

Refining Pre-Trained Motion Models

Xinglong Sun¹, Adam W. Harley¹, and Leonidas J. Guibas¹

Abstract—Given the difficulty of manually annotating motion in video, the current best motion estimation methods are trained with synthetic data, and therefore struggle somewhat due to a train/test gap. Self-supervised methods hold the promise of training directly on real video, but typically perform worse. These include methods trained with warp error (i.e., color constancy) combined with smoothness terms, and methods that encourage cycle-consistency in the estimates (i.e., tracking backwards should yield the opposite trajectory as tracking forwards). In this work, we take on the challenge of improving state-of-the-art supervised models with self-supervised training. We find that when the initialization is supervised weights, most existing self-supervision techniques actually make performance worse instead of better, which suggests that the benefit of seeing the new data is overshadowed by the noise in the training signal. Focusing on obtaining a “clean” training signal from real-world unlabelled video, we propose to separate label-making and training into two distinct stages. In the first stage, we use the pre-trained model to estimate motion in a video, and then select the subset of motion estimates which we can verify with cycle-consistency. This produces a sparse but accurate pseudo-labelling of the video. In the second stage, we fine-tune the model to reproduce these outputs, while also applying augmentations on the input. We complement this bootstrapping method with simple techniques that densify and re-balance the pseudo-labels, ensuring that we do not merely train on “easy” tracks. We show that our method yields reliable gains over fully-supervised methods in real videos, for both short-term (flow-based) and long-range (multi-frame) pixel tracking. Our code can be found here: <https://github.com/AlexSunNik/refining-motion-code>.

I. INTRODUCTION

Given the difficulty of manually annotating motion in video, the current best motion estimation methods are trained with synthetic data [1, 2, 3, 4, 5]. This synthetic training data is diverse, but highly unrealistic, forcing the methods to contend with a sim-to-real gap.

Self-supervised (or “unsupervised”) models hold the promise of training directly on real video, yet lag behind their supervised counterparts. State-of-the-art self-supervised motion methods typically rely on warp error (i.e., color constancy) combined with smoothness terms [6, 7, 8], and cycle-consistency (i.e., tracking backwards should yield the opposite trajectory as tracking forwards) [9, 10, 11, 12]. In current experimental setups, fully-supervised versions are typically positioned in a league ahead, reflecting in some way an unfair advantage. This divide between supervised and unsupervised methods is purely academic; if synthetic supervision is available and helpful, it makes sense to use it.

In this work, we improve state-of-the-art supervised models with self-supervised training in an unlabelled test domain.

¹Stanford University, Computer Science Department. {xs15, harleya, guibas}@stanford.edu

Unlike typical self- or semi-supervised motion estimation approaches which are trained from scratch, we refine an off-the-shelf pre-trained model for a given test domain without requiring any additional data.

We argue that this setup reflects a very practical use-case: a user notices some errors in the model’s output, and is willing to invest some GPU cycles into tuning up the results—provided that the tuning is effective.

In this setup, the most straightforward solution is to directly optimize standard self-supervision objectives in the unlabelled data. We find that this reliably makes the pre-trained model perform even worse, suggesting that the self-supervised training signals are too imprecise and noisy to provide useful gradients to an already well-optimized motion model. This observation is consistent with the fact that self-supervised models perform worse than supervised ones in general.

In order to obtain a “clean” training signal from real-world unlabelled video, we propose to separate labeling and training into two distinct stages. In the first stage, we use the pre-trained model to estimate motion in a video, and then select the subset of motion estimates which we can verify with cycle-consistency. This produces a sparse but accurate pseudo-labeling of the video. In the second stage, we fine-tune the model to reproduce these outputs, while also applying augmentations on the input. We complement this self-training method with simple techniques that densify and re-balance the pseudo-labels, ensuring that we do not merely train on “easy” tracks.

In our experiments, we demonstrate that our method yields reliable gains over fully-supervised methods in real videos. Our experiments cover two types of motion models: optical flow (i.e., RAFT [13]) and long-term point tracking (i.e., PIPs [14]). We use standard benchmarks for the analysis: MPI-Sintel [3] for flow; CroHD [15] and TAP-Vid-DAVIS [16] for point tracking. Despite starting from pre-trained state-of-the-art weights, our self-supervision framework produces consistent improvements in results. We hope that our technique, and our semi-supervised setup, will inspire future work into refining pre-trained motion models.

II. RELATED WORK

Learning motion from color and smoothness. Soon after the introduction of deep-learned flow methods [1], there has been interest in learning optical flows in self-supervised manner [6], converting classic flow assumptions [17], color constancy, and motion smoothness, into supervision objectives. Color constancy implies that when flow is accurate, corresponding pixels in consecutive frames should have the same color. By warping images to align with estimated flow,

reducing per-pixel differences can improve flow accuracy, particularly when computed across multiple scales [18]. The smoothness assumption helps to resolve ambiguities and occlusions which can be simply implemented by penalizing spatial gradients on the output flow. Minute details in this setup can affect performance greatly [19]. Instead of relying on color and smoothness, which can be understood as indirectly encouraging correct motion, we compute pseudo-labels for unlabelled data, providing direct regression targets.

Bootstrapping with student-teacher setups. Some recent works stabilize the losses from color-matching by using student-teacher setups. In this setting, the teacher is a moving average copy of the student, and the student receives harder data than the teacher [7, 8, 20]. Our work can also be considered a student-teacher setup, but the teacher is a frozen copy of a pre-trained model, and the student is a trainable version of the same model. Our setup is in line with knowledge distillation [21], except our pair of models begins with the same architecture and weights (rather than the student being a smaller model).

Learning motion from cycle-consistency. Color matching is known to break down under some conditions (e.g., specularities and occlusions), so many works focus instead on learning directly from cycle-consistency [9, 10, 11]. The key idea is that after tracking a target from a given startpoint to an estimated endpoint, reversing the video and re-applying the tracker from the endpoint should lead back to the original startpoint. This core idea is typically combined with patch-level affinity matrices, which can be traversed with spatial transformers [11], region-level motion averages [10], or random walks [12, 22, 23]. Unlike existing self-supervision works which initialize weights from ImageNet [24] or randomly, our method begins with state-of-the-art models whose architectures and weights are pre-optimized for motion estimation [13, 14] and attempts to further refine them on a new test domain. Compared to previous works which directly leverage cycle consistency as part of training objective [11], we use it to design a filter to select cycle-consistent motion estimates for the finetuning.

Semi-supervised motion estimation. Semi-supervised methods in motion estimation typically exploit a mix of labeled and unlabelled data, hoping to achieve better performance than exclusively using labeled data. Lai et al. [25] propose to learn a discriminator on warp errors from ground truth warps to provide feedback on the overall quality of the warps, replacing the noisy per-pixel color cues. Jeong et al. [26] add a separate segmentation module to a flow network, supervise this module with occlusion masks applied onto unlabelled images, and also ask flows to be consistent across transformations applied to the input, which is similar to our goal but requires modified architecture for segmentation.

Training from pseudo labels. The technique of training on self-generated estimates (i.e., pseudo labels) was first proposed by Lee et al. [27]. This core idea is often paired with methods to filter the training targets to “confident” ones, where confidence may be approximated with the help of ensemble methods [28, 29]. Our work computes confidence with the

help of domain knowledge: if a trajectory is within some cycle-consistency margin and color constancy, we consider it a pseudo-label. Similar to Chen et al. [29], we assume that it is possible to train on a small number of confident samples and generalize to other regions of the domain.

III. METHOD

Our method begins with a pre-trained motion estimator and unlabelled videos from test domain, and attempts to improve the estimator performance with self-supervision on the input videos. Our overall framework, illustrated in Fig. 1, consists of two stages. In the first stage, we generate pseudo-labels, by running the pre-trained model on the videos and selecting a subset of model estimates with cycle-consistency-based filters. In the second stage, we refine the model by fine-tuning it on those pseudo-labels, challenging it to reproduce those estimates under augmentations. In this section, we step through the method: we first describe our motion models, then discuss pseudo-labeling, and end with refinement.

A. Preliminaries: RAFT and PIPs

Our approach can be applied to any motion model, but we focus here on two popular methods: RAFT [13] and PIPs [14]. These are fine-grained general-purpose motion models, in the sense that they can track arbitrary small elements (i.e., points on the scene surface, specified by an (x, y) coordinate in a given frame) in arbitrary videos (i.e., not restricted to particular domains, though potentially suffering a generalization gap). We use pre-trained weights publicly released by the authors, which were produced by training in synthetic data. In general, any motion model could be used, but we choose RAFT and PIPs as they represent the state-of-the-art for optical flow and multi-frame point tracking.

RAFT is a *2-frame dense-motion model*. It consumes two consecutive RGB frames as input, and produces a dense flow map as output. The flow map is a spatial displacement map, indicating the frame-2 correspondence for each pixel in frame-1. Optical flow can be linked across time to form longer-range correspondences, but these trajectories are sensitive to occlusions, and slow to compute. RAFT works by iteratively refining an estimated flow map, using feature correlations indexed in a local window around the current estimates.

PIPs is an *8-frame sparse-motion model*. It consumes a sequence of 8 consecutive frames as input, along with a list of target coordinates to track, and it produces an 8-frame trajectory for each target. These 8-frame trajectories are less sensitive to occlusions than optical flow, but are typically sparse since computing these trajectories densely is expensive. PIPs’ design is inspired by RAFT, and similarly operates by iteratively refining its estimates, using feature correlations indexed in a local window around the current estimates.

B. First Stage: Pseudo-label Generation

In the first stage, we generate pseudo-labels, by applying the pre-trained model to unlabeled videos and collecting a reliable subset of the motion estimates, where reliability is estimated with cycle-consistency.

Cycle-consistency leverages the fact that if we track forwards in time without error, then start at the predicted endpoint and track backward in time without error, we will end up at the original startpoint. Errors in tracking typically break this cycle. We define the cycle consistency similarly for 2-frame optical flow (RAFT) and 8-frame tracking (PIPs).

2-Frame Cycle-Consistency Filter. For optical flow, we compute consistency by warping the backward flow map into alignment with the forward flow map and then measure per-pixel disagreements with a sum (since opposite flows should cancel out) [30]. Denoting the forward flow as w and the aligned backward flow as \hat{w} , \forall pixel locations i, j , we check:

$$\|w_{ij} + \hat{w}_{ij}\|_2^2 < \alpha(\|w_{ij}\|_2^2 + \|\hat{w}_{ij}\|_2^2 + \beta), \quad (1)$$

which linearly scales the tolerance threshold according to the motion magnitude. Flows that satisfy Eqn. 1 are considered cycle-consistent.

Following the classical optical flow assumption of color constancy [17, 6, 19], we also ask for the tracks to be color-consistent. Denoting the zeroth frame as r and the flow-aligned first frame as \hat{r} , we use:

$$\|r - \hat{r}\|_2^2 < \gamma, \quad (2)$$

where the image colors are normalized to the range $[-0.5, 0.5]$. This filter demands that corresponding pixels should have similar colors. Only tracks that are both color and cycle-consistent will be kept by the filter.

8-Frame Cycle-Consistency Filter. For 8-frame sparse motion tracking, we measure the maximum per-timestep distance between the forward trajectory v and the time-flipped backward trajectory \hat{v} :

$$\max_t \|v_t - \hat{v}_t\|_2^2 < \tau, \quad (3)$$

where a threshold of $\tau = 1$ means that the trajectories may only drift a maximum of 1 pixel apart. Tracks that satisfy Eqn. 3 are considered cycle-consistent.

For both 2-frame and 8-frame tracking, the threshold values (τ , α , and β) depend on meta information such as the spatial resolution of the video. We measure the effect of these thresholds in our ablation study, and find that there is a Goldilocks zone: with thresholds set too small, yielding a strict filter, there are too few good tracks for the pre-trained model to finetune; with thresholds set too large, yielding a loose filter, there are too many noisy tracks which corrupt the pre-trained weights.

We run the motion model with temporal stride 2 (to reduce data redundancy), and use our cycle-consistency filters to select a small subset (approximately 10%) of the resulting tracks. We use this subset as pseudo-labels for the next stage.

C. Second Stage: Motion Refinement

In the second stage, we try to improve the pre-trained motion model by fine-tuning it on the batched and augmented pseudo labels, challenging the model to reproduce those motion estimates under more difficult settings. It may seem strange to fine-tune the model using its own estimates, but there are several reasons this may be helpful. First, the




Fig. 1: **Refining a pre-trained motion model.** We apply a pre-trained motion model on an unlabelled video, yielding a dense set of tracks. We filter this down to a sparse set of cycle-consistent tracks, creating pseudo-labels. We then train on the pseudo-labelled data using augmentations, to improve the model.

model is iterative (in the case of both RAFT and PIPs), and we are fine-tuning *all* iterations with the model’s *final* iteration, making the early iterations more confident to jump toward the answer. Second, taking gradient steps on these pseudo-labels may lead the model to “forget” behavior which is inconsistent with the pseudo-labels, which is a form of domain adaptation. Third, we apply augmentations while fine-tuning, which widens the domain where the model should produce these cycle-consistent outputs. We use the default augmentations from RAFT and PIPs, which include color jittering, cropping, resizing, and random occlusions.

If the test domain is a single video, we simply fine-tune on it directly, as shown in Fig. 1. If the test domain has multiple videos, we find that the algorithm is robust to either mixing all video clips together, or fine-tuning in an online fashion, continuing the optimization for each video in sequence.

An important hyperparameter in our approach is the number finetuning iterations κ , which controls how long the model is trained on the pseudo-labels. In the ablation study we measure the performance at different settings of κ , and find that as κ goes beyond its optimal value, performance degrades slightly but does not collapse.

IV. EXPERIMENTS

We present experiments on three widely-used datasets: CroHD [15], Tap-Vid-DAVIS [31, 16], and MPI-Sintel [3]. We use Sintel to measure performance on 2-frame dense optical flow estimation, and use CroHD and Tap-Vid-DAVIS to measure performance on 8-frame sparse motion tracking. We use RAFT [13] for optical flow, and PIPs [14] for multi-frame tracking. We use the publicly released checkpoints from the authors. Both models are (pre-)trained on the synthetic FlyingThings [2] dataset, which is noticeably unrealistic but produces strong models. In addition to evaluating the effect of our fine-tuning scheme, we evaluate alternative optimization objectives and also additional supervised and self-supervised baselines, contextualizing our results with existing works.

A. Datasets

In this section we describe each dataset in detail.

CroHD [15]. This dataset focuses on pedestrian head-tracking in long videos of crowded environments. Ground truth labels are provided on the train-split of the videos, thus we perform our self-supervised refinement framework on

pre-trained PIPs [14] on these videos then validate using the ground truth labels. Videos from CroHD are very similar to each other, because all demonstrate pedestrians walking either indoors or outdoors. In this dataset, we fine-tune on multiple videos, continuing optimization from one video to the next.

Tap-Vid-DAVIS [16]. This is a point tracking dataset based on the DAVIS 2017 validation set [31]. The dataset contains 30 challenging videos featuring different scenes and objects ranging from people bicycling to horse jumping captured with different camera motion. Each video is quite short: around 5 seconds on average. Ground truth labels are provided for approximately 26 points per video. We refine PIPs [14] on these videos then validate using the sparse ground truth. In this dataset we fine-tune on each video individually, since each video is unique.

MPI-Sintel [3]. This dataset consists of 23 animated scenes, with per-pixel ground truth flow. The videos vary widely in speed of motion, occlusion ratio, and scene contents. Two versions of the data are provided (‘clean’ and ‘final’), which vary in postprocessing effects (e.g., motion blur). In this data we refine pre-trained RAFT [13], and validate using the ground truth. In this dataset we fine-tune on each video individually, since each video is unique.

B. Evaluation Metrics

With three different datasets and tracking tasks, we leverage the most suitable evaluation metric to gauge the performance for each task focusing on different aspects.

CroHD [15]. For CroHD, we used five different evaluation metrics to comprehensively evaluate the model performance. ATE-Vis and ATE-Occ evaluate the average track error in pixels each time with an 8-frames video, for targets that stay fully-visible and ones that undergo occlusion, respectively. The other three metrics MTE, δ , and Survival Rate, focus on longer-frame tracking. MTE denotes the median track error in pixels; δ is an accuracy metric measuring the ratio of points in our predictions that are within a threshold distance of ground truth, *averaged* over thresholds (1, 2, 4, 8, 16). Survival measures the average time that a track manages to stay “on target” (within a threshold distance 16) without falling off the trajectory and switching to other object or points, measured as a ratio of video length [5].

Tap-Vid-DAVIS [16]. For Tap-Vid-DAVIS, we use the two evaluation metrics from the benchmark [16]: *delta*, measuring tracking position accuracy (as done in CroHD), and Average Jaccard (AJ), which is a metric that incorporates both position accuracy and occlusion accuracy.

MPI-Sintel [3]. For MPI-Sintel, we use the widely-adopted end-point error metric (EPE), which measures the average ℓ_2 error of the flow estimates.

C. Baselines

In our experiments, we consider the following baselines.

DINO [32]. This is a vision transformer (ViT-S [33] with patch size 8) trained with contrastive self-supervision setup on ImageNet [24]. We use the authors’ publicly released code for

TABLE I: **Evaluation in CroHD.** Ours clearly surpasses baselines in most of the metrics and improves ATE of pre-trained PIPs[14] by 10%. Existing self-supervision methods make performance worse.

Method	ATE-Vis. ↓	ATE-Occ. ↓	MTE ↓	δ (%) ↑	Survival (%) ↑
Supervised Out of Domain					
RAFT [13]	8.04	13.21	82.76	15.82	62.22
DINO [32]	23.04	25.91	116.80	8.46	37.11
PIPs [14]	4.57	7.71	8.35	50.89	82.37
Supervised Out of Domain + Self-Supervised					
PIPs+Color	4.57	7.75	8.36	50.75	82.55
PIPs+Feat.	4.52	7.70	8.37	49.76	82.90
PIPs+Ours	4.02	6.79	7.93	49.88	83.96

TABLE II: **Evaluation in Tap-Vid-DAVIS.** Ours clearly surpasses all baselines in both δ and AJ with the proposed self-supervised finetuning.

Method	δ (%) ↑	AJ↑
Supervised Out of Domain		
COTR [36]	51.3	35.4
RAFT [13]	46.3	79.6
PIPs [14]	59.4	42.0
TAP-Net [16]	53.1	38.4
Supervised Out of Domain + Self-Supervised		
PIPs+Ours	60.0	42.5

converting this model into a tracker, which works per-pixel feature matching, with the help of a memory pool [34].

Tap-Net [16]. This is a recent method designed specifically for multi-frame point tracking. It estimates pixel correspondences by taking the argmax of frame-by-frame cost maps computed efficiently using time-shifted convolutions [35].

Self-supervision with color constancy. This has been previously explored in methods that train from scratch [6, 19, 7, 8], but to our knowledge has not been applied to refine pre-trained supervised models. We consider this as a baseline in our work, by fine-tuning RAFT and PIPs with classic color constancy and edge-aware smoothness objectives.

D. Results

We present results of our method on each dataset. Overall, our results demonstrate that our method yields reliable gains over fully-supervised methods, and significantly surpasses the baselines.

CroHD [15]. We use PIPs [14] pre-trained on FlyingThings [2] dataset and run CroHD inference during sampling with a spatial resolution of (768, 1280). The 8-frame cycle consistency filter (Eqn. 3) and multiple-video fine-tuning (Sec. III-C) are adopted. We choose 2.5 for the threshold τ and 3000 for finetuning iterations κ .

Results and comparisons are provided in Table I. We compare with state-of-the-art methods supervised out-of-domain [13, 32, 14], and also methods that perform self-supervised finetuning on a pre-trained model like ours. Our method clearly surpasses all of the baselines with better scores on most of the metrics. Notably, compared with pre-trained




Fig. 2: **Evaluation in Tap-Vid-DAVIS.** Performance change in δ for each of the video compared with the baseline pre-trained PIPs model. *Positive value denotes improvement (i.e., accuracy increasing)*




Fig. 3: **Evaluation in MPI-Sintel.** Percent change in EPE for each video compared with the pre-trained RAFT [13] model. *Negative value denotes improvement (i.e., error decreasing).*

PIPs where our optimization begins, we reduce ATE-Vis and ATE-Occ from 4.57 and 7.71 to 4.02 and 6.79, improving by around 10% with our self-supervised refinement procedure.

For baselines, we also try several self-supervised techniques to finetune the pre-trained model as ours, namely the PIPs+Color and PIPs+Feat shown in Table I. The idea of them is to use a self-supervised loss measuring the difference between the color or the image features at each predicted track position. We could observe that both these two techniques are not effective, demonstrating even worse scores on many evaluation metrics than the pre-trained model. This further implies the effectiveness of our proposed procedure in finetuning a pre-trained motion model.

TAP-Vid-DAVIS [16]. We use PIPs [14] pre-trained on FlyingThings [2] dataset and run DAVIS inference during sampling with a spatial resolution of (288, 512). The 8-frame cycle consistency filter (Eqn. 3) and single-video finetuning (Sec. III-C) are adopted. Specifically, we choose 1 for the threshold τ and 100 for finetuning iterations κ for each video. For a single-video finetuning setup, we show performance change on each individual video with the baseline pre-trained model. To compare with other methods on the same scale, we also report an average of our performance over all videos.

Results and comparisons are provided in Table II and Fig. 2. We compare with state-of-the-art methods [36, 13, 14, 16]




Fig. 4: Comparison of optical flow visualizations on Sintel produced by the pre-trained RAFT [13] and RAFT with ours, with difference highlighted in red boxes. Ours refine the flows by cleaning noisy tracks and completing missing objects.

TABLE III: **Evaluation in MPI-Sintel.** A clear advantage is observed for our method, improving pre-trained RAFT by around 10%. Existing self-supervised methods do not help.

Method	EPE↓	
	Clean	Final
Supervised Out of Domain		
FlowNet2 [37]	2.02	3.14
PWC-Net [38]	2.55	3.93
VCN [39]	2.21	3.62
RAFT [13]	1.46	2.72
Self-Supervised		
SelFlow [7]	2.88	3.87
UFlow [19]	3.01	4.09
SMURF-test [8]	1.99	2.80
Supervised Out of Domain + Self-Supervised		
RAFT+Color	1.48	2.73
RAFT+Col.+Smooth	1.47	2.74
RAFT+Col.+Edge-Sm.	1.51	2.84
RAFT+Ours	1.32	2.46

supervised out of domain including the latest works like Tap-Net [16]. We beat all of the baselines with better δ and AJ. Notably, compared with PIPs where we start from, we improve δ and AJ from 59.4 and 42.0 to 60.0 and 42.5. Moreover, we also show the performance change in δ_{avg} for each video in Fig. 2. We could observe positive change through our self-supervised refinement on almost all of the 30 videos except two videos ‘parkour’ and ‘scooter-black’.

Compared with the CroHD results in Table I where we improve the pre-trained model by 10%, results on Tap-Vid-DAVIS do not show as much improvement. However, we note that each Tap-Vid-DAVIS video is only 5 seconds long on average (not much space to extract pseudo-labels), and ground truth is only available for a small number of points (less than 30 per video).

MPI-Sintel. For MPI-Sintel, we use RAFT [13] pre-trained on FlyingThings datasets and run Sintel inference during sampling with a spatial resolution of (436, 1024). We use the

2-frame cycle consistency filter (Eqn. 1) augmented by color consistency (Eqn. 2) and fine-tune on each video individually. We use the threshold hyperparameters $\alpha = 0.005$, $\beta = 0.25$, and $\gamma = 0.1$. We choose 100 for finetuning iterations κ for each video. We evaluate performance change on each individual scene, and also measure the average across videos.

Quantitative comparisons are provided in Table III and Fig. 3, and qualitative comparisons are in Fig. 4. We compare with state-of-the-art methods supervised out-of-domain, self-supervised methods, and self-supervised techniques to finetune the pre-trained model. We surpass all of the baselines, with better EPE on both clean and final splits of MPI-Sintel. Notably, compared with the pre-trained RAFT model, we reduce the EPE for both clean and final from 1.46 and 2.72 to 1.32 and 2.46, demonstrating around 10% improvement. Since evaluation variance is quite huge on Sintel, we visualize the percentage EPE change for each of the videos in Figure 3. We can clearly observe that our self-supervised refinement improves the model performance on all of the scenes except ‘ambush_6’ in the clean split. On the other scenes, we observe obvious and significant improvements, even 30% on scenes like ‘ambush_7’. We visualize the optical flow results in Fig. 4, and highlight the noticeably different regions in red boxes. We observe that our proposed method helps refine the optical flow output by completing missing objects, and improving smoothness in large regions.

As baselines to our optimization strategy, we try using color constancy as a self-supervised loss to finetune the pre-trained model, optionally combined with a smoothness regularizer or edge-aware smoothness regularizer. We list these in Table III as RAFT+Color, RAFT+Col.+Smooth and RAFT+Col.+Edge-Sm. In practice, we found that none of these techniques help, and EPE simply rises as training progresses.

E. Ablation Studies

We conduct two ablation studies to investigate the effect of the design choices and hyperparameters. First, we study the effect of the threshold τ and finetuning iterations κ . Results are shown in Fig. 5. We conduct this study on CroHD. We collect the best ATE_OCC and ATE_VIS for various τ s, and as an additional variant, randomly sample the tracks to choose the pseudo-label finetuning set, denoted as ‘No Threshold’ in the figure. We find that random sampling leads to almost no performance change, while the cycle-consistency-based filtering leads to improvements, across various values of τ . We observe that the optimal performance is achieved at τ of 2.5 in this dataset, with a U-shaped curve around this value. This follows expectations, since a loose filter (with large τ) includes too many noisy motion estimates, which corrupts the pre-trained weights; a strict filter (with small τ) contains too few good estimates to provide useful guidance to the pre-trained model for the new context.

Fig. 5-(c) and Fig. 5-(d) measure the effect of fine-tuning iterations κ . The plot illustrates that $\kappa = 3000$ achieves the best performance. Performance degrades slightly with either smaller or larger κ , which either insufficiently finetunes the model on the pseudo-label set, or leads to slight overfitting.

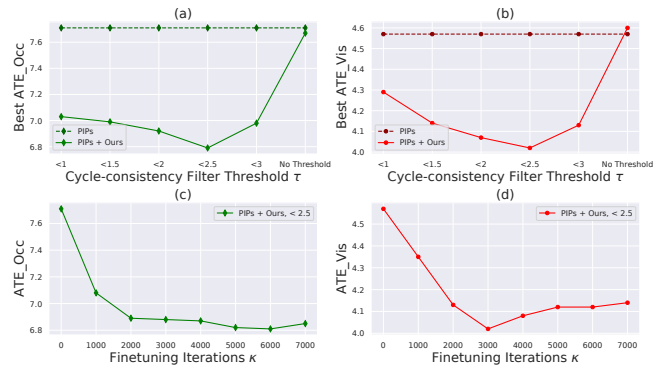


Fig. 5: **Ablation study in CroHD.** We show ATE_VIS and ATE_OCC for picking pseudo-labeled tracks with different threshold τ (a and b) and fine-tuned for different iterations κ (c and d). All metrics are the smaller the better. Optimal performance is observed at τ of 2.5 and κ of 3000.

TABLE IV: **Ablation study in MPI-Sintel-Clean.** Results are evaluated on 8 videos from the Sintel-Clean train split. Both Cycle Consistency and Color Constancy self-supervised metrics lead to improvements from the pre-trained RAFT. Combining both metrics works the best.

Method	Color Constancy	Cycle Consistency	EPE↓
RAFT	n/a	n/a	1.05
RAFT + Ours	✓	✗	0.98
RAFT + Ours	✗	✓	0.99
RAFT + Ours	✓	✓	0.93

Table IV measures the effectiveness of different strategies for choosing pseudo-labels for optical flow. We conduct this study in Sintel [3]. Both color constancy and cycle consistency alone are helpful, with cycle consistency slightly outperforming color constancy. Combining the two filtering strategies yields the best performance.

V. DISCUSSION AND CONCLUSION

In this paper, we present a procedure to improve state-of-the-art supervised motion models with only self-supervised fine-tuning. Our procedure is a two-stage method where we first use a pre-trained motion model to sample and filter cycle-consistent tracks in the test domain, and then we fine-tune on these tracks, with added augmentations. We compare our method with state-of-the-art supervised models, self-supervised training methods, and self-supervised finetuning techniques. We find our method improves the pre-trained model performance, on three different datasets, with two different tracking models, whereas existing self-supervised techniques worsen performance. We hope our method will inspire future work on refining pre-trained motion models.

VI. ACKNOWLEDGMENTS.

This work was supported by the Toyota Research Institute under the University 2.0 program, ARL grant W911NF-21-2-0104, and a Vannevar Bush Faculty Fellowship. The authors thank Leonid Keselman for suggesting the title.

REFERENCES

- [1] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. Van Der Smagt, D. Cremers, and T. Brox, "FlowNet: Learning optical flow with convolutional networks," in *ICCV*, 2015, pp. 2758–2766.
- [2] N. Mayer, E. Ilg, P. Häusser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox, "A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation," in *CVPR*, 2016.
- [3] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black, "A naturalistic open source movie for optical flow evaluation," in *ECCV*, 2012, pp. 611–625.
- [4] W. Wang, D. Zhu, X. Wang, Y. Hu, Y. Qiu, C. Wang, Y. Hu, A. Kapoor, and S. Scherer, "Tartanair: A dataset to push the limits of visual slam," in *2020 IEEE/RISJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 4909–4916.
- [5] Y. Zheng, A. W. Harley, B. Shen, G. Wetzstein, and L. J. Guibas, "Pointodyssey: A large-scale synthetic dataset for long-term point tracking," in *ICCV*, 2023.
- [6] J. J. Yu, A. W. Harley, and K. G. Derpanis, "Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness," in *ECCVW*. Springer, 2016, pp. 3–10.
- [7] P. Liu, M. Lyu, I. King, and J. Xu, "Selflow: Self-supervised learning of optical flow," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4571–4580.
- [8] A. Stone, D. Maurer, A. Ayvaci, A. Angelova, and R. Jonschkowski, "Smurf: Self-teaching multi-frame unsupervised raft with full-image warping," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 3887–3896.
- [9] N. Wang, Y. Song, C. Ma, W. Zhou, W. Liu, and H. Li, "Unsupervised deep tracking," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1308–1317.
- [10] X. Li, S. Liu, S. De Mello, X. Wang, J. Kautz, and M.-H. Yang, "Joint-task self-supervised learning for temporal correspondence," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [11] X. Wang, A. Jabri, and A. A. Efros, "Learning correspondence from the cycle-consistency of time," in *CVPR*, 2019.
- [12] A. Jabri, A. Owens, and A. A. Efros, "Space-time correspondence as a contrastive random walk," *Advances in Neural Information Processing Systems*, 2020.
- [13] Z. Teed and J. Deng, "RAFT: Recurrent all-pairs field transforms for optical flow," in *ECCV*, 2020.
- [14] A. W. Harley, Z. Fang, and K. Fragkiadaki, "Particle video revisited: Tracking through occlusions using point trajectories," in *ECCV*, 2022.
- [15] R. Sundararaman, C. De Almeida Braga, E. Marchand, and J. Pettre, "Tracking pedestrian heads in dense crowd," in *CVPR*, 2021, pp. 3865–3875.
- [16] C. Doersch, A. Gupta, L. Markeeva, A. Recasens, L. Smaira, Y. Aytar, J. Carreira, A. Zisserman, and Y. Yang, "Tap-vid: A benchmark for tracking any point in a video," 2022.
- [17] B. D. Lucas, T. Kanade, *et al.*, *An iterative image registration technique with an application to stereo vision*. Vancouver, 1981, vol. 81.
- [18] P. J. Burt and E. H. Adelson, "The laplacian pyramid as a compact image code," in *Readings in computer vision*. Elsevier, 1987, pp. 671–679.
- [19] R. Jonschkowski, A. Stone, J. T. Barron, A. Gordon, K. Konolige, and A. Angelova, "What matters in unsupervised optical flow," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*. Springer, 2020, pp. 557–572.
- [20] W. Im, S. Lee, and S.-E. Yoon, "Semi-supervised learning of optical flow by flow supervisor," in *ECCV*. Springer, 2022, pp. 302–318.
- [21] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [22] Y. Tang, Z. Jiang, Z. Xie, Y. Cao, Z. Zhang, P. H. Torr, and H. Hu, "Breaking shortcut: Exploring fully convolutional cycle-consistency for video correspondence learning," *arXiv preprint arXiv:2105.05838*, 2021.
- [23] Z. Bian, A. Jabri, A. A. Efros, and A. Owens, "Learning pixel trajectories with multiscale contrastive random walks," *CVPR*, 2022.
- [24] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *CVPR09*, 2009.
- [25] W.-S. Lai, J.-B. Huang, and M.-H. Yang, "Semi-supervised learning for optical flow with generative adversarial networks," *Advances in neural information processing systems*, vol. 30, 2017.
- [26] J. Jeong, J. M. Lin, F. Porikli, and N. Kwak, "Imposing consistency for optical flow estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 3181–3191.
- [27] D.-H. Lee *et al.*, "Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks," in *Workshop on challenges in representation learning, ICML*, vol. 3, no. 2, 2013, p. 896.
- [28] A. Carlson, J. Betteridge, B. Kisiel, B. Settles, E. R. Hruschka, and T. M. Mitchell, "Toward an architecture for never-ending language learning," in *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, ser. AAAI'10. AAAI Press, 2010, p. 1306–1313.
- [29] X. Chen, A. Shrivastava, and A. Gupta, "NEIL: Extracting visual knowledge from web data," in *ICCV*, 2013, pp. 1409–1416.
- [30] N. Sundaram, T. Brox, and K. Keutzer, "Dense point trajectories by GPU-accelerated large displacement optical flow," in *ECCV*, 2010.
- [31] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool, "The 2017 DAVIS challenge on video object segmentation," *arXiv:1704.00675*, 2017.
- [32] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," in *ICCV*, 2021.
- [33] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *ICLR*, 2021.
- [34] Z. Lai, E. Lu, and W. Xie, "MAST: A memory-augmented self-supervised tracker," in *CVPR*, 2020.
- [35] J. Lin, C. Gan, and S. Han, "Tsm: Temporal shift module for efficient video understanding," in *ICCV*, 2019, pp. 7083–7093.
- [36] W. Jiang, E. Trulls, J. Hosang, A. Tagliasacchi, and K. M. Yi, "COTR: Correspondence Transformer for Matching Across Images," in *ICCV*, 2021.
- [37] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox, "FlowNet 2.0: Evolution of optical flow estimation with deep networks," in *CVPR*, 2017.
- [38] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, "PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume," in *CVPR*, 2018.
- [39] G. Yang and D. Ramanan, "Volumetric correspondence networks for optical flow," *Advances in neural information processing systems*, vol. 32, 2019.