

Self-Supervised Any-Point Tracking by Contrastive Random Walks

Ayush Shrivastava and Andrew Owens

University of Michigan

{ayshrv,ahowens}@umich.edu

Abstract. We present a simple, self-supervised approach to the Tracking Any Point (TAP) problem. We train a global matching transformer to find cycle consistent tracks through video via contrastive random walks, using the transformer’s attention-based global matching to define the transition matrices for a random walk on a space-time graph. The ability to perform “all pairs” comparisons between points allows the model to obtain high spatial precision and to obtain a strong contrastive learning signal, while avoiding many of the complexities of recent approaches (such as coarse-to-fine matching). To do this, we propose a number of design decisions that allow global matching architectures to be trained through self-supervision using cycle consistency. For example, we identify that transformer-based methods are sensitive to shortcut solutions, and propose a data augmentation scheme to address them. Our method achieves strong performance on the TapVid benchmarks, outperforming previous self-supervised tracking methods, such as DIFT, and is competitive with several supervised methods.

1 Introduction

The problem of finding space-time correspondences underlies a number of computer vision tasks. An emerging line of work on the Tracking Any Point (TAP) problem [7] has addressed the specific challenges of tracking over long time horizons: estimating all future and past positions of any given physical point in a video. This problem addresses the shortcomings of traditional formulations of long-range tracking, such as chained optical flow and sparse tracking, enabling applications in animation [7] and robotics [43]. Yet the difficulty of acquiring labeled training data has restricted the capabilities of these models. Existing models are thus limited to training on small, synthetic datasets.

This is in contrast to many other areas of computer vision [6, 16, 17], where self-supervised methods have arisen as a powerful way to learn from unlabeled data. While a number of such methods have been proposed for space-time correspondence [2, 19, 44, 46, 52], these methods are not well suited to the challenge of tracking physical points over long time horizons, the core challenge of Tracking Any Point (Fig. 1). Self-supervised optical flow methods [2, 22, 53], for example, obtain dense short-range motion fields but struggle to track over long time horizons, while methods that excel at semantic tracking, such as recent methods




Fig. 1: Global Matching Random Walks. We present a self-supervised method for tracking all physical points over the course of a video, i.e., the Tracking Any Point problem [7]. Our model uses a global matching transformer [49] to track points cycle consistently over time, using the contrastive random walk [19]. Our approach outperforms self-supervised tracking methods, such as self-supervised DIFT [40] and supervised optical flow methods, like RAFT [42], on the TAP-Vid benchmark [7].

that repurpose text-to-image diffusion features [40], match together points that belong to the same object category, not necessarily those that correspond to the same physical points.

In this paper, we propose a simple and effective self-supervised approach to the Tracking Any Point problem. We adapt the global matching transformer architecture [49] to learn through cycle consistency [19, 46, 56]: i.e., tracking forward in time, then backward, should take us back to where we started. In lieu of labeled data, we supervise the model via the contrastive random walk [19], using the self-attention from global matching to define the transition matrix for a random walk that moves between points in adjacent frames. This “all pairs” matching mechanism allows us to define transition matrices that consider large numbers of points at once, thereby increasing spatial precision and enabling us to obtain a richer learning signal by considering a large number of paths through the space-time graph on which the random walk is performed. Additionally, we identify that global matching architectures are susceptible to shortcut solutions (e.g., due to their use of positional encodings), and that previously proposed methods for addressing these shortcuts are insufficient [41]. We therefore propose a type of data augmentation that removes these shortcuts.

Our approach obtains strong performance on the TAP-Vid [7] benchmark, significantly outperforming previous self-supervised tracking methods on TAPVid-DAVIS and Kubric. Through experiments, we show:

- Self-supervised models can obtain strong performance on the Tracking Any Point task, e.g., obtaining competitive performance to TAP-Net [7] on many metrics.

- The contrastive random walk can successfully be extended to long-range point tracking.
- Global matching transformers can be trained through cycle consistency.
- Data augmentation can remove shortcut solutions to cycle consistent training.

2 Related Work

Space-time representation learning. A variety of recent methods have been proposed for learning to track pixels through video via self-supervision. Vondrick et al. [44] learned a representation in which pixels were photoconsistent, providing the model with only grayscale images during training and using the held-out colors to assess the quality of the match. While this approach is effective, it implicitly assumes that objects are photo-consistent over time, an assumption that is frequently violated [22]. Another line of work proposes to use cycle consistency. Wang et al. [46] used a model based on spatial transformers to track pixels forward in time, then backwards, learning a representation that minimized the distance from their point of origin. Other work has combined these two approaches together [27] or use two-stage matching [25]. Jabri et al. [19] formulated video cycle consistency as a random walk on a graph containing space-time patches, providing dense supervision to the model. We extend this approach and use a transformer architecture that allows for global matching and finer correspondences. Bian et al. [2] showed that the contrastive random walk can obtain pixel-accurate matches through multi-scale matching, and proposed several extensions that unify it with self-supervised optical flow methods [22]. In addition to focusing on obtaining spatially precise matches, our goal is to predict tracks in long videos. However, we note that multi-scale matching could be combined with our method to obtain more accurate results. Tang et al. [41] proposed an extension that allowed for fully convolutional training, avoiding shortcut solutions. They proposed a data augmentation technique to use different image crops for the forward and backward cycles of the random walk. However, this does not avoid shortcut solutions for transformer-based methods. Recent work has learned features that change slowly over time [10, 52] or by adding temporal alignment [13], and other work has proposed standardized tracking benchmarks [32]. Another recent work [12] extends the Masked Autoencoders [16] to videos. They mask large fractions of patches in the video and learn visual representations by reconstructing the missing patches.

Optical flow. A parallel line of work has focused on creating unsupervised models for optical flow [22, 29, 35, 37, 47, 54]. Typically these models combine simple photometric losses (e.g., using hand-crafted features) with heavy augmentation and physical constraints, such as smoothness. RAFT [42] proposed a recurrent architecture that updates the flow field through iterative all-pairs matching and regression. Different from prior flow approaches, GMFlow [49] formulated optical flow prediction as a global matching problem that identifies correspondences by directly comparing feature similarities. Extending GMFlow, Xu et al. [51] proposed a unified model for flow, rectified stereo matching and unrectified stereo

depth estimation from posed images. Rather than predicting two-frame velocity estimates which most optical flow methods do, we learn probabilistic matches between frames. To do this, we adopt the transformer-based global matching network architecture of Xu et al. [49], which performs non-parametric matching. However, instead of supervising the model to generate flow estimates, we train it to perform a contrastive random walk.

Cycle consistency. Zhou et al. [56] proposed to use cycle consistency as a supervisory signal across different instances of the same category to train correspondence models. Dwibedi et al [9] used temporal cycle consistency between multiple varying videos to learn representations useful for fine-grained temporal understanding in videos. Other methods have used cycle-consistency to detect occlusions [1, 18, 26, 38], such as within unsupervised flow models [20, 22, 47, 57]. Other work uses cycle consistency for semi-supervised learning [14]. [55] represents the video as two sub-graphs, one which connects inter-frame nodes (similar to Jabri et al. [19]) and another which connects intra-frame nodes located in local neighborhood and use cycle-consistency to perform random walk on this graph.

Tracking. Recently, many works have focused on supervised methods for long-term pixel tracking along with the introduction of new benchmarks. TAP-Vid [7] proposed a new test-bed for tracking any point in a video. They released four benchmarks based on real and synthetic videos for evaluating tracking methods. They also provided a tracking model called TAP-Net that compares the features of a query point with all frames to predict tracks independent of time. Concurrently, Persistent Independent Particles (PIPs) [15] introduced a tracking method that searches over a local neighborhood and smooths the estimates over time by iterative refinement. They also release a dataset called FlyingThings++ based on FlyingThings [31]. Neoral et al [33] proposed a supervised tracker that exploits optical flows from consecutive frames and pairs of frames at logarithmically scaled intervals. TAPIR [8] uses a two-stage approach: matching stage that provides an initial location of the query point in every frame and a refinement stage that updates the tracks based on local neighborhoods, essentially combining techniques from TAP-Net and PIPs models. CoTracker [23] introduced a transformer architecture that jointly tracks multiple points throughout an entire video. OmniMotion [45] introduces a test-time optimization method by building globally consistent motion representations. All these methods use ground-truth tracks or trajectories computed from pretrained optical flow methods for training, whereas we use self-supervision to train our model provided by cycle-consistency as the supervisory signal.

3 Method: Global Matching Random Walk

We propose a self-supervised method for tracking points in video. We use the contrastive random walk [19] to learn cycle-consistent track, using an architecture based on self-attention from a global matching transformer. Our use of a transformer-based architecture, which requires addressing additional shortcut solutions that were not present in previous work on cycle consistent tracking, which




Fig. 2: Model Architecture. Our model takes a pair of images I_t and I_{t+1} as input over which it computes correspondences. We extract visual features from a CNN, add positional encodings, and pass them as tokens to our global matching transformer. The transformer consisting of 6 stacked layers of self-attention, cross-attention and feed-forward networks, processes these features and produces correlated features F_t and F_{t+1} . We compute self-attention over F_t and F_{t+1} and use the attention as the transition matrix for performing contrastive random walks. To compute tracks during evaluation, we can take an expectation over the affinity matrix to get coordinates (x, y) .

we address through an augmentation scheme. We train on unlabeled videos and perform point tracking in a strided fashion to get pixel-level features for matching.

Problem setup. In the Tracking Any Point (TAP) problem, we are given a video V with frames $\{I_t\}_{t=1}^T$ where $I_t \in \mathbb{R}^{H \times W \times 3}$ and query points $\{\mathbf{q}_k\}_{k=1}^N$ where each $\mathbf{q}_k = (t_k, x_k, y_k)$ represents a timestep in the video and its position in the frame I_{t_k} . The goal is to produce point trajectories, $\mathbf{p}_t^k = (x_t^k, y_t^k)$ for all $t \in T$ for each query point \mathbf{q}_k . In addition to positions, the visibility $v_t^k \in \{0, 1\}$ indicate whether the point \mathbf{p}_t^k is occluded in frame I_t . Note that the query point can be provided for any timestep in the video, not necessarily the first frame, and that at that timestep the query point is assumed to be visible.

Following recent supervised approaches [7], we create a model that operates on a pair of (not necessarily temporally adjacent) video frames. At test time, this model can be used to track points over long time horizons, either by directly matching pairs of frames or by chaining.

3.1 Global matching architecture

In order to obtain spatially precise correspondences, we require an architecture (Fig. 2) that can perform efficient, global matching. We adapt the recent GMFlow architecture of Xu et al. [49, 50] for self-supervised tracking. This architecture has previously been applied to optical flow and stereo-matching in supervised settings.

A key advantage of this design is that it finds a match by “all pairs” matching through self-attention. This is in contrast to alternative architectures, which rely

on regressing motion [8, 15, 39, 42] and thus provide only a single point estimate. We use global matching to define a transition matrix for the contrastive random walk, which allows us to model the motion of a very large number of points at once. This has several advantages to previous approaches to the contrastive random walk [2, 19]. First, the motion can be captured at a much finer-grained level. For example, in our experiments we divide an 256×256 resolution image into a 64×64 grid, whereas Jabri et al. [19] was limited to a coarse 7×7 grid and Bian et al. [2] required coarse-to-fine matching. Second, it allows us to obtain more supervision per iteration through contrastive learning, since the model can simultaneously explore a larger number of paths through the space-time graph on which we will perform the random walk.

Image features. We extract d -dimensional features $\phi(I_t) \in \mathbb{R}^{\frac{H}{c} \times \frac{W}{c} \times d}$ for each image I_t from a convolutional neural network where $c = 4$. We add 2D positional encoding to these visual features. While the visual encoder used in Xu et al. [49] computes features at two scales, we only perform single-scale matching for simplicity.

Global correlation matching. For each pair of consecutive video frames, I_t and I_{t+1} , we process the image features through six layers of stacked self-, cross-attention and feed-forward networks to get correlation features, F_t, F_{t+1} . The keys and values for cross-attention layers come from the same feature, but queries come from the other feature in the pair. We also use Swin-style shifted local windows to improve computation efficiency [30, 49].

Computing a random walk transition matrix. Once we have correlation features F_t, F_{t+1} , we compute the transition matrix as $A_t^{t+1} = \text{softmax}(F_t F_{t+1}^\top / \tau)$. This matrix represents the probability of a patch in frame t matching to frame $t + 1$, and is used to define the transition probability for a contrastive random walk. In contrast to previous contrastive random walk models, which (following work in contrastive learning [4, 17, 48]) L_2 -normalize the embedding features and apply a small temperature parameter [19], we follow the standard practice in transformers and use unnormalized features with a large normalization constant $\tau = \sqrt{d}$. We provide full architectural details in the supplementary material.

Sampling stride. Our model uses probabilities from the affinity matrix to find correspondences. The resolution of the affinity matrix is determined by the spatial resolution of the correlation features F_t, F_{t+1} . To compute pixel-level features, we use different feature strides to sample features from the images following previous works [23, 33]. We implement this by upsampling the original image by the stride. We vary our sampling strides among $s = \{1, 2, 4\}$ and mention the stride used during training and evaluation in our experiments.

Estimating the motion and visibility. Given a transition matrix $A_{s,t}$ between frames s and t , we can compute the expected change in position, following [2, 49]. In other words, we take a weighted sum of the positions of the points we match in frame t :

$$\mathbf{f}_{s,t} = \mathbb{E}_{A_{s,t}}[A_{s,t}D - D], \quad (1)$$




Fig. 3: Label Warping. We propose label warping as a remedy to avoid shortcut solutions that arise when we use transformer-based models for contrastive random walks. Instead of warping the last feature to match the first feature, we propose to warp the label used for cycle consistency. For an image pair I_t, I_{t+1} , we apply different affine transformations T^f, T^b to the forward and backward cycle. We then compute A_t^{t+1}, A_{t+1}^t and chain them together to get the affinity matrix for cycle consistency. We then supervise it with the warped identity matrix $T_f^b(I)$ where T_f^b represents the transformation to go from T^f to T^b .

where $\mathbf{f}_{s,t} \in \mathbb{R}^{n \times 2}$ is the matrix of predicted optical flows, $D \in \mathbb{R}^{n \times 2}$ is the (constant) matrix containing of pixel coordinates for each point, and $A_{s,t}D$ is the expected position from frame s to t .

To estimate the visibility v_i for a given point, we perform a cycle consistency test: we check whether the predicted motion fields $\mathbf{f}_{s,t}$ and \mathbf{t}, \mathbf{s} produce a motion that is within a threshold τ_{cyc} (we use $\tau_{cyc} = 3$ pixels) of the original point. Those that exceed this threshold are marked as being invisible.

3.2 Learning the model

Cycle-consistency by contrastive random walks. We use the cycle consistency objective introduced contrastive random walks in Jabri et al. [19] to supervise our model. We treat the video as a space-time graph and train a model to perform random walks. We use our model to compute the transition matrices for the walker, which walks from patches in frame t to $t + 1$, then back to t . This sequence is known as a *palindrome* sequence. We supervise the model by maximizing the probability that the walker returns to its initial location, which amounts to minimizing the loss:

$$\mathcal{L}_{\text{crw}} = \mathcal{L}_{\text{CE}}(A_t^{t+1} A_{t+1}^t, I), \quad (2)$$

where $A_t^{t+1} A_{t+1}^t$ is the (chained) transition matrix for the random walk, and \mathcal{L}_{CE} is cross-entropy loss, and I is the identity matrix. In other words, the model penalizes random walks whose transition matrices differ from the identity, and which thus are not cycle-consistent.

Shortcut solutions for global matching models. Many cycle consistency learning schemes are susceptible to a trivial solution [2, 19]: ignore the visual content of the patch, and match solely based on its position. Since the transformer model has a positional encoding, and our architecture has a large number of global self-attention layers, this solution is easy to find. Tang et al. [41] showed that this shortcut could be removed in simple fully convolutional CNN models by randomly cropping and resizing the video frames, using different (but consistent) augmentations in the forward and backward directions of the walk. When computing the loss (Eq. 2), they undo the crop by warping the last feature in the affinity matrix with the inverse transformation, and supervise the affinity matrix to be the identity I . We found this method still led to shortcut solutions. In our case, the CNN features are processed by a transformer, which is capable of undoing the warp and trivially matching based on the positional information. To address this issue, we perform *label warping* (Fig. 3). Instead of warping the network features, we warp the labels to match the transformation between the forward and backward cycle. Let T^f and T^b be different resize-crop transformations applied to the forward and backward cycles of the contrastive random walk. Specifically, for the image sequence (I_1, I_2, I_1) , we construct the frame sequence $[T^f(I_1), T^f(I_2), T^b(I_1)]$ and compute the loss:

$$\mathcal{L}_{\text{crw}} = \mathcal{L}_{\text{CE}}(A_s, T_f^b(I)), \quad (3)$$

where A_s is the chained transition matrices for the full random walk, and $T_f^b(I)$ denotes that we apply the same spatial transformation to the label set.

Smoothness loss. We also evaluate model variations that impose spatial smoothness. To ensure the movement of random walkers is smooth, we also consider variations of the model that follow Bian et al. [2] and add an edge-aware smoothness loss [22]:

$$\mathcal{L}_{\text{smooth}} = \mathbb{E}_p \sum_{d \in \{x, y\}} \exp(-\lambda_c I_d(p)) \left| \frac{\partial^2 \mathbf{f}_{s,t}(p)}{\partial d^2} \right| \quad (4)$$

where the p is a pixel, $I_d(p) = \frac{1}{3} \sum_c \left| \frac{\partial I_c}{\partial d} \right|$ is the spatial derivative averaged over all color channels I_c in direction d , and where we use the second derivatives of the estimated flow field. The parameter λ_c controls the influence of pixels with similar colors. When both losses are used together, the model minimizes:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{crw}} + \lambda_s \mathcal{L}_{\text{smooth}}, \quad (5)$$

where λ_s is a constant.

4 Experiments

Our model estimates space-time correspondences between a pair of video frames, from which we can compute expected coordinates for each track. We evaluate our method on the four Tap-Vid benchmarks and use the standard evaluation metrics [7].

4.1 Training

Following Doersch et al. [7], we use the Kubric dataset [11] as our main dataset for training. The TapVid-Kubric dataset consists of synthetic images of a few objects randomly moving around in the video. It contains a training set of 38,325 videos of 256×256 resolution and 799 validation videos. During training, we randomly sample 2 frames separated by a few timesteps from the video and create a palindrome sequence ($I_1 \rightarrow I_2 \rightarrow I_1'$). We then apply different random resized crops to forward (I_1, I_2) and backward images (I_1') and train for cycle consistency and smoothness loss. Unlike other models that address TAP, we do not use any labels from the training data, as our method is entirely self-supervised. To evaluate the ability of our model to generalize to “in the wild” internet video, we also train a version of the model using the Kinetics 400 [24] dataset using a similar 2-frame sampling strategy.

4.2 Evaluation

To validate our approach and to test the generalization of our trained models, we run evaluations on the TapVid benchmarks [7], namely on Kubric, DAVIS, Kinetics, and RGB-Stacking. TapVid-DAVIS is a real dataset of 30 videos from DAVIS 2017 validation [34] with videos from 34-104 frames. Similarly, TapVid-Kinetics is a real dataset of 1000+ videos from Kinetics [24] with 250 frames each. TapVid-RGB-Stacking is a dataset of 50 synthetic videos with 250 frames each from robotic hand object manipulation task. We mainly use TapVid-DAVIS for our testbed as it is based on real videos.

Metrics. In TapVid benchmarks, the query points are sampled using a special strategy (when using *strided* query method) [7] and then are tracked forward and backward in time. Performance is evaluated on: (1) the positional accuracy $< \delta_{\text{avg}}^x$ metric for frames in which the point is visible, the fraction of points that are within a threshold over the ground truth, averaged over several thresholds. (2) *Occlusion Accuracy* which is a classification accuracy for predicting where a point is visible or occluded. (3) *Average Jaccard* (AJ), the fraction of true positives (points within a threshold of visible ground truth points), divided by true positives and false positives, averaged over multiple thresholds [7].

5 Results

We compare our method to previous supervised and self-supervised approaches, and evaluate a number of different design decisions.

5.1 Comparison to other methods

In Table 1, we compare our method to previously proposed self-supervised tracking and supervised approaches.

Method	Multi-Frame	Kubric			DAVIS			Kinetics			RGB-Stacking			
		AJ \uparrow	$< \delta_{\text{avg}}^x$ \uparrow	OA \uparrow	AJ \uparrow	$< \delta_{\text{avg}}^x$ \uparrow	OA \uparrow	AJ \uparrow	$< \delta_{\text{avg}}^x$ \uparrow	OA \uparrow	AJ \uparrow	$< \delta_{\text{avg}}^x$ \uparrow	OA \uparrow	
Supervised	RAFT-C [42]	41.2	58.2	86.4	30.7	46.6	80.2	31.7	51.7	84.3	42.0	56.4	91.5	
	Kubric-VFS-Like [11]	51.9	69.8	84.6	33.1	48.5	79.4	40.5	59.0	80.0	57.9	72.6	91.9	
	RAFT-D [42]	61.8	79.1	87.9	34.1	48.9	76.1	72.1	85.1	92.1	50.6	66.9	85.5	
	COTR [21]	40.1	60.7	78.6	35.4	51.3	80.2	19.0	38.8	57.4	6.8	13.5	79.1	
	TAP-Net [7]	65.4	77.7	93.0	38.4	53.1	82.3	46.6	60.9	85.0	59.9	72.8	90.4	
	PIPs [15]	✓	59.1	74.8	88.6	42.0	59.4	82.1	35.3	54.8	77.4	37.3	51.0	91.6
	TAPIR [8]	✓	84.7	92.1	95.8	61.3	73.6	88.8	57.2	70.1	87.8	62.7	74.6	91.6
	CoTracker [23]	✓	—	—	—	64.8	79.1	88.7	—	—	—	—	—	—
Self-supervised	CRW-C [19]	31.4	48.1	76.3	7.7	13.5	72.9	20.2	33.6	70.6	25.3	35.2	70.1	
	CRW-D [19]	35.8	52.4	80.9	23.6	38.0	77.2	21.9	36.8	70.4	13.1	23.0	83.4	
	DIFT-C [40]	28.3	45.2	69.0	18.1	33.0	68.8	19.8	33.7	68.7	13.2	21.9	56.3	
	DIFT-D [40]	41.6	59.8	83.9	29.7	48.2	77.2	19.5	34.4	70.1	24.4	38.9	89.9	
	Flow-Walk-C [2]	49.4	66.7	82.7	35.2	51.4	80.6	40.9	55.5	84.5	41.3	55.7	92.2	
	Flow-Walk-D [2]	51.1	68.1	80.3	24.4	40.9	76.5	46.9	65.9	81.8	66.3	82.7	91.2	
	ARFlow-C [28]	52.3	68.1	81.4	35.0	51.8	79.7	27.3	44.3	79.5	33.0	47.2	91.9	
	Ours - GMRW-C	54.2	72.4	82.6	41.8	60.9	78.3	31.9	52.3	72.9	39.8	56.5	90.8	
	Ours - GMRW-D	51.4	71.7	83.9	30.3	49.4	77.3	36.3	59.2	71.0	56.4	74.1	90.9	

Table 1: Comparison of our method and baselines on the Tap-Vid benchmark. We show strong performance on all four Tap-Vid benchmarks. We outperform self-supervised methods on Kubric and DAVIS and are comparable with several supervised methods.

Baselines. We compare our method with several self-supervised baselines. Following Wang et al. [45], we run all self-supervised methods in Chained (-C) and Direct (-D) settings and show their performance in Table 1 (bottom section). In the chained setting, the predictions are made for adjacent frames and they are chained together over time to form long-range tracks. In the direct setting, the query point frame is compared directly with all frames and the motion is computed for each pair independent of time. For all self-supervised methods (none of which explicitly predict occlusion masks), we test for occlusion using cycle consistency over the predicted tracks using the same approach used in our method. We evaluate the following methods:

- **Diffusion Features (DIFT):** Next, we consider recent work that uses diffusion features [40] for tracking as a baseline. This model extracts features from pretrained Stable Diffusion model [36] and Ablated Diffusion [5] model, and successfully uses them for a variety of correspondence tasks. For real images, they add a noise of a specific time step t , feed it to the network together with t to extract intermediate layer activations. These activations are then used as features. We use the Ablated Diffusion variation of the model, since it performs better for temporal correspondence tasks.
- **Contrastive Random Walk (CRW):** We first evaluate the CRW model of Jabri et al. [19] on the TapVid benchmarks. In CRW, the images are partitioned into small patches, and the model is trained to learn representations

Method	Kubric			DAVIS		
	AJ \uparrow	$< \delta_{\text{avg}}^x \uparrow$	OA \uparrow	AJ \uparrow	$< \delta_{\text{avg}}^x \uparrow$	OA \uparrow
Supervised	63.7	83.2	83.9	39.1	59.6	77.3
CRW	25.4	39.3	83.3	10.4	19.2	76.9
+ Label Warping	45.3	62.2	83.1	32.1	48.9	78.2
+ Smoothness loss	49.0	66.7	84.4	33.0	50.5	79.4
+ Train stride $s = 2$	47.7	65.6	84.4	34.5	52.1	79.3
Trained w/ Kinetics	47.5	65.0	83.8	34.6	52.6	78.7
Eval stride $s = 4$	37.8	53.8	84.1	23.5	38.4	78.8
Eval stride $s = 2$	47.5	65.0	83.8	34.6	52.6	78.7
Eval stride $s = 1$	54.2	72.4	82.6	41.8	60.9	78.3

Table 2: Model variations. We evaluate several variations of our model. We consider a version of the architecture trained via supervised learning, ablated versions of the model, a version trained on “in-the-wild” Kinetics [24] videos, and several different stride configurations during evaluation.

for these patches through cycle consistency. The model is trained for cycle consistent videos of length 10 and sub-cycles are also supervised. Since the CRW operates at the patch level, the correspondences computed are very coarse.

- **FlowWalk:** We also evaluate FlowWalk [2] on these tracking benchmarks, using the reported results from OmniMotion [45]. FlowWalk is another contrastive random walk-based method that is trained for cycle consistency at multiple scales, and which can perform optical flow estimation. The output from their model is pixel-level flow values.
- **Supervised methods:** We also show several supervised methods on the TapVid benchmark in the top section of Table 1. Out of these, PIPs, TAPIR, and CoTracker are trained with multiple frames and use local, spatial-temporal information to refine their tracks.

Quantitative comparisons. We compare our method qualitatively to self-supervised baselines in Table 1. We find that it outperformed other self-supervised baselines on TapVid-Kubric and TapVid-DAVIS and have competitive numbers with several supervised methods, such as TAP-Net. We evaluate our method at a stride $s = 1$ for chained setting and $s = 2$ for direct setting, except for the TapVid-RGB-Stacking benchmark, where we find the best performance at $s = 4$ (perhaps because the highly synthetic nature of the videos leads to local ambiguity). Among other self-supervised methods, CRW operates at the patch level, so the correspondence we get is very coarse and does not work well when evaluated for tracking. DIFT outperforms CRW, but fails to outperform other self-supervised methods, possibly because DIFT’s text-to-image diffusion features rely on the semantics of the image to find correspondence, but tracking requires relying on low-level motion cues to find correspondences.



Fig. 4: Qualitative results. We show qualitative results for TapVid-DAVIS videos and compare them with DIFT and RAFT. DIFT relies on semantic correspondences and often loses the point of interest when motion occurs in the video. RAFT produces accurate movements for several tracks but suffers from drifting of points when the predictions are chained over a long period. In the first video, our method can track points accurately over the long timesteps. RAFT, on the other hand, loses current locations for 2 query points and latches on points on the ground and starts tracking them. In the other 2 videos as well, our method works better than RAFT and DIFT. DIFT produces inaccurate tracks that do not capture motion well. RAFT being accurate most of the time, loses track of points close to the boundary.




Fig. 5: Optical flow visualization. Although our method is not trained for the optical flow prediction task, it is able to produce reasonable flow outputs over multiple timesteps. RAFT produces high quality flows as it is an optical flow method trained for this objective. DIFT predicts inaccurate flow which are spotty in nature, suggesting that it relies on finding semantic correspondence for certain points in the image, instead of relying local motion cues.

Qualitative results. We show qualitative results on tracking in Figure 4 and optical flow in Figure 5 and compare them with DIFT and RAFT. Our method is able to produce reasonable tracks for long time horizon and works better than DIFT and RAFT. Even though our method is not supervised for the optical flow task, it is still able to produce reasonable-quality flow maps.

5.2 Model variations and ablations

We investigate our model’s components in Table 2. First, we distinguish between the performance of our architecture versus our learning procedure. To test this, we train our global matching-based architecture (Sec. 3.1) by training a *supervised* variant of our model (evaluated at stride $s = 4$). Our training setup closely follows TAP-Net [7]. We randomly sample 2 frames from a TapVid-Kubric training video and supervise the estimated motion (Eq. 1) using the ground truth, using a scaled Huber loss. We compare with TAP-Net as it is modeled with a

similar setup as ours without multi-scale features and spatial-temporal iterative refinement. We see that the performance of our supervised model is better than TAP-Net on TAPVid-Kubric and TAPVid-DAVIS. These results suggest that the architecture is capable of obtaining tracking results that are on par with other supervised architectures that have been proposed for TAP.

Next, we evaluate the different components of our approach. We train our self-supervised model on the Kubric dataset without label warping, finding that it performs poorly. This suggests that it can find shortcuts using positional embeddings without learning meaningful representations. After adding label warping, we see a large boost in performance. Next, we train our model with smoothness loss in addition to cycle consistency and see a small improvement in performance. Interestingly, this is in contrast to Bian et al. [2], which found the smoothness loss to be critically important. We hypothesize that this is due to our use of global matching, rather than the coarse-to-fine search used in [2]. The latter may implicitly require neighboring pixels to match to similar locations, since the finer scales are obtained by warping a feature map [3] using the estimated optical flow of each pixel’s neighbors. This highlights a potential advantage of our network architecture.

Next, to test the effect of stride in training, we lower the training stride from $s = 4$ to $s = 2$ and see that we obtain a minor performance improvement. Finally, we trained our model with all components on the Kinetics 400 dataset and see that performance improves slightly on TAPVid-DAVIS, while decreasing for the TAPVid-Kubric dataset. This suggests, first, that our model can successfully be trained using unlabeled in-the-wild video, rather than the synthetic datasets used in existing supervised learning work. Second, since the improvement was for a dataset of real videos (DAVIS) and decreased performance was for a synthetic dataset (Kubric), this may suggest that it is beneficial to match the distribution by training on real video.

6 Conclusion

In this paper, we present a simple, self-supervised model for addressing long-range tracking for the Tracking Any Point task. We adapt the global matching transformer architecture to training via the contrastive random walk. Our approach significantly outperforms previous self-supervised approaches on the TAP-Vid benchmark, obtaining performance that on some metrics is close to that of supervised TAP-Net [7]. While our model can handle occlusions to some degree using cycle-consistency, we did not design our model to handle occlusion explicitly. It also does not have the capability to track pixels through occlusions. We see this work as a step towards building tracking methods trained using self-supervised learning, and to creating computer vision models that learn at scale from unlabeled data.

Acknowledgements. This work was supported by Toyota Research Institute and Cisco Systems. We would like to thank Adam Harley, Allan Jabri, Daniel Geng and Ziyang Chen for helpful discussions and feedback on the paper.

References

1. Baker, S., Scharstein, D., Lewis, J., Roth, S., Black, M.J., Szeliski, R.: A database and evaluation methodology for optical flow. *International journal of computer vision* **92**, 1–31 (2011) [4](#)
2. Bian, Z., Jabri, A., Efros, A.A., Owens, A.: Learning pixel trajectories with multi-scale contrastive random walks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6508–6519 (2022) [1](#), [3](#), [6](#), [8](#), [10](#), [11](#), [14](#)
3. Brox, T., Bruhn, A., Papenberg, N., Weickert, J.: High accuracy optical flow estimation based on a theory for warping. In: *Computer Vision-ECCV 2004: 8th European Conference on Computer Vision, Prague, Czech Republic, May 11-14, 2004. Proceedings, Part IV* 8. pp. 25–36. Springer (2004) [14](#)
4. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International conference on machine learning*. pp. 1597–1607. PMLR (2020) [6](#)
5. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021) [10](#)
6. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1422–1430 (2015) [1](#)
7. Doersch, C., Gupta, A., Markeeva, L., Recasens, A., Smaira, L., Aytar, Y., Carreira, J., Zisserman, A., Yang, Y.: Tap-vid: A benchmark for tracking any point in a video. *arXiv preprint arXiv:2211.03726* (2022) [1](#), [2](#), [4](#), [5](#), [8](#), [9](#), [10](#), [13](#), [14](#)
8. Doersch, C., Yang, Y., Vecerik, M., Gokay, D., Gupta, A., Aytar, Y., Carreira, J., Zisserman, A.: Tapir: Tracking any point with per-frame initialization and temporal refinement. *arXiv preprint arXiv:2306.08637* (2023) [4](#), [6](#), [10](#)
9. Dwibedi, D., Aytar, Y., Tompson, J., Sermanet, P., Zisserman, A.: Temporal cycle-consistency learning. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 1801–1810 (2019) [4](#)
10. Gordon, D., Ehsani, K., Fox, D., Farhadi, A.: Watching the world go by: Representation learning from unlabeled videos. *arXiv preprint arXiv:2003.07990* (2020) [3](#)
11. Greff, K., Belletti, F., Beyer, L., Doersch, C., Du, Y., Duckworth, D., Fleet, D.J., Gnanapragasam, D., Golemo, F., Herrmann, C., et al.: Kubric: A scalable dataset generator. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3749–3761 (2022) [9](#), [10](#)
12. Gupta, A., Wu, J., Deng, J., Li, F.F.: Siamese masked autoencoders. *Advances in Neural Information Processing Systems* **36** (2024) [3](#)
13. Hadji, I., Derpanis, K.G., Jepson, A.D.: Representation learning via global temporal alignment and cycle-consistency. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11068–11077 (2021) [3](#)
14. Haeusser, P., Mordvintsev, A., Cremers, D.: Learning by association—a versatile semi-supervised training method for neural networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 89–98 (2017) [4](#)
15. Harley, A.W., Fang, Z., Fragkiadaki, K.: Particle video revisited: Tracking through occlusions using point trajectories. In: *European Conference on Computer Vision*. pp. 59–75. Springer (2022) [4](#), [6](#), [10](#)
16. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 16000–16009 (2022) [1](#), [3](#)

17. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020) [1](#), [6](#)
18. Hur, J., Roth, S.: Mirrorflow: Exploiting symmetries in joint optical flow and occlusion estimation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 312–321 (2017) [4](#)
19. Jabri, A., Owens, A., Efros, A.: Space-time correspondence as a contrastive random walk. *Advances in neural information processing systems* **33**, 19545–19560 (2020) [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#), [10](#)
20. Janai, J., Guney, F., Ranjan, A., Black, M., Geiger, A.: Unsupervised learning of multi-frame optical flow with occlusions. In: Proceedings of the European conference on computer vision (ECCV). pp. 690–706 (2018) [4](#)
21. Jiang, W., Trulls, E., Hosang, J., Tagliasacchi, A., Yi, K.M.: Cotr: Correspondence transformer for matching across images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6207–6217 (2021) [10](#)
22. Jonschkowski, R., Stone, A., Barron, J.T., Gordon, A., Konolige, K., Angelova, A.: What matters in unsupervised optical flow. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16. pp. 557–572. Springer (2020) [1](#), [3](#), [4](#), [8](#)
23. Karaev, N., Rocco, I., Graham, B., Neverova, N., Vedaldi, A., Rupprecht, C.: Cotracker: It is better to track together. arXiv preprint arXiv:2307.07635 (2023) [4](#), [6](#), [10](#)
24. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017) [9](#), [11](#)
25. Lai, Z., Lu, E., Xie, W.: Mast: A memory-augmented self-supervised tracker. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6479–6488 (2020) [3](#)
26. Lei, C., Yang, Y.H.: Optical flow estimation on coarse-to-fine region-trees using discrete optimization. In: 2009 IEEE 12th International Conference on Computer Vision. pp. 1562–1569. IEEE (2009) [4](#)
27. Li, X., Liu, S., De Mello, S., Wang, X., Kautz, J., Yang, M.H.: Joint-task self-supervised learning for temporal correspondence. *Advances in Neural Information Processing Systems* **32** (2019) [3](#)
28. Liu, L., Zhang, J., He, R., Liu, Y., Wang, Y., Tai, Y., Luo, D., Wang, C., Li, J., Huang, F.: Learning by analogy: Reliable supervision from transformations for unsupervised optical flow estimation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [10](#)
29. Liu, P., King, I., Lyu, M.R., Xu, J.: Ddflow: Learning optical flow with unlabeled data distillation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 8770–8777 (2019) [3](#)
30. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021) [6](#)
31. Mayer, N., Ilg, E., Haussler, P., Fischer, P., Cremers, D., Dosovitskiy, A., Brox, T.: A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4040–4048 (2016) [4](#)

32. McKee, D., Zhan, Z., Shuai, B., Modolo, D., Tighe, J., Lazebnik, S.: Transfer of representations to video label propagation: Implementation factors matter. arXiv preprint arXiv:2203.05553 (2022) [3](#)
33. Neoral, M., Šerých, J., Matas, J.: Mft: Long-term tracking of every pixel. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 6837–6847 (2024) [4](#), [6](#)
34. Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Van Gool, L.: The 2017 davis challenge on video object segmentation. arXiv preprint arXiv:1704.00675 (2017) [9](#)
35. Ren, Z., Yan, J., Ni, B., Liu, B., Yang, X., Zha, H.: Unsupervised deep learning for optical flow estimation. In: Thirty-First AAAI Conference on Artificial Intelligence (2017) [3](#)
36. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022) [10](#)
37. Stone, A., Maurer, D., Ayvaci, A., Angelova, A., Jonschkowski, R.: Smurf: Self-teaching multi-frame unsupervised raft with full-image warping. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3887–3896 (2021) [3](#)
38. Sun, D., Vlasic, D., Herrmann, C., Jampani, V., Krainin, M., Chang, H., Zabih, R., Freeman, W.T., Liu, C.: Autoflow: Learning a better training set for optical flow. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10093–10102 (2021) [4](#)
39. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8934–8943 (2018) [6](#)
40. Tang, L., Jia, M., Wang, Q., Phoo, C.P., Hariharan, B.: Emergent correspondence from image diffusion. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=yp0iXjdfnU> [2](#), [10](#)
41. Tang, Y., Jiang, Z., Xie, Z., Cao, Y., Zhang, Z., Torr, P.H., Hu, H.: Breaking shortcut: Exploring fully convolutional cycle-consistency for video correspondence learning. arXiv preprint arXiv:2105.05838 (2021) [2](#), [3](#), [8](#)
42. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16. pp. 402–419. Springer (2020) [2](#), [3](#), [6](#), [10](#)
43. Vecerik, M., Doersch, C., Yang, Y., Davchev, T., Aytar, Y., Zhou, G., Hadsell, R., Agapito, L., Scholz, J.: Robotap: Tracking arbitrary points for few-shot visual imitation. arXiv preprint arXiv:2308.15975 (2023) [1](#)
44. Vondrick, C., Shrivastava, A., Fathi, A., Guadarrama, S., Murphy, K.: Tracking emerges by colorizing videos. In: Proceedings of the European conference on computer vision (ECCV). pp. 391–408 (2018) [1](#), [3](#)
45. Wang, Q., Chang, Y.Y., Cai, R., Li, Z., Hariharan, B., Holynski, A., Snavely, N.: Tracking everything everywhere all at once. arXiv preprint arXiv:2306.05422 (2023) [4](#), [10](#), [11](#)
46. Wang, X., Jabri, A., Efros, A.A.: Learning correspondence from the cycle-consistency of time. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2566–2576 (2019) [1](#), [2](#), [3](#)
47. Wang, Y., Yang, Y., Yang, Z., Zhao, L., Wang, P., Xu, W.: Occlusion aware unsupervised learning of optical flow. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4884–4893 (2018) [3](#), [4](#)

48. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3733–3742 (2018) [6](#)
49. Xu, H., Zhang, J., Cai, J., Rezatofghi, H., Tao, D.: Gmflow: Learning optical flow via global matching. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8121–8130 (2022) [2](#), [3](#), [4](#), [5](#), [6](#)
50. Xu, H., Zhang, J., Cai, J., Rezatofghi, H., Yu, F., Tao, D., Geiger, A.: Unifying flow, stereo and depth estimation. arXiv preprint arXiv:2211.05783 (2022) [5](#)
51. Xu, H., Zhang, J., Cai, J., Rezatofghi, H., Yu, F., Tao, D., Geiger, A.: Unifying flow, stereo and depth estimation. IEEE Transactions on Pattern Analysis and Machine Intelligence (2023) [3](#)
52. Xu, J., Wang, X.: Rethinking self-supervised correspondence learning: A video frame-level similarity perspective. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10075–10085 (2021) [1](#), [3](#)
53. Yu, J.J., Harley, A.W., Derpanis, K.G.: Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In: Computer Vision–ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part III 14. pp. 3–10. Springer (2016) [1](#)
54. Yu, J.J., Harley, A.W., Derpanis, K.G.: Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In: Computer Vision - ECCV 2016 Workshops, Part 3 (2016) [3](#)
55. Zhao, Z., Jin, Y., Heng, P.A.: Modelling neighbor relation in joint space-time graph for video correspondence learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9960–9969 (2021) [4](#)
56. Zhou, T., Krahenbuhl, P., Aubry, M., Huang, Q., Efros, A.A.: Learning dense correspondence via 3d-guided cycle consistency. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 117–126 (2016) [2](#), [4](#)
57. Zou, Y., Luo, Z., Huang, J.B.: Df-net: Unsupervised joint learning of depth and flow using cross-task consistency. In: Proceedings of the European conference on computer vision (ECCV). pp. 36–53 (2018) [4](#)