

Real-World Point Tracking with Verifier-Guided Pseudo-Labeling

Görkay Aydemir¹ Fatma Güney^{1,2 †} Weidi Xie^{3 †}

¹Department of Computer Engineering, Koç University ²KUIS AI Center

³School of Artificial Intelligence, Shanghai Jiao Tong University

gorkayaydemir@gmail.com

Abstract

*Models for long-term point tracking are typically trained on large synthetic datasets. The performance of these models degrades in real-world videos due to different characteristics and the absence of dense ground-truth annotations. Self-training on unlabeled videos has been explored as a practical solution, but the quality of pseudo-labels strongly depends on the reliability of teacher models, which vary across frames and scenes. In this paper, we address the problem of real-world fine-tuning and introduce **verifier**, a meta-model that learns to assess the reliability of tracker predictions and guide pseudo-label generation. Given candidate trajectories from multiple pretrained trackers, the verifier evaluates them per frame and selects the most trustworthy predictions, resulting in high-quality pseudo-label trajectories. When applied for fine-tuning, verifier-guided pseudo-labeling substantially improves the quality of supervision and enables data-efficient adaptation to unlabeled videos. Extensive experiments on four real-world benchmarks demonstrate that our approach achieves state-of-the-art results while requiring less data than prior self-training methods. Project page: kuis-ai.github.io/track_on_r.*

1. Introduction

Motion estimation is a longstanding problem in computer vision, aiming to reliably track physical points across video frames. While short-term correspondences can be well handled by optical flow [39, 41], extending this capability over long temporal horizons, often referred to as point tracking [11, 18], has recently drawn increased attention. Point tracking is a fundamental primitive for long-term visual understanding and control, enabling a wide range of applications where motion is crucial, including video editing [16], robot manipulation and perception [5, 43, 45], 4D scene understanding [4, 14, 24, 36, 44], and medical analysis [3, 23].

Recently, transformer-based point trackers [1, 22, 29]

have advanced the state of the art, yet their training regimes remain synthetic due to the prohibitive cost of dense, frame-accurate long-term trajectory annotation in real videos. As a result, models often inherit a sim-to-real discrepancy: appearance statistics, nonrigid motion, occlusion patterns, lighting changes, and sensor artifacts in natural footage degrade reliability over extended sequences. Bridging this gap requires training that can exploit unlabeled, large-scale real-world videos without access to dense ground-truth trajectories.

Self-training via pseudo-labels is an attractive path forward: predictions from a pretrained tracker (or an ensemble) serve as supervision to adapt on real data [21]. However, naïve pseudo-labeling is brittle, as the teacher predictions are not uniformly reliable. Different trackers excel in different regimes, some resist fast motion but drift under low texture; others handle occlusion better but suffer from identity switches or jitter. Fixed heuristics or global confidence thresholds cannot reconcile these heterogeneous error profiles and often propagate systematic errors during adaptation.

We argue that effective real-world training for point tracking hinges on reliability estimation: models must learn when and where to trust tracker outputs. To this end, we introduce a **verifier**, a learned meta-model that scores the framewise reliability of candidate trajectories produced by multiple pretrained, off-the-shelf trackers. Given a query point and several candidate tracks, the verifier predicts which candidate most faithfully follows the underlying motion at each frame, enabling dynamic selection and seamless switching as conditions change (Fig. 1, left). Trained entirely on synthetic data with ground-truth trajectories, the verifier learns from deliberately perturbed candidates that emulate realistic errors (drift, jumps, occlusions, re-appearances), using a contrastive objective to rank correct versus corrupted alternatives. Crucially, this training requires no real-world annotations yet teaches the verifier to recognize consistency cues that transfer across domains.

During adaptation on real videos, the verifier acts as a supervision selector: it filters and fuses per-frame reliability scores to form robust pseudo-label trajectories, reducing error accumulation and preventing collapse when any single

[†]Equal supervision

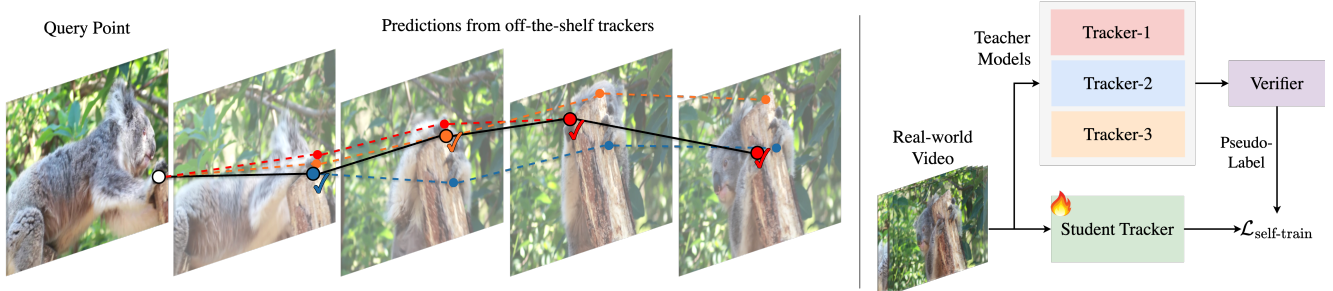


Figure 1. **Verifier-guided real-world adaptation.** (Left) Given a query point in a real-world video, multiple off-the-shelf trackers produce alternative trajectory hypotheses. **Verifier** evaluates these per-frame predictions and selects the most reliable ones, forming a refined pseudo-label trajectory. (Right) Unlike naïve self-training, which randomly selects a teacher model for pseudo-label generation, the verifier adaptively combines predictions from multiple teachers, providing cleaner supervision for the student tracker during real-world fine-tuning.

teacher fails (Fig. 1, right). The same mechanism can be used at inference time as a plug-and-play ensemble module, combining complementary trackers on the fly based on learned reliability cues rather than fixed weighting. This unified approach turns model diversity into a strength, yielding annotation-free adaptation, improved robustness to distribution shift, and better long-term coherence.

We validate our approach across diverse real-world datasets and tracking regimes. Verifier-guided pseudo-labeling consistently improves effectiveness during fine-tuning on unlabeled videos, and verifier-based ensembling at inference provides additional gains, particularly under challenging motion and occlusion. Ablations confirm that the verifier effectively exploits the complementary strengths of different trackers while remaining robust to individual failures. Together, these results show that the verifier provides a unified, data-efficient framework for reliable pseudo-labeling and model coordination in real-world point tracking.

In summary, our contributions are threefold: (i) a verifier, that learns to score and select reliable per-frame predictions from multiple trackers, enabling both training-time supervision selection and optional inference-time ensembling; (ii) a verifier-guided pseudo-labeling framework that scales fine-tuning on real videos while mitigating the failure modes of naïve pseudo-labeling; (iii) extensive experiments and ablations showing state-of-the-art performance and robust gains across real-world point-tracking benchmarks, highlighting the data efficiency and practicality of our approach.

2. Related Work

Point tracking. Tracking arbitrary points across long videos requires maintaining fine-grained correspondences under motion, occlusion, and reappearance. PIPs [18] introduced iterative refinement within temporal windows, and TAP-Vid [11] established a large-scale benchmark. TAPIR [12] improved temporal precision, while CoTracker [22] reformulated tracking as joint multi-point reasoning with trans-

formers. Subsequent works extended these ideas through query-based [28, 29] or region-level designs [8], later unified in CoTracker3 [21]. Track-On [1] recast tracking as patch classification for online inference, and Track-On2 [2] enhanced its efficiency. TAPNext [48] explored a complementary state-space approach. In contrast, we focus on improving real-world reliability through ensemble-based adaptation rather than new architecture design.

Pseudo-labeling for unlabeled data. Pseudo-labeling exploits model predictions as supervision for unlabeled data [27, 38, 40, 46], while ensemble learning [6, 15] combines predictors to improve generalization. Recent variants [7, 17, 26, 32] show that structured consensus across models yields more reliable supervision than single-teacher schemes. In point tracking, pseudo-labels bridge the gap between synthetic and real domains: BootsTAPIR [13] applies large-scale self-distillation, CoTracker3 [21] integrates real pseudo-labeled data, and AnthroTAP [25] uses 3D mesh priors for human motion. Our verifier extends this paradigm into a learnable ensemble, reasoning across trackers to select reliable labels via spatio-temporal cues.

3. Problem Formulation

Given an RGB video of T frames, $\mathcal{V} = \{\mathbf{I}_1, \dots, \mathbf{I}_T\} \in \mathbb{R}^{T \times H \times W \times 3}$, and a query point $\mathbf{q}_{t_0} \in \mathbb{R}^2$ at time t_0 , the goal of point tracking is to predict its trajectory and visibility across subsequent frames:

$$\{(\hat{\mathbf{p}}_t, \hat{v}_t)\}_{t=t_0+1}^T = \Phi(\mathcal{V}, \mathbf{q}_{t_0}) \quad (1)$$

where $\hat{\mathbf{p}}_t \in \mathbb{R}^2$ denotes the predicted 2D coordinates and $\hat{v}_t \in \{0, 1\}$ indicates whether the point is visible at time t .

Data setup & training regime. From the data perspective, we consider two domains: \mathcal{D}_{syn} and $\mathcal{D}_{\text{real}}$. The labeled synthetic dataset $\mathcal{D}_{\text{syn}} = \{(\mathcal{V}, \mathcal{Y})\}$ contains videos and their ground-truth point trajectories, where $\mathcal{Y} = \{(\mathbf{p}_t, v_t)\}_{t=1}^T$ denotes the trajectory and visibility of one queried point. In practice, multiple query points can be sampled per video,

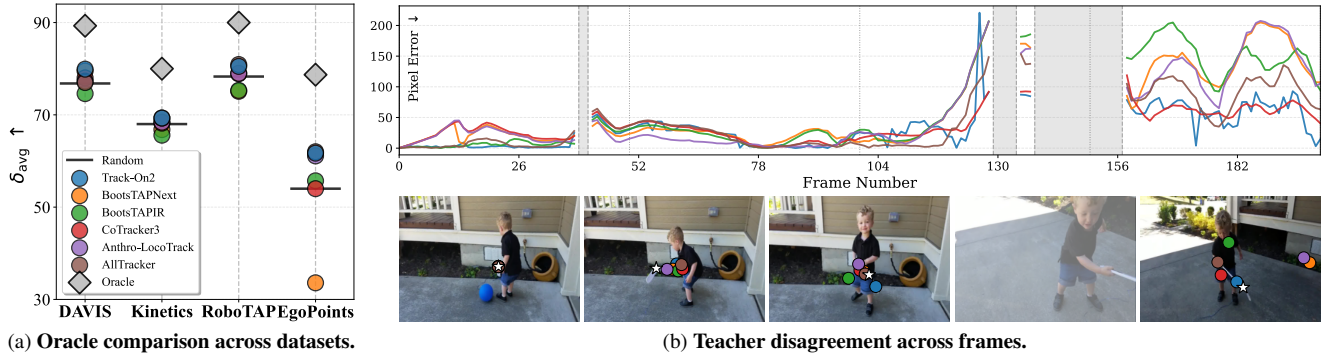


Figure 2. **Teacher inconsistency and oracle performance.** (a) Across 4 real-world datasets, six off-the-shelf teacher models (shown on the legend) are compared against an oracle that, at each frame, selects the most accurate teacher prediction. Individual teachers (colored circles) cluster below the oracle (diamonds), while the black horizontal line marks the performance of random teacher selection. The large gap between the oracle and both individual models and random selection highlights the substantial headroom available for adaptive, per-frame selection. (b) Example from TAP-Vid Kinetics [11]: Teacher predictions whose pixel errors fluctuate across time. The upper plot shows per-frame pixel error curves, with occluded frames shaded in gray. Colored lines correspond to the same trackers as in (a), illustrating that accuracy varies across time. The lower row shows uniformly sampled frames with teacher predictions and the ground-truth point (white star).

each with its own trajectory. The unlabeled real-world dataset $\mathcal{D}_{\text{real}} = \{\mathcal{V}\}$ consists of videos without trajectory annotations, *i.e.* $\mathcal{Y} = \emptyset$.

Existing approaches fall into two regimes: (i) training directly on \mathcal{D}_{syn} ; and (ii) pretraining on \mathcal{D}_{syn} followed by self-training on $\mathcal{D}_{\text{real}}$ using pseudo-labels.

Sim2Real training on \mathcal{D}_{syn} . Most existing point tracking models are trained on large-scale synthetic datasets \mathcal{D}_{syn} , such as TAP-Vid Kubric [11], and evaluated on both synthetic [20, 47] and real-world benchmarks [11, 43]. This setting assumes that models trained on synthetic videos can generalize to real ones, yet in practice, differences in texture, illumination, motion, and occlusion patterns introduce a domain gap, potentially leading to degraded performance.

Naïve self-training on $\mathcal{D}_{\text{real}}$. To bridge the domain gap, existing work performs self-training on real videos using pseudo-labels generated by pretrained teacher models [21]. Formally, a video $\mathcal{V} \sim \mathcal{D}_{\text{real}}$ is sampled, and a teacher $\Phi^{(m)}$ is randomly selected from a set of M pretrained models $\{\Phi^{(1)}, \dots, \Phi^{(M)}\}$. For randomly chosen query point \mathbf{q}_{t_0} , pseudo-labels $\tilde{\mathcal{Y}} = \{(\hat{\mathbf{p}}_t, \hat{v}_t)\}_{t=1}^T = \Phi^{(m)}(\mathcal{V}, \mathbf{q}_{t_0})$ are then used as supervision in place of the ground-truth \mathcal{Y} .

Discussion. A single teacher rarely provides uniformly reliable predictions across an entire video; accuracy fluctuates frame to frame, with intermittent successes and failures (Fig. 2b). In practice, different trackers exhibit error patterns that are inconsistent and only weakly correlated, yielding temporally unstable pseudo-labels rather than truly complementary signals. Naïvely sampling from, or averaging across, such predictions tends to amplify noise and drift, undermining the benefits of self-training on real videos (Fig. 2a). We address these limitations with a verifier that explicitly es-

timates per-frame reliability and selects among candidate trajectories accordingly. By conditioning supervision on learned reliability rather than fixed heuristics, the verifier stabilizes pseudo-labels, suppresses failure cascades, and turns model diversity into a practical advantage for adaptation.

4. Verifier

We introduce a **verifier**, trained on labeled synthetic videos, to assess the reliability of tracker predictions. At inference, it operates on real-world videos to select, at each time step, the most reliable estimate among multiple pretrained teacher trackers. Unlike conventional trackers, the verifier is a meta-model, *i.e.*, it does not predict trajectories directly but learns to evaluate which tracker is most reliable for each frame. Specifically, it compares the visual features around the predicted locations with those of the initial query point, reasoning across teachers and time to estimate a per-frame reliability score for each candidate trajectory. These scores are then used to construct a refined trajectory by combining the most reliable predictions over time, producing stable and accurate pseudo-labels for real-world fine-tuning.

An oracle test on existing tracking models. To measure the available headroom when combining multiple teachers, we construct an *oracle* that, at each time step, selects the prediction closest to the ground truth (Fig. 2a). As expected, the oracle performs best among all settings, but the large gap between it and the individual teachers highlights a significant potential for improvement. Across datasets, this gap remains consistent and is particularly pronounced on challenging egocentric videos such as EgoPoints [9]. These results indicate that the best-performing teacher varies across frames and scenes, making static or random selection suboptimal.

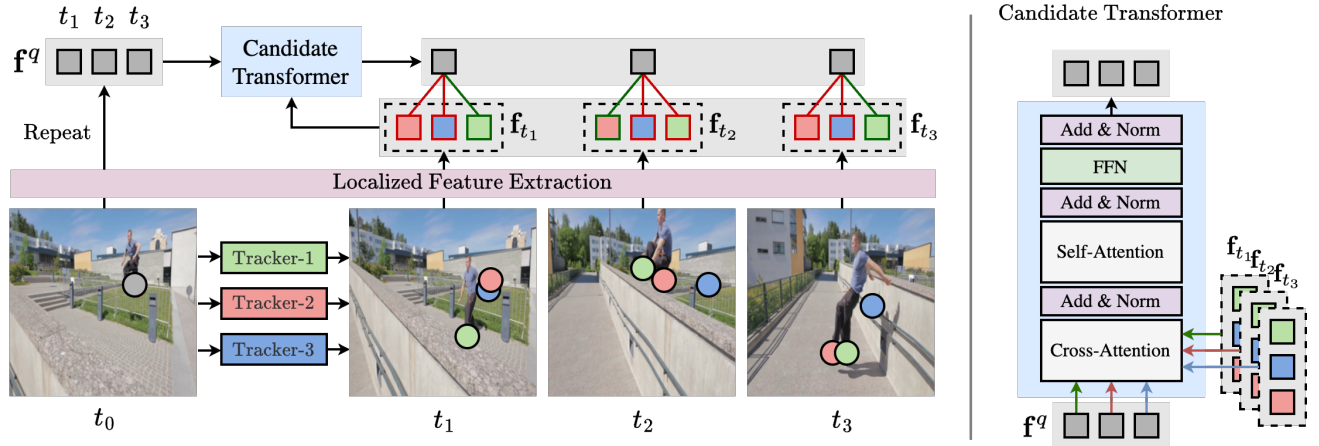


Figure 3. **Verifier overview.** Given query points at frame t_0 and their candidate predictions (teacher outputs during inference or randomly augmented trajectories during training), we extract local features for both the query and each candidate, producing query features \mathbf{f}^q (replicated across time) and candidate features \mathbf{f}_t (a vector for each candidate, per frame). The query features are then decoded by the Candidate Transformer (right), which consists of restricted cross-attention, where each frame-level query attends only to its corresponding candidates, followed by self-attention and feed-forward layers. The transformer outputs per-frame reliability distributions over candidates, capturing spatial and temporal consistency based on feature similarity.

Motivated by this observation, we propose to train a verifier that adaptively identifies the most reliable prediction among off-the-shelf trackers, producing more stable and accurate pseudo-labels $\hat{\mathcal{Y}}$ on $\mathcal{D}_{\text{real}}$.

Formally, the verifier takes as input a query point \mathbf{q}_{t_0} in a video \mathcal{V} and a tensor of M **candidate trajectories** that provide alternative motion estimates for the same target. We represent these trajectories as $\mathbf{C} \in \mathbb{R}^{L \times M \times 2}$, where $L = T - t_0$ is the trajectory length and each $\hat{\mathbf{p}}_t^{(m)}$ specifies the 2D location of candidate m at frame t . The verifier assigns a reliability score to each candidate at every time step, producing a set of per-frame score vectors:

$$\hat{\mathcal{S}} = \Phi_{\text{ver}}(\mathcal{V}, \mathbf{q}_{t_0}, \mathbf{C}) \quad (2)$$

where $\hat{\mathcal{S}} = \{\hat{\mathbf{s}}_t \mid t \in (t_0, T]\}$, each $\hat{\mathbf{s}}_t \in \mathbb{R}^M$ contains the reliability scores of all candidates at frame t , and $\hat{s}_t^{(m)}$ denotes the score assigned to candidate m . We describe how \mathbf{C} is constructed in Sec. 4.1. As illustrated in Fig. 3 (left), the verifier first extracts localized features for the query and all candidate trajectories at each frame (Sec. 4.2). The candidate transformer then reasons jointly over candidates and time to identify the most consistent motion hypothesis (Sec. 4.3). Finally, the verifier outputs per-frame reliability scores, selecting the most plausible candidate at each frame. These selected predictions are subsequently used to fine-tune a tracking model for real-world adaptation (Sec. 4.5).

4.1. Candidate Trajectories

Training setup. During training, the verifier is supervised on labeled synthetic videos $(\mathcal{V}, \mathcal{Y}) \sim \mathcal{D}_{\text{syn}}$. Each ground-truth

trajectory is represented by $\mathbf{p} \in \mathbb{R}^{L \times 2}$, where L is the trajectory length and \mathbf{p}_t gives the 2D location of the tracked point at frame t . The candidate trajectories \mathbf{C} are generated by applying random perturbations to \mathbf{p} , producing M (typically 6–12) perturbed versions, *i.e.* $\mathbf{p} \in \mathbb{R}^{L \times 2} \rightarrow \mathbf{C} \in \mathbb{R}^{L \times M \times 2}$, that simulate common prediction errors such as drift, occlusion, and jitter (see supplementary for details). These augmentations create diverse motion hypotheses resembling the errors observed at inference, allowing the verifier to learn to distinguish reliable trajectories from unreliable ones. Each candidate’s reliability target is computed from its distance to the corresponding ground-truth point, providing explicit supervision for training the verifier to rank predictions by reliability, as detailed in Sec. 4.4.

Inference setup. At inference and during real-world fine-tuning, \mathbf{C} consists of trajectories predicted by $M=6$ pretrained teacher trackers: Track-On2 [2], BootsTAPIR [13], BootsTAPNext [48], Anthro-LoCoTrack [25], AllTracker [19], and CoTracker3 (window variant) [21]. For each query point, the verifier evaluates these trajectories and assigns frame-wise reliability scores, selecting the highest-confidence prediction as the pseudo-label used for fine-tuning on $\mathcal{D}_{\text{real}}$ without manual annotations (Sec. 4.5).

4.2. Localized Feature Extraction

The verifier determines, at each frame, which candidate trajectory best continues the motion of the query point. To make this judgment, it compares the local visual evidence around the query with the appearance context around each candidate prediction. The localized feature extraction there-

fore converts image regions near these 2D locations into compact descriptors that are directly comparable.

Frame-wise feature extraction. We compute dense visual features for all frames using the frozen CNN encoder of pre-trained CoTracker3 [21], projected to dimension D through a linear layer:

$$\mathbf{F}_t = \phi_{\text{enc}}(\mathbf{I}_t), \quad \mathbf{F}_t \in \mathbb{R}^{H' \times W' \times D}, \quad (3)$$

where (H', W') are the downsampled spatial resolutions.

Query and candidate representations. The verifier compares each candidate prediction against a stable visual reference of the target. The query point serves as this reference, encoding the appearance and context of the target at its first visible frame t_0 . We first obtain this reference embedding by bilinearly sampling the feature map at \mathbf{q}_{t_0} :

$$\mathbf{q}_{\text{sample}} = \text{sample}(\mathbf{F}_{t_0}, \mathbf{q}_{t_0}) \in \mathbb{R}^D \quad (4)$$

For each frame t , the verifier measures how well the local region around each candidate in $\mathbf{C}_t \in \mathbb{R}^{M \times 2}$ matches the query reference. To capture this local context, deformable attention is applied at both the query and candidate locations, producing locally aggregated descriptors:

$$\mathbf{h}_{t_0}^q = \phi_{\text{def}}(\mathbf{q}_{\text{sample}}, \mathbf{F}_{t_0}, \mathbf{q}_{t_0}) \in \mathbb{R}^D, \quad (5)$$

$$\mathbf{h}_t = \phi_{\text{def}}(\mathbf{q}_{\text{sample}}, \mathbf{F}_t, \mathbf{C}_t) \in \mathbb{R}^{M \times D}, \quad t > t_0. \quad (6)$$

Here, $\phi_{\text{def}}(\cdot)$ uses the reference feature $\mathbf{q}_{\text{sample}}$ as the query input, the frame feature map as key-value pairs, and the spatial coordinates from \mathbf{C}_t as attention centers, ensuring that all candidate descriptors are computed relative to the same query appearance. In essence, these features capture how the appearance of the initial query is expressed across all candidate locations. This operation aggregates adaptive contextual information around each candidate, enabling the verifier to compare appearance similarity within local neighborhoods rather than relying on single-point features.

Position and identity embeddings. The extracted features are purely visual and lack explicit spatial context. To provide positional context, we apply a sinusoidal embedding $\eta(\cdot)$ to displacement vectors, and append a learned identity embedding that distinguishes the query (ID_0) from the candidate predictions (ID_1). After concatenation, a projection layer ϕ_{proj} maps the combined representation to the model width:

$$\mathbf{f}_{t_0}^q = \phi_{\text{proj}}(\mathbf{h}_{t_0}^q, \eta(\mathbf{0}), \text{ID}_0) \in \mathbb{R}^D, \quad (7)$$

$$\mathbf{f}_t = \phi_{\text{proj}}(\mathbf{h}_t, \eta(\mathbf{\Delta}_t), \text{ID}_1) \in \mathbb{R}^{M \times D}, \quad t > t_0, \quad (8)$$

where $\mathbf{\Delta}_t = \mathbf{C}_t - \mathbf{q}_{t_0}$ represents the displacement of each candidate from the initial query location, and ϕ_{proj} aligns the concatenated representation to the model dimension.

Output. The local encoder produces temporally aligned feature descriptors for both the query and candidate trajectories. The query feature $\mathbf{f}_{t_0}^q$ is replicated across time, *i.e.*

$\mathbf{f}_{t_0}^q \rightarrow \{\mathbf{f}_{t_0+1}^q, \dots, \mathbf{f}_T^q\}$, forming a tensor $\mathbf{f}^q \in \mathbb{R}^{L \times D}$ that encodes the reference appearance of the target at each frame, where each $\mathbf{f}_t^q \in \mathbb{R}^D$ corresponds to the query embedding at frame t . The candidate descriptors across all frames are stacked into a tensor $\mathbf{f} \in \mathbb{R}^{L \times M \times D}$, with $\mathbf{f}_t \in \mathbb{R}^{M \times D}$ representing the M candidate features at frame t . Together, \mathbf{f}^q and \mathbf{f} form the input to the candidate transformer.

4.3. Candidate Transformer

The candidate transformer decodes the query descriptors \mathbf{f}^q with the candidate descriptors \mathbf{f} to produce temporally informed embeddings for reliability estimation. It extends a standard transformer decoder [42] to reason jointly over candidates and time. Each layer consists of localized cross-attention, temporal self-attention, and a feed-forward network (Fig. 3, right).

At each frame t , the query embedding $\mathbf{f}_t^q \in \mathbb{R}^D$ attends only to its corresponding candidate features $\mathbf{f}_t \in \mathbb{R}^{M \times D}$, performing cross-attention along the candidate dimension M . During this step, the temporal dimension L is treated as a batch axis, so attention is computed independently for each frame. This step allows the query to integrate information from all candidate trajectories at that frame and identify the most consistent motion hypothesis. The resulting per-frame query embeddings are then connected across time through self-attention along the temporal dimension L , allowing information flow between frames. This temporal reasoning maintains consistency and helps resolve ambiguities: if a prediction is uncertain at one frame, contextual evidence from nearby frames with confident predictions can refine its reliability estimate. The outputs of these attention stages are passed through a feed-forward network to produce updated query representations for the next layer. No attention mask is applied, and all candidate predictions are included even when the point is occluded at that frame.

Verifier output. After decoding the query embeddings, the candidate transformer outputs a reliability distribution over candidates at each time step. Formally, for each frame t , we compute a temperature-scaled softmax over cosine similarities between the decoded query feature \mathbf{f}_t^q and the candidate features \mathbf{f}_t :

$$\hat{\mathbf{s}}_t = \text{Softmax}(\mathbf{f}_t^q \cdot \mathbf{f}_t / \tau), \quad \hat{\mathbf{s}}_t \in \mathbb{R}^M. \quad (9)$$

Here, $\mathbf{f}_t \in \mathbb{R}^{M \times D}$ represents all candidate features at frame t , $\mathbf{f}_t^q \in \mathbb{R}^D$ is the decoded query feature, and \cdot denotes cosine similarity applied row-wise between \mathbf{f}_t^q and each row of \mathbf{f}_t , scaled by the temperature $\tau = 0.1$. This yields a per-frame reliability distribution $\hat{\mathbf{s}}_t$ over the M candidates.

4.4. Training

Given the ground-truth trajectory $\mathbf{p} \in \mathbb{R}^{L \times 2}$ and its visibility sequence $\mathbf{v} \in \{0, 1\}^L$, the verifier is trained with a soft contrastive objective. Intuitively, candidate predictions that




Figure 4. **Verifier as inference time ensemble.** Comparison of the verifier ensemble against individual teacher models and the random-teacher baseline on real-world datasets. All teacher results are reproduced using their official checkpoints. The verifier consistently achieves the best performance across datasets, demonstrating its ability to exploit the complementary strengths of different models.

lie closer to the ground-truth locations should receive higher reliability, while those farther away should be downweighted. At each frame t , the verifier receives a set of M candidate predictions $\mathbf{C}_t \in \mathbb{R}^{M \times 2}$ and the corresponding ground-truth point $\mathbf{p}_t \in \mathbb{R}^2$. We define the per-frame target reliability distribution \mathbf{s}_t as a softmax over the negative distances between all candidates in \mathbf{C}_t and \mathbf{p}_t :

$$\mathbf{s}_t = \text{Softmax}(-\|\mathbf{C}_t - \mathbf{p}_t\|/\tau_s), \quad \mathbf{s}_t \in \mathbb{R}^M, \quad (10)$$

where $\|\mathbf{C}_t - \mathbf{p}_t\|$ computes the Euclidean distance from each candidate to the ground truth, and $\tau_s = 0.3$ controls the sharpness of the distribution. The predicted reliability scores $\hat{\mathbf{s}}_t$ are supervised using a cross-entropy loss $\mathcal{L} = \sum_{t=1}^L v_t \text{CE}(\hat{\mathbf{s}}_t, \mathbf{s}_t)$, where v_t masks out occluded frames, ensuring that only visible points contribute to the objective.

4.5. Fine-tuning with Verified Pseudo-Labels

We start from a point tracking model pretrained on labeled synthetic data \mathcal{D}_{syn} and adapt it to real-world videos from $\mathcal{D}_{\text{real}}$. In our case, the base model is Track-On2 [2], pretrained on TAP-Vid Kubric [11]. For real-world adaptation, $\mathcal{D}_{\text{real}}$ is constructed from in-the-wild object tracking and segmentation datasets, namely TAO [10], OVIS [35], and VSPW [33]. We retain videos longer than 48 frames without additional filtering, resulting in 4864 real-world sequences with diverse motion patterns and scene content.

To fine-tune the model on $\mathcal{D}_{\text{real}}$, we employ the trained verifier to generate reliable pseudo-labels. For each video $\mathcal{V} \sim \mathcal{D}_{\text{real}}$, we sample query points based on both appearance and motion cues. Specifically, 2/3 of the queries are drawn from SIFT [31] detections, while the remaining 1/3 are selected from motion-salient regions obtained via grayscale frame differencing with mild spatial smoothing. Queries are sampled from the first half of the video, using four uniformly spaced frames (*e.g.*, frames 1, 5, ...).

Each query is tracked by the ensemble of teacher models, producing a set of candidate trajectories. At every frame t , the candidate with the highest reliability score in $\hat{\mathbf{s}}_t$, *i.e.* the

prediction corresponding to $\arg \max_m \hat{s}_t^{(m)}$, is selected as the pseudo-label. Visibility is estimated via majority voting across teacher predictions. This verifier-guided pseudo-labeling enables efficient adaptation to real-world domains without requiring manual annotations or labeled data.

In practice, we fine-tune using both \mathcal{D}_{syn} (with ground-truth annotations) and $\mathcal{D}_{\text{real}}$, gradually increasing the loss weight of real-world samples while reducing that of synthetic ones during training. This schedule yields slightly better performance than training on $\mathcal{D}_{\text{real}}$ alone, although using only $\mathcal{D}_{\text{real}}$ still produces competitive results (see Sec. 5.4).

5. Experiments

5.1. Setup

Datasets. For verifier training, we use K-EPIC [9], a synthetic dataset of 11K videos, each 24 frames long, containing similar synthetic objects to those in TAP-Vid Kubric [11]. For real-world fine-tuning with the verifier, we use 8K videos from our filtered collection described in Sec. 4.5. We evaluate our model on four public real-world datasets with diverse characteristics: TAP-Vid DAVIS, 30 real-world videos from the DAVIS dataset; TAP-Vid Kinetics, 1000 videos from the validation split of Kinetics-700-2020; RoboTAP [43], 265 robotic sequences averaging over 250 frames each; and EgoPoints [9], long egocentric videos spanning up to several thousand frames.

Metrics. We evaluate tracking performance on TAP-Vid subsets and RoboTAP using three standard metrics: Occlusion Accuracy (OA), visibility prediction accuracy; δ_{avg}^x , the average proportion of visible points tracked within thresholds of 1, 2, 4, 8, and 16 pixels; and Average Jaccard (AJ), a combined measure of localization and visibility. On EgoPoints, we report δ_{avg} and OA following the official evaluation.

Evaluation details. We follow the standard TAP-Vid benchmarking protocol, downsampling all videos to 256×256 resolution except for EgoPoints. Models are evaluated in the queried-first setting, corresponding to the causal tracking

Table 1. **Quantitative results on real-world datasets.** Comparison with prior work on EgoPoints, RoboTAP, TAP-Vid Kinetics, and TAP-Vid DAVIS in terms of AJ, δ_{avg}^x , and OA (higher is better). Models are grouped into **synthetic pretraining** and **real-world fine-tuning**. The latter includes methods fine-tuned on additional real-world videos using semi-supervised strategies. We report officially published results when available; if no reported values exist, we evaluate the authors’ official checkpoints. Missing checkpoints are shown as “—”.

[†] AllTracker leverages additional optical-flow datasets (synthetic and real-world with ground-truth), without pseudo-label fine-tuning.

Model	EgoPoints		RoboTAP			Kinetics			DAVIS		
	δ_{avg}^x	OA	AJ	δ_{avg}^x	OA	AJ	δ_{avg}^x	OA	AJ	δ_{avg}^x	OA
Synthetic Pretraining											
TAPIR [12]	50.2	79.9	59.6	73.4	87.0	49.6	64.2	85.0	56.2	70.0	86.5
LocoTrack [8]	59.6	88.6	62.3	76.2	87.1	52.9	66.8	85.3	63.0	75.3	87.2
TAPNext-B [48]	31.8	73.8	59.5	72.8	88.0	53.3	67.9	87.0	62.4	76.6	90.5
CoTracker3 (Window) [21]	—	—	60.8	73.7	87.1	54.1	66.6	87.1	64.5	76.7	89.7
Track-On2 [2]	61.7	<u>89.9</u>	68.1	80.5	<u>93.4</u>	55.3	69.3	<u>89.6</u>	<u>67.0</u>	<u>79.9</u>	<u>92.0</u>
Real-World Fine-Tuning											
BootsTAPIR [13]	55.7	78.2	64.9	80.1	86.3	54.6	68.4	86.5	61.4	73.6	88.7
Anthro-LoCoTrack [25]	61.1	89.5	64.7	79.2	88.4	53.9	68.4	86.4	64.8	77.3	89.1
BootsTAPNext-B [48]	33.6	68.5	64.0	75.0	88.7	<u>57.3</u>	<u>70.6</u>	87.4	65.2	78.5	91.2
CoTracker3 (Window) [21]	54.0	84.4	66.4	78.8	90.8	55.8	68.5	88.3	63.8	76.3	90.2
AllTracker [†] [19]	<u>62.0</u>	87.1	<u>68.8</u>	<u>80.9</u>	92.2	56.8	69.3	89.1	63.7	77.0	88.7
Track-On-R (Ours)	67.3	90.2	70.9	82.6	94.0	57.8	71.0	90.5	68.1	80.3	92.5

scenario: the first visible point in each trajectory is used as the query, and the model tracks that point in subsequent frames. For EgoPoints, frames are resized to 384×512 , however evaluation is done on 256×256 , consistent with its official evaluation setup.

5.2. Results: Verifier as Inference-time Ensemble

We first evaluate the verifier as an inference-time ensemble to assess its ability to identify the most reliable tracker per frame. As shown in Fig. 4, we compare the **reproduced** performance of each teacher model using their official checkpoints, the random-teacher baseline (where a teacher is randomly selected per video), and the verifier-based selection across four real-world datasets. Reproduced results are used for all comparisons. These correspond to the exact candidate trajectories provided to the verifier, ensuring that frame-wise selection is evaluated fairly.

The verifier consistently outperforms both the random baseline and the strongest individual teacher. This demonstrates its ability to adaptively select the most reliable prediction at each time step. The gains are particularly pronounced on challenging datasets such as EgoPoints, where model reliability fluctuates significantly across motion patterns and scene content.

Performance rankings also vary substantially across datasets. For example, BootsTAPNext [48] ranks last on RoboTAP but second on DAVIS, highlighting that no single tracker is universally dominant. These results confirm that the verifier provides a principled mechanism for aggregat-

ing complementary teacher models and serves as a robust adaptive ensemble strategy, validating its design prior to real-world fine-tuning.

5.3. Results: Verifier-Guided Adaptation

In this section, we compare our model **Track-On-R**, fine-tuned on real-world videos using the verifier-guided approach, against prior work. As summarized in Table 1, we categorize existing approaches into two groups: **synthetic pretraining** and **real-world fine-tuning**, presented side by side for direct comparison. Models in the first group are trained exclusively on synthetic data from TAP-Vid Kubric [11], whereas those in the second group leverage additional real-world videos for fine-tuning to enhance generalization. When official results are available, we report the numbers published by the original authors. If a method has no officially reported values, we reproduce the results using the authors’ released checkpoints. For methods without any publicly released checkpoints or reported values, we indicate missing entries with “—”.

Baseline models. Existing real-world adaptation strategies differ in both supervision and data scale. BootsTAPIR [13] and BootsTAPNext [48] rely on large-scale student–teacher distillation over millions of real videos. AnthroTAP [25] instead uses domain-specific human-mesh annotations to derive pseudo-labels, which limits its applicability outside human-centric datasets. AllTracker [19] does not perform real-world fine-tuning, but is trained on high-quality synthetic optical flow with ground-truth as well as annotated

Table 2. **Effect of teacher composition on verifier performance.** We report δ_{avg}^x for the random-teacher baseline (Rand.) and the verifier-based selection (Ver.) across different teacher subsets.

	Teacher Models					DAVIS		RoboTAP	
	A	B	C	D	E	Rand.	Ver.	Rand.	Ver.
(i)	✓	✓	✗	✗	✗	79.5	80.6	77.4	81.8
(ii)	✓	✓	✓	✗	✗	79.4	80.6	77.3	82.4
(iii)	✓	✓	✓	✓	✗	77.5	80.8	77.4	82.8
(iv)	✓	✓	✓	✓	✓	77.7	81.1	78.0	83.1

real-world flow data. The most similar approach is CoTracker3 [21], which applies a generic random-teacher pseudo-labeling strategy on real-world videos. Our method also follows a teacher-based pseudo-labeling paradigm, but introduces a verifier that enables reliable adaptation to arbitrary real-world video collections using only raw videos.

Comparison on EgoPoints. EgoPoints presents a particularly challenging setting due to long sequences and extended occlusion periods. Our model significantly improves over the synthetic baseline, increasing δ_{avg}^x from 61.7 to 67.3. It also widens the gap to the closest competitor, AllTracker [19], by +5.3 in δ_{avg}^x . These results demonstrate that our pseudo-labels transfer effectively to complex ego-centric motion, even though the training data are not explicitly ego-centric.

Comparison on RoboTAP. Compared to the synthetic baseline, our model achieves consistent gains across all metrics, including a +2.8 improvement in AJ. Our adaptation attains the highest scores on the benchmark, outperforming both synthetic and real-world baselines. Notably, our real-world training set contains no robotic sequences, indicating that the verifier-guided pseudo-labels generalize well and transfer effectively to robotic manipulation scenarios.

Comparison on Kinetics. On TAP-Vid Kinetics, our model improves over the synthetic baseline by +2.5 in AJ. It also surpasses the highest AJ score of 57.3 reported by BootsTAPNext [48], while achieving 3.1 points higher in OA, indicating more accurate visibility prediction.

Comparison on DAVIS. Our model achieves the highest scores across all metrics, outperforming both synthetic and real-world baselines. Compared to the baseline Track-On2 [2], our fine-tuned variant improves AJ by +1.1, establishing a new state of the art on this benchmark.

5.4. Ablation Study

Teacher models. To analyze how the composition of teacher models influences performance, Table 2 compares the random selection (round-robin selection) and our verifier-based selection over different teacher subsets in terms of δ_{avg}^x on TAP-Vid DAVIS and RoboTAP. Models A–E correspond to Track-On2 [2], BootsTAPNext [48], BootsTAPIR [13],

Table 3. **Synthetic vs. real data during fine-tuning.** We compare three configurations: Real (only $\mathcal{D}_{\text{real}}$), Mix ($\mathcal{D}_{\text{real}} + \mathcal{D}_{\text{syn}}$), and Mix + Schedule, where the loss weight of real-world samples is gradually increased during training.

Data	EgoPoints		RoboTAP		Kinetics		DAVIS	
	δ_{avg}^x	OA	δ_{avg}^x	OA	δ_{avg}^x	OA	δ_{avg}^x	OA
Real	66.9	90.2	82.7	93.8	70.9	90.4	80.3	92.4
Mix	65.6	90.7	82.3	93.9	70.8	90.3	80.7	92.5
Mix + Schedule	67.3	90.2	82.6	94.0	71.0	90.5	80.3	92.5

CoTracker [21], and Anthro-LoCoTrack [25]. The verifier consistently outperforms the random baseline and remains robust even when weaker teachers are included. Adding or removing a model that lowers the baseline average either preserves or improves verifier accuracy. For example, in DAVIS, adding D to {A,B,C} (rows ii–iii) lowers the random baseline but increases verifier accuracy. These trends indicate that the verifier effectively exploits complementary teacher behaviors rather than being diluted by weaker models.

Synthetic–real training schedule. We study the effect of training data composition during fine-tuning in Table 3. We compare three settings: using only real-world videos from $\mathcal{D}_{\text{real}}$, mixing $\mathcal{D}_{\text{real}}$ with synthetic data \mathcal{D}_{syn} (with ground-truth supervision), and the same mixture with a schedule that gradually increases the loss weight of real-world samples (our default). The mixture improves OA due to visibility supervision from synthetic data, while real-only training yields slightly better localization in terms of δ_{avg}^x . The scheduled mixture combines both advantages and achieves the best overall performance, although adding synthetic data on top of real-only training provides only marginal gains.

6. Conclusion & Limitation

In this paper, we have introduced a verifier module that learns to assess the reliability of point tracker predictions, enabling robust real-world adaptation without manual annotations. When used for fine-tuning, it produces high-quality pseudo-labels that substantially improve tracking performance across diverse domains and achieve data-efficient adaptation. Beyond pseudo-label generation, the verifier also functions as an effective inference-time ensemble mechanism, offering a general framework for model selection and uncertainty estimation in video correspondence tasks. However, the effectiveness of fine-tuning still depends on the quality and diversity of the available video data, which highlights the importance of curated real-world collections. In addition, the verifier’s upper bound is constrained by the quality of its teacher trackers, suggesting that future work should focus on developing stronger and more robust pretrained models to further enhance verifier-guided learning.

Acknowledgement

This project is funded by the European Union (ERC, ENSURE, 101116486) with additional compute support from Leonardo Booster (EuroHPC Joint Undertaking, EHPC-AI-2024A05-028). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. Weidi would like to acknowledge the funding from Scientific Research Innovation Capability Support Project for Young Faculty (ZY-GXQNJSKYCXNLZCXM-I22).

References

- [1] Görkay Aydemir, Xiongyi Cai, Weidi Xie, and Fatma Güney. Track-On: Transformer-based online point tracking with memory. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2025. [1](#), [2](#)
- [2] Görkay Aydemir, Weidi Xie, and Fatma Güney. Track-on2: Enhancing online point tracking with memory. *arXiv preprint arXiv:2509.19115*, 2025. [2](#), [4](#), [6](#), [7](#), [8](#)
- [3] Md Abulkalam Azad, John Nyberg, Havard Dalen, Bjørnar Grenne, Lasse Lovstakken, and Andreas Østvik. Taming modern point tracking for speckle tracking echocardiography via impartial motion. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV) Workshops*, 2025. [1](#)
- [4] Abhishek Badki, Hang Su, Bowen Wen, and Orazio Gallo. L4p: Low-level 4d vision perception unified. *arXiv preprint arXiv:2502.13078*, 2025. [1](#)
- [5] Homanga Bharadhwaj, Roozbeh Mottaghi, Abhinav Gupta, and Shubham Tulsiani. Track2Act: Predicting point tracks from internet videos enables diverse zero-shot robot manipulation. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2024. [1](#)
- [6] Leo Breiman. Bagging predictors. In *Machine Learning*, 1996. [2](#)
- [7] Liangyu Chen, Bo Li, Sheng Shen, Jingkan Yang, Chunyuan Li, Kurt Keutzer, Trevor Darrell, and Ziwei Liu. Large language models are visual reasoning coordinators. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. [2](#)
- [8] Seokju Cho, Jiahui Huang, Jisu Nam, Honggyu An, Seungryong Kim, and Joon-Young Lee. Local all-pair correspondence for point tracking. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2024. [2](#), [7](#)
- [9] Ahmad Darkhalil, Rhodri Guerrier, Adam W Harley, and Dima Damen. Egopoints: Advancing point tracking for ego-centric videos. In *Proc. of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2025. [3](#), [6](#), [1](#)
- [10] Achal Dave, Tarasha Khurana, Pavel Tokmakov, Cordelia Schmid, and Deva Ramanan. Tao: A large-scale benchmark for tracking any object. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2020. [6](#)
- [11] Carl Doersch, Ankush Gupta, Larisa Markeeva, Adria Recasens, Lucas Smaira, Yusuf Aytar, Joao Carreira, Andrew Zisserman, and Yi Yang. TAP-Vid: A benchmark for tracking any point in a video. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [12] Carl Doersch, Yi Yang, Mel Vecerik, Dilara Gokay, Ankush Gupta, Yusuf Aytar, Joao Carreira, and Andrew Zisserman. TAPIR: Tracking any point with per-frame initialization and temporal refinement. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2023. [2](#), [7](#)
- [13] Carl Doersch, Pauline Luc, Yi Yang, Dilara Gokay, Skanda Koppula, Ankush Gupta, Joseph Heyward, Ignacio Rocco, Ross Goroshin, João Carreira, and Andrew Zisserman. BootstAP: Bootstrapped training for tracking-any-point. *Proc. of the Asian Conf. on Computer Vision (ACCV)*, 2024. [2](#), [4](#), [7](#), [8](#)
- [14] Haiwen Feng, Junyi Zhang, Qianqian Wang, Yufei Ye, Pengcheng Yu, Michael J Black, Trevor Darrell, and Angjoo Kanazawa. St4rtrack: Simultaneous 4d reconstruction and tracking in the world. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2025. [1](#)
- [15] Yoav Freund, Robert E Schapire, et al. Experiments with a new boosting algorithm. In *Proc. of the International Conf. on Machine Learning (ICML)*, 1996. [2](#)
- [16] Daniel Geng, Charles Herrmann, Junhua Hur, Forrester Cole, Serena Zhang, Tobias Pfaff, Tatiana Lopez-Guevara, Yusuf Aytar, Michael Rubinstein, Chen Sun, et al. Motion prompting: Controlling video generation with motion trajectories. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2025. [1](#)
- [17] Fredrik K Gustafsson, Martin Danelljan, and Thomas B Schon. Evaluating scalable bayesian deep learning methods for robust computer vision. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2020. [2](#)
- [18] Adam W Harley, Zhaoyuan Fang, and Katerina Fragkiadaki. Particle video revisited: Tracking through occlusions using point trajectories. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2022. [1](#), [2](#)
- [19] Adam W. Harley, Yang You, Xinglong Sun, Yang Zheng, Nikhil Raghuraman, Yunqi Gu, Sheldon Liang, Wen-Hsuan Chu, Achal Dave, Pavel Tokmakov, Suya You, Rares Ambrus, Katerina Fragkiadaki, and Leonidas J. Guibas. AllTracker: Efficient dense point tracking at high resolution. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2025. [4](#), [7](#), [8](#), [2](#)
- [20] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. DynamicStereo: Consistent dynamic depth from stereo videos. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2023. [3](#), [2](#)
- [21] Nikita Karaev, Iurii Makarov, Jianyuan Wang, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. CoTracker3: Simpler and better point tracking by pseudo-labelling real videos. *arXiv preprint arXiv:2410.11831*, 2024. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#), [8](#)
- [22] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. CoTracker: It is better to track together. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2024. [1](#), [2](#)
- [23] Mert Asim Karaoglu, Wenbo Ji, Ahmed Abbas, Nassir Navab, Benjamin Busam, and Alexander Ladikos. Litetracker: Lever-

- aging temporal causality for accurate low-latency tissue tracking. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2025. 1
- [24] Yoni Kasten, Wuyue Lu, and Haggai Maron. Fast encoder-based 3d from casual videos via point track processing. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 1
- [25] Inès Hyeonsu Kim, Seokju Cho, Jahyeok Koo, Junghyun Park, Jiahui Huang, Joon-Young Lee, and Seungryong Kim. Learning to track any points from human motion. *arXiv preprint arXiv:2507.06233*, 2025. 2, 4, 7, 8
- [26] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 2
- [27] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Proc. of the International Conf. on Machine learning (ICML) Workshops*, 2013. 2
- [28] Hongyang Li, Hao Zhang, Shilong Liu, Zhaoyang Zeng, Feng Li, Tianhe Ren, Bohan Li, and Lei Zhang. TAP-TRv2: Attention-based position update improves tracking any point. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 2
- [29] Hongyang Li, Hao Zhang, Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, and Lei Zhang. TAPTR: Tracking any point with transformers as detection. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2024. 1, 2
- [30] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2019. 1, 2
- [31] David G Lowe. Object recognition from local scale-invariant features. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 1999. 6, 2
- [32] Zhihe Lu, Jiawang Bai, Xin Li, Zeyu Xiao, and Xinchao Wang. Beyond sole strength: Customized ensembles for generalized vision-language models. In *Proc. of the International Conf. on Machine learning (ICML)*, 2024. 2
- [33] Jiaxu Miao, Yunchao Wei, Yu Wu, Chen Liang, Guangrui Li, and Yi Yang. Vspw: A large-scale dataset for video scene parsing in the wild. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 6
- [34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 2
- [35] Jiyang Qi, Yan Gao, Yao Hu, Xinggang Wang, Xiaoyu Liu, Xiang Bai, Serge Belongie, Alan Yuille, Philip Torr, and Song Bai. Occluded video instance segmentation: A benchmark. In *International Journal of Computer Vision (IJCV)*, 2022. 6
- [36] Chris Rockwell, Joseph Tung, Tsung-Yi Lin, Ming-Yu Liu, David F Fouhey, and Chen-Hsuan Lin. Dynamic camera poses and where to find them. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2025. 1
- [37] Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. DINOv3. *arXiv preprint arXiv:2406.08085*, 2025. 2
- [38] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [39] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [40] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 2
- [41] Zachary Teed and Jia Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2020. 1
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 5
- [43] Mel Vecerik, Carl Doersch, Yi Yang, Todor Davchev, Yusuf Aytar, Guangyao Zhou, Raia Hadsell, Lourdes Agapito, and Jon Scholz. RoboTAP: Tracking arbitrary points for few-shot visual imitation. In *Proc. IEEE International Conf. on Robotics and Automation (ICRA)*, 2023. 1, 3, 6
- [44] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. VggT: Visual geometry grounded transformer. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2025. 1
- [45] Chuan Wen, Xingyu Lin, John So, Kai Chen, Qi Dou, Yang Gao, and Pieter Abbeel. Any-point trajectory modeling for policy learning. *arXiv preprint arXiv:2401.00025*, 2023. 1
- [46] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [47] Yang Zheng, Adam W Harley, Bokui Shen, Gordon Wetzstein, and Leonidas J Guibas. PointOdyssey: A large-scale synthetic dataset for long-term point tracking. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2023. 3, 2
- [48] Artem Zhohlov, Carl Doersch, Yi Yang, Skanda Koppula, Viorica Patraucean, Xu Owen He, Ignacio Rocco, Mehdi SM Sajjadi, Sarath Chandar, and Ross Goroshin. Tapnext: Tracking any point (tap) as next token prediction. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2025. 2, 4, 7, 8