond the threshold.
- **Occlusion Accuracy (OA)** evaluates the accuracy of the visibility/occlusion prediction at each frame.
- **Temporal Coherence (TC)** evaluates the temporal coherence of the tracks by measuring the  $L_2$  distance between

the acceleration of groundtruth tracks and predicted tracks. The acceleration is measured as the flow difference between two adjacent frames for visible points.

### 6.2. Baselines

We compare OmniMotion to various types of dense correspondence methods, including optical flow, feature matching, and multi-frame trajectory estimation as follows:

**RAFT** [64] is a state-of-the-art two-frame flow method. We consider two ways to use RAFT to generate multi-frame trajectories at test time: 1) chaining RAFT predictions between consecutive frames into longer tracks, which we call **RAFT-C**, and 2) directly computing RAFT flow between any (non-adjacent) query and target frames (**RAFT-D**). When generating trajectories using RAFT-D, we always use the previous flow prediction as initialization for the current frame.

**PIPs** [22] is a method for estimating multi-frame point trajectories that can handle occlusions. By default, the method uses a temporal window of 8 frames, and longer trajectories must be generated through chaining. We used the official implementation of PIPs to perform chaining.

**Flow-Walk** [5] uses a multi-scale contrastive random walk to learn space-time correspondences by encouraging cycle consistency across time. Similar to RAFT, we report both chained and direct correspondence computation as **Flow-Walk-C** and **Flow-Walk-D**, respectively.

**TAP-Net** [14] uses a cost volume to predict the location of a query point in a single target frame, along with a scalar occlusion logit.

**Deformable Sprites** [79] is a layer-based video decomposition method. Like our method, it uses a per-video test-time optimization. However, it does not directly produce frame-to-frame correspondence, as the mapping from each frame to a canonical texture image is non-invertible. A nearest neighbor search in texture image space is required to find correspondence. Layered Neural Atlases [31] shares similarities to Deformable Sprites, but requires semantic segmentation masks as input, so we opt to compare to Deformable Sprites.

PIPs, TAP-Net and Deformable Sprites directly predict occlusion, but RAFT and Flow-Walk do not. Therefore we follow prior work [14, 76] and use a cycle consistency check with a threshold of 48 pixels to produce occlusion predictions for these methods. For our method, we detect occlusion by first mapping the query point to its corresponding 3D location in the target frame, then checking the transmittance of that 3D location in the target frame.

### 6.3. Comparisons

**Quantitative comparisons.** We compare our method to baselines on the TAP-Vid benchmark in Tab. 1. Our method achieves the best position accuracy, occlusion accuracy, and






Figure 3: *Qualitative comparison of our method and baselines on DAVIS [49].* The leftmost image shows query points in the first frame, and the right three images show tracking results over time. Notably, our method tracks successfully through the occlusion events in *swing* and *india*, while baseline methods fail. Our method additionally detects occlusion (marked as cross “+”) and gives plausible location estimates even when a point is occluded. Please refer to the supplemental video for better visualizations of tracking accuracy and coherence.

temporal coherence consistently across different datasets. Our method works well with different input pairwise correspondences from RAFT and TAP-Net, and provides consistent improvements over both of these base methods.

Compared to approaches that directly operate on (non-adjacent) pairs of query and target frames like RAFT-D, TAP-Net, and Flow-Walk-D, our method achieves significantly better temporal coherence due to our globally consistent representation. Compared to flow chaining methods like RAFT-C, PIPs, and Flow-Walk-C, our method has better tracking performance, especially on longer videos. Chaining methods accumulate errors over time and are not robust to oc-

clusion. Although PIPs considers a wider temporal window (8 frames) for better occlusion handling, it fails to track a point if it stays occluded beyond the entire temporal window. In contrast, OmniMotion can track points through extended occlusions. Our method also outperforms the test-time optimization approach Deformable Sprites [79]. Deformable Sprites decomposes a video using a predefined two or three layers with fixed ordering and models the background as a 2D atlas with a per-frame homography, limiting its ability to fit to videos with complex camera and scene motion.

**Qualitative comparisons.** We compare our method qualitatively to our baselines in Fig. 3. We highlight our ability to

Method	Kinetics				DAVIS				RGB-Stacking			
	AJ $\uparrow$	$< \delta_{\text{avg}}^x \uparrow$	OA $\uparrow$	TC $\downarrow$	AJ $\uparrow$	$< \delta_{\text{avg}}^x \uparrow$	OA $\uparrow$	TC $\downarrow$	AJ $\uparrow$	$< \delta_{\text{avg}}^x \uparrow$	OA $\uparrow$	TC $\downarrow$
RAFT-C [64]	31.7	51.7	84.3	0.82	30.7	46.6	80.2	0.93	42.0	56.4	91.5	0.18
RAFT-D [64]	50.6	66.9	85.5	3.00	34.1	48.9	76.1	9.83	72.1	<u>85.1</u>	92.1	1.04
TAP-Net [14]	48.5	61.7	86.6	6.65	38.4	53.4	81.4	10.82	61.3	73.7	91.5	1.52
PIPs [22]	39.1	55.3	82.9	1.30	39.9	56.0	81.3	1.78	37.3	50.6	89.7	0.84
Flow-Walk-C [5]	40.9	55.5	84.5	<u>0.77</u>	35.2	51.4	80.6	0.90	41.3	55.7	<u>92.2</u>	<u>0.13</u>
Flow-Walk-D [5]	46.9	65.9	81.8	3.04	24.4	40.9	76.5	10.41	66.3	82.7	91.2	0.47
Deformable-Sprites [79]	25.6	39.5	71.4	1.70	20.6	32.9	69.7	2.07	45.0	58.3	84.0	0.99
Ours (TAP-Net)	<u>53.8</u>	<u>68.3</u>	<u>88.8</u>	<u>0.77</u>	<u>50.9</u>	<u>66.7</u>	<b>85.7</b>	<u>0.86</u>	<u>73.4</u>	84.1	<u>92.2</u>	<b>0.11</b>
Ours (RAFT)	<b>55.1</b>	<b>69.6</b>	<b>89.6</b>	<b>0.76</b>	<b>51.7</b>	<b>67.5</b>	<u>85.3</u>	<b>0.74</b>	<b>77.5</b>	<b>87.0</b>	<b>93.5</b>	<u>0.13</u>

Table 1: Quantitative comparison of our method and baselines on the TAP-Vid benchmark [14]. We refer to our method as *Ours*, and present two variants, *Ours (TAP-Net)* and *Ours (RAFT)*, which are optimized using input pairwise correspondences from TAP-Net [14] and RAFT [64], respectively. Both *Ours* and *Deformable Sprites* [79] estimate global motion via test-time optimization on each individual video, while all other methods estimate motion locally in a feed-forward manner. Our method notably improves the quality of the input correspondences, achieving the best position accuracy, occlusion accuracy, and temporal coherence among all methods tested.

Method	AJ $\uparrow$	$< \delta_{\text{avg}}^x \uparrow$	OA $\uparrow$	TC $\downarrow$
No invertible	12.5	21.4	76.5	0.97
No photometric	42.3	58.3	84.1	0.83
Uniform sampling	47.8	61.8	83.6	0.88
Full	<b>51.7</b>	<b>67.5</b>	<b>85.3</b>	<b>0.74</b>

Table 2: Ablation study on DAVIS [49].

identify and track through (long) occlusion events while also providing plausible locations for points during occlusion, as well as handling large camera motion parallax. Please refer to the supplementary video for animated comparisons.

#### 6.4. Ablations and analysis

**Ablations.** We perform ablations to verify the effectiveness of our design decisions in Tab. 2. *No invertible* is a model variant that replaces our invertible mapping network  $M_\theta$  with separate forward and backward mapping networks between local and canonical frames (i.e., without the strict cycle consistency guarantees of our proposed bijections). While we additionally add a cycle consistency loss for this ablation, it still fails to construct a meaningful canonical space, and can only represent simple motions of the static background. *No photometric* is a version that omits the photometric loss  $\mathcal{L}_{\text{pho}}$ ; the reduced performance suggests the importance of photoconsistency for refining motion estimates. *Uniform sampling* replaces our hard-mining sampling strategy with a uniform sampling strategy, which we found leads to an inability to capture fast motion.

**Analysis.** In Fig. 4, we show pseudo-depth maps generated from our model to demonstrate the learned depth ordering. Note that these maps do not correspond to physical depth, nonetheless, they demonstrate that using only photometric and flow signals, our method is able to sort out the relative




Figure 4: Pseudo-depth maps extracted from our representation, where blue indicates closer objects and red indicates further.

ordering between different surfaces, which is crucial for tracking through occlusions. More ablations and analyses are provided in the supplemental material.

## 7. Limitations

Like many motion estimation methods, our method struggles with rapid and highly non-rigid motion as well as thin structures. In these scenarios, pairwise correspondence methods can fail to provide enough reliable correspondences for our method to compute accurate global motion.

In addition, due to the highly non-convex nature of the underlying optimization problem, we observe that our optimization process can be sensitive to initialization for certain difficult videos. This can result in sub-optimal local minima, e.g., incorrect surface ordering or duplicated objects in canonical space that can sometimes be hard to correct through optimization.

Finally, our method in its current form can be computationally expensive. First, the flow collection process involves computing all pairwise flows exhaustively, which scales quadratically with respect to the sequence length. However, we believe the scalability of this process can be improved by



exploring more efficient alternatives to exhaustive matching, e.g., vocabulary trees or keyframe-based matching, drawing inspiration from the Structure from Motion and SLAM literature. Second, like other methods that utilize neural implicit representations [43], our method involves a relatively long optimization process. Recent research in this area [44, 62] may help accelerate this process and allow further scaling to longer sequences.

## 8. Conclusion

We proposed a new test-time optimization method for estimating complete and globally consistent motion for an entire video. We introduced a new video motion representation called OmniMotion which includes a quasi-3D canonical volume and per-frame local-canonical bijections. OmniMotion can handle general videos with varying camera setups and scene dynamics, and produce accurate and smooth long-range motion through occlusions. Our method achieves significant improvements over prior state-of-the-art methods both qualitatively and quantitatively.

**Acknowledgements.** We thank Jon Barron, Richard Tucker, Vickie Ye, Zekun Hao, Xiaowei Zhou, Steve Seitz, Brian Curless, and Richard Szeliski for their helpful input and assistance. This work was supported in part by an NVIDIA academic hardware grant and by the National Science Foundation (IIS-2008313).

## References

- [1] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building Rome in a day. *Communications of the ACM*, 54(10):105–112, 2011. [2](#)
- [2] Vijay Badrinarayanan, Fabio Galasso, and Roberto Cipolla. Label propagation in video sequences. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 3265–3272. IEEE, 2010. [2](#)
- [3] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. mip-NeRF 360: Unbounded anti-aliased neural radiance fields. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 5470–5479, 2022. [5](#)
- [4] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer vision and image understanding*, 110(3):346–359, 2008. [2](#)
- [5] Zhangxing Bian, Allan Jabri, Alexei A Efros, and Andrew Owens. Learning pixel trajectories with multiscale contrastive random walks. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 6508–6519, 2022. [2](#), [6](#), [8](#)
- [6] Michael J Black and Padmanabhan Anandan. A framework for the robust estimation of optical flow. In *Proc. Int. Conf. on Computer Vision (ICCV)*, pages 231–236. IEEE, 1993. [2](#)
- [7] Thomas Brox, Christoph Bregler, and Jitendra Malik. Large displacement optical flow. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 41–48, 2009. [2](#)
- [8] Thomas Brox and Jitendra Malik. Large displacement optical flow: descriptor matching in variational motion estimation. *Trans. Pattern Analysis and Machine Intelligence*, 33(3):500–513, 2010. [2](#)
- [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proc. Int. Conf. on Computer Vision (ICCV)*, pages 9650–9660, 2021. [2](#)
- [10] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 6299–6308, 2017. [6](#)
- [11] Jason Chang, Donglai Wei, and John W Fisher. A video representation using temporal superpixels. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 2051–2058, 2013. [2](#)
- [12] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proc. Computer Vision and Pattern Recognition Workshops*, pages 224–236, 2018. [2](#)
- [13] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using Real NVP. *arXiv preprint arXiv:1605.08803*, 2016. [4](#)
- [14] Carl Doersch, Ankush Gupta, Larisa Markeeva, Adria Recasens Contente, Kucas Smaira, Yusuf Aytar, Joao Carreira, Andrew Zisserman, and Yi Yang. Tap-vid: A benchmark for tracking any point in a video. In *NeurIPS Datasets Track*, 2022. [2](#), [5](#), [6](#), [8](#)
- [15] 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. *arXiv preprint arXiv:2306.08637*, 2023. [2](#)
- [16] Hugh Durrant-Whyte and Tim Bailey. Simultaneous localization and mapping: part i. *IEEE Robotics & Automation Magazine*, 13(2):99–110, 2006. [2](#)
- [17] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-net: A trainable cnn for joint description and detection of local features. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 8092–8101, 2019. [2](#)
- [18] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative filter networks. In *International Conference on Learning Representations*, 2021. [5](#)
- [19] Philipp Fischer, Alexey Dosovitskiy, Eddy Ilg, Philip Häusser, Caner Hazırbaş, Vladimir Golkov, Patrick Van der Smagt, Daniel Cremers, and Thomas Brox. Flownet: Learning optical flow with convolutional networks. *arXiv preprint arXiv:1504.06852*, 2015. [2](#)
- [20] Hang Gao, Ruilong Li, Shubham Tulsiani, Bryan Russell, and Angjoo Kanazawa. Monocular dynamic view synthesis: A reality check. In *Neural Information Processing Systems*, 2022. [3](#)
- [21] 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 *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 3749–3761, 2022. 6
- [22] Adam W Harley, Zhaoyuan Fang, and Katerina Fragkiadaki. Particle video revisited: Tracking through occlusions using point trajectories. In *Proc. European Conf. on Computer Vision (ECCV)*, pages 59–75. Springer, 2022. 1, 2, 6, 8
- [23] Berthold KP Horn and Brian G Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. 2
- [24] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. Liteflownet: A lightweight convolutional neural network for optical flow estimation. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 8981–8989, 2018. 2
- [25] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 2462–2470, 2017. 2
- [26] Michal Irani. Multi-frame optical flow estimation using sub-space constraints. In *Proc. Int. Conf. on Computer Vision (ICCV)*, volume 1, pages 626–633, 1999. 2
- [27] Allan Jabri, Andrew Owens, and Alexei Efros. Space-time correspondence as a contrastive random walk. In *Neural Information Processing Systems*, pages 19545–19560, 2020. 2
- [28] Joel Janai, Fatma Guney, Anurag Ranjan, Michael Black, and Andreas Geiger. Unsupervised learning of multi-frame optical flow with occlusions. In *Proc. European Conf. on Computer Vision (ECCV)*, pages 690–706, 2018. 2
- [29] Wei Jiang, Eduard Trulls, Jan Hosang, Andrea Tagliasacchi, and Kwang Moo Yi. Cotr: Correspondence transformer for matching across images. In *Proc. Int. Conf. on Computer Vision (ICCV)*, pages 6207–6217, 2021. 2
- [30] 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
- [31] Yoni Kasten, Dolev Ofri, Oliver Wang, and Tali Dekel. Layered neural atlases for consistent video editing. In *ACM Trans. Graphics (SIGGRAPH Asia)*, 2021. 3, 6
- [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [33] Alex X Lee, Coline Manon Devin, Yuxiang Zhou, Thomas Lampe, Konstantinos Bousmalis, Jost Tobias Springenberg, Arunkumar Byravan, Abbas Abdolmaleki, Nimrod Gileadi, David Khosid, et al. Beyond pick-and-place: Tackling robotic stacking of diverse shapes. In *Proc. Conference on Robot Learning*, 2021. 6
- [34] Jiahui Lei and Kostas Daniilidis. Cadex: Learning canonical deformation coordinate space for dynamic surface representation via neural homeomorphism. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 6624–6634, 2022. 4
- [35] José Lezama, Karteek Alahari, Josef Sivic, and Ivan Laptev. Track to the future: Spatio-temporal video segmentation with long-range motion cues. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 3369–3376. IEEE, 2011. 2
- [36] Xueting Li, Sifei Liu, Shalini De Mello, Xiaolong Wang, Jan Kautz, and Ming-Hsuan Yang. Joint-task self-supervised learning for temporal correspondence. In *Neural Information Processing Systems*, 2019. 2
- [37] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 6498–6508, 2021. 3, 5
- [38] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 2041–2050, 2018. 2
- [39] Zhengqi Li, Qianqian Wang, Forrester Cole, Richard Tucker, and Noah Snavely. Dynibar: Neural dynamic image-based rendering. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [40] Ce Liu, Jenny Yuen, and Antonio Torralba. Sift flow: Dense correspondence across scenes and its applications. *Trans. Pattern Analysis and Machine Intelligence*, 33(5):978–994, 2010. 2
- [41] David G Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. of Computer Vision*, 60:91–110, 2004. 2
- [42] Bruce D Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. In *International Joint Conference on Artificial Intelligence*, volume 2, pages 674–679, 1981. 2
- [43] 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. 3, 5, 9
- [44] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. In *ACM Trans. Graphics (SIGGRAPH)*, 2022. 9
- [45] Michal Neoral, Jonáš Šerých, and Jiří Matas. Mft: Long-term tracking of every pixel. *arXiv preprint arXiv:2305.12998*, 2023. 2
- [46] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 5865–5874, 2021. 3
- [47] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. HyperNeRF: A higher-dimensional representation for topologically varying neural radiance fields. In *ACM Trans. Graphics (SIGGRAPH)*, 2021. 3
- [48] Despoina Paschalidou, Angelos Katharopoulos, Andreas Geiger, and Sanja Fidler. Neural parts: Learning expressive 3d shape abstractions with invertible neural networks. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 3204–3215, 2021. 4
- [49] 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:1704.00675*, 2017. 6, 7, 8
- [50] 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 *Proc. Winter Conf. on Computer Vision (WACV)*, pages 2077–2086, 2019. 2
- [51] Ignacio Rocco, Mircea Cimpoi, Relja Arandjelović, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Neighbourhood consensus networks. In *Neural Information Processing Systems*, volume 31, 2018. 2
- [52] Michael Rubinstein, Ce Liu, and William T Freeman. Towards longer long-range motion trajectories. In *Proc. British Machine Vision Conf. (BMVC)*, 2012. 2
- [53] Peter Sand and Seth Teller. Particle video: Long-range motion estimation using point trajectories. *Int. J. of Computer Vision*, 80:72–91, 2008. 1, 2
- [54] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 4104–4113, 2016. 2
- [55] Jianbo Shi et al. Good features to track. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 593–600. IEEE, 1994. 2
- [56] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Neural Information Processing Systems*, volume 33, pages 7462–7473, 2020. 3
- [57] Noah Snavely, Steven M Seitz, and Richard Szeliski. Modeling the world from internet photo collections. *Int. J. of Computer Vision*, 80:189–210, 2008. 2
- [58] Deqing Sun, Erik Sudderth, and Michael Black. Layered image motion with explicit occlusions, temporal consistency, and depth ordering. In *Neural Information Processing Systems*, volume 23, 2010. 2
- [59] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 8934–8943, 2018. 2
- [60] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 8922–8931, 2021. 2
- [61] Narayanan Sundaram, Thomas Brox, and Kurt Keutzer. Dense point trajectories by GPU-accelerated large displacement optical flow. In *Proc. European Conf. on Computer Vision (ECCV)*, pages 438–451. Springer, 2010. 2
- [62] Matthew Tancik, Vincent Casser, Xinchen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 8248–8258, 2022. 9
- [63] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Neural Information Processing Systems*, volume 33, pages 7537–7547, 2020. 3, 5
- [64] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Proc. European Conf. on Computer Vision (ECCV)*, pages 402–419, 2020. 2, 5, 6, 8
- [65] Carlo Tomasi and Takeo Kanade. Detection and tracking of point. *Int. J. of Computer Vision*, 9:137–154, 1991. 2
- [66] Prune Truong, Martin Danelljan, and Radu Timofte. Glu-net: Global-local universal network for dense flow and correspondences. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 6258–6268, 2020. 2
- [67] Prune Truong, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning accurate dense correspondences and when to trust them. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 5714–5724, 2021. 2
- [68] Sebastian Volz, Andres Bruhn, Levi Valgaerts, and Henning Zimmer. Modeling temporal coherence for optical flow. In *Proc. Int. Conf. on Computer Vision (ICCV)*, pages 1116–1123. IEEE, 2011. 2
- [69] Carl Vondrick, Abhinav Shrivastava, Alireza Fathi, Sergio Guadarrama, and Kevin Murphy. Tracking emerges by coloring videos. In *Proc. European Conf. on Computer Vision (ECCV)*, pages 391–408, 2018. 2
- [70] John YA Wang and Edward H Adelson. Layered representation for motion analysis. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 361–366, 1993. 2
- [71] Qianqian Wang, Xiaowei Zhou, Bharath Hariharan, and Noah Snavely. Learning feature descriptors using camera pose supervision. In *Proc. European Conf. on Computer Vision (ECCV)*, pages 757–774, 2020. 2
- [72] Xiaolong Wang, Allan Jabri, and Alexei A Efros. Learning correspondence from the cycle-consistency of time. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 2566–2576, 2019. 2
- [73] Philippe Weinzaepfel, Jerome Revaud, Zaid Harchaoui, and Cordelia Schmid. DeepFlow: Large displacement optical flow with deep matching. In *Proc. Int. Conf. on Computer Vision (ICCV)*, pages 1385–1392, 2013. 2
- [74] Yuefan Wu, Zeyuan Chen, Shaowei Liu, Zhongzheng Ren, and Shenlong Wang. Casa: Category-agnostic skeletal animal reconstruction. *Advances in Neural Information Processing Systems*, 35:28559–28574, 2022. 3
- [75] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 9421–9431, 2021. 3
- [76] Jiarui Xu and Xiaolong Wang. Rethinking self-supervised correspondence learning: A video frame-level similarity perspective. In *Proc. Int. Conf. on Computer Vision (ICCV)*, pages 10075–10085, 2021. 2, 6
- [77] 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. 3
- [78] Gengshan Yang, Deqing Sun, Varun Jampani, Daniel Vlasic, Forrester Cole, Ce Liu, and Deva Ramanan. Viser: Video-specific surface embeddings for articulated 3d shape recon-



- struction. *Advances in Neural Information Processing Systems*, 34:19326–19338, 2021. [3](#)
- [79] Vickie Ye, Zhengqi Li, Richard Tucker, Angjoo Kanazawa, and Noah Snavely. Deformable sprites for unsupervised video decomposition. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 2657–2666, 2022. [3](#), [6](#), [7](#), [8](#)
- [80] Wang Zhao, Shaohui Liu, Hengkai Guo, Wenping Wang, and Yong-Jin Liu. ParticleSfM: Exploiting dense point trajectories for localizing moving cameras in the wild. In *Proc. European Conf. on Computer Vision (ECCV)*, pages 523–542, 2022. [2](#)
- [81] Tinghui Zhou, Philipp Krahenbuhl, Mathieu Aubry, Qixing Huang, and Alexei A Efros. Learning dense correspondence via 3d-guided cycle consistency. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 117–126, 2016. [2](#)