

## Learning Pixel Trajectories with Multiscale Contrastive Random Walks

Zhangxing Bian<sup>1,3</sup>   Allan Jabri<sup>2</sup>   Alexei A. Efros<sup>2</sup>   Andrew Owens<sup>1</sup>  
 University of Michigan<sup>1</sup>   UC Berkeley<sup>2</sup>   Johns Hopkins University<sup>3</sup>

### Abstract

*A range of video modeling tasks, from optical flow to multiple object tracking, share the same fundamental challenge: establishing space-time correspondence. Yet, approaches that dominate each space differ. We take a step towards bridging this gap by extending the recent contrastive random walk formulation to much denser, pixel-level space-time graphs. The main contribution is introducing hierarchy into the search problem by computing the transition matrix between two frames in a coarse-to-fine manner, forming a multiscale contrastive random walk when extended in time. This establishes a unified technique for self-supervised learning of optical flow, keypoint tracking, and video object segmentation. Experiments demonstrate that, for each of these tasks, the unified model achieves performance competitive with strong self-supervised approaches specific to that task.<sup>1</sup>*

### 1. Introduction

Temporal correspondence underlies a range of video understanding tasks, from optical flow to object tracking. At the core, the challenge is to estimate the motion of some entity as it persists in the world, by searching in space and time. For historical reasons, the practicalities differ substantially across tasks: optical flow aims for dense correspondences but only between neighboring pairs of frames, whereas tracking cares about longer-range correspondences but is spatially sparse. We argue that the time might be right to try and re-unify these different takes on temporal correspondence.

An emerging line of work in self-supervised learning has shown that generic representations pretrained on unlabeled images and video can lead to strong performance across a range of tracking tasks [20, 26, 35, 75, 77]. The key idea is that if tracking can be formulated as label propagation [86] on a space-time graph, all that is needed is a good measure of similarity between nodes. Indeed, the recent *contrastive random walk* (CRW) formulation [26] shows how such a similarity measure can be learned for temporal cor-

respondence problems, suggesting a path towards a unified solution. However, scaling this perspective to pixel-level space-time graphs holds challenges. Since computing similarity between frames is quadratic in the number of nodes, estimating dense motion is prohibitively expensive. Moreover, there is no way of explicitly estimating the motion in ambiguous cases, like occlusion. In parallel, the unsupervised optical flow community has adopted highly effective methods for dense matching [83], which use *multiscale* representations [9, 45, 61, 84] to reduce the large search space, and smoothness priors to deal with ambiguity and occlusion. But, in contrast to the self-supervised tracking methods, they rely on hand-crafted distance functions, such as the Census Transform [32, 47]. Furthermore, because they focus on producing point estimates of motion, they may be less robust under long-term dynamics.

In this work, we take a step toward bridging the gap between tracking and optical flow by extending the contrastive random walk formulation [26] to much denser, pixel-level space-time graphs. The main contribution is introducing hierarchy into the search problem, i.e. the *multiscale* contrastive random walk. By integrating local attention in a coarse-to-fine manner, the model can efficiently consider a distribution over pixel-level trajectories. Through experiments across optical flow and video label propagation benchmarks, we show:

- This provides a unified technique for self-supervised optical flow, pose tracking, and video object segmentation.
- For optical flow, the model is competitive with many recent unsupervised optical flow methods, despite using a novel loss function (without hand-crafted features).
- For tracking, the model outperforms existing self-supervised approaches on pose tracking, and is competitive on video object segmentation.
- Cycle-consistency provides a complementary learning signal to photo-consistency.
- Multi-frame training improves two-frame performance.

### 2. Related work

**Space-time representation learning.** Recent work has proposed methods for tracking objects in video through

<sup>1</sup>Project page at <https://jasonbian97.github.io/flowwalk>

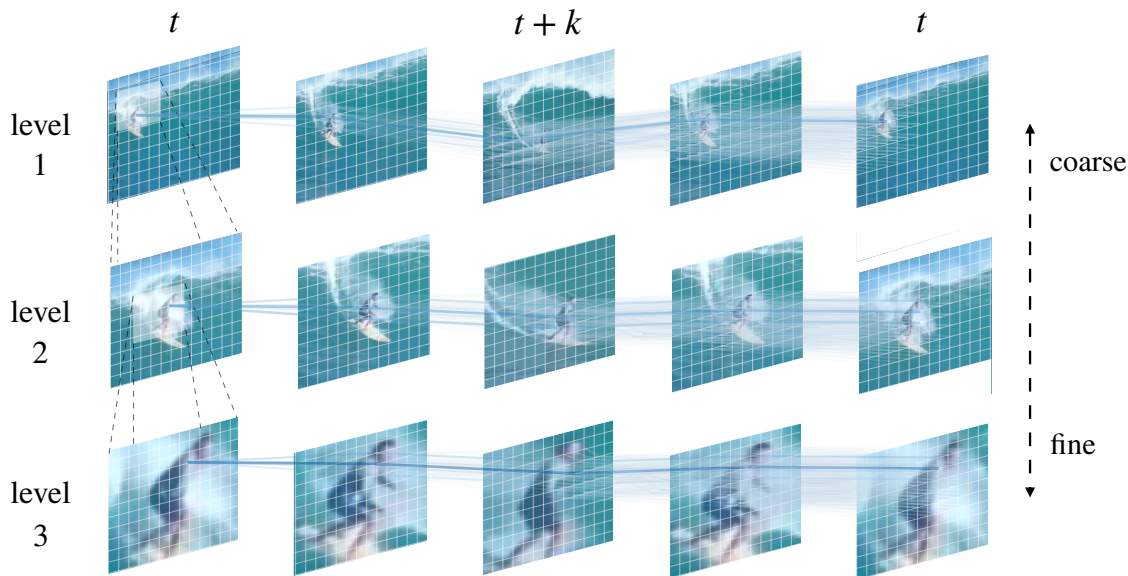


Figure 1. **Multiscale contrastive random walks.** We learn representations for dense, fine-grained matching using *multiscale contrastive random walks*. At each scale, we create a space-time graph in which each pixel is a node, and where nodes close in space-time are connected. Transition probabilities are based on similarity in a learned representation. We train the network to maximize the probability of obtaining a cycle-consistent random walk. Here, we illustrate a pixel’s walk over 3 spatial scales and 3 frames. The walker is initialized using its position at the previous, coarser scale.

self-supervised representation learning. Vondrick *et al.* [71] posed tracking as a colorization problem, by training a network to match pixels in grayscale images that have the same held-out colors. The assumption that matching pixels have the same color may break down over long time horizons, and the method is limited to grayscale inputs. This approach was extended to obtain higher-accuracy matches with two-stage matching [35,38]. In contrast to our approach, the predictions are coarse, patch-level associations. Another line of work learns representations by maximizing cycle consistency. These methods track patches forward, then backward, in time and test whether they end up where they began. Wang *et al.* [74] proposed a method based on hard attention and spatial transformers [27]. Jabri *et al.* [26] posed cycle-consistency as a random walk problem, allowing the model to obtain dense supervision from multi-frame video. Tang *et al.* [67] proposed an extension that allowed for fully convolutional training. These approaches are trained on sparse patches and learn coarse-grained correspondences. In contrast, we learn pixel-to-pixel correspondences. Other work has encouraged patches in the same position in neighboring frames to be close in an embedding space [17,79].

**Optimization-based optical flow.** Lucas and Kanade [2,45] used Gauss-Newton optimization to minimize a brightness constancy objective. In their seminal work, Horn and Schunck [21] combined brightness constancy with a spatial smoothness criteria, and estimated flow using variational methods. Later research improved flow estimation using robust penalties [4,6,40,61,62], coarse-to-fine esti-

mation [8], discrete optimization [13,55,58], feature matching [5,53,78], bilateral filtering [1], and segmentation [87]. In contrast, our model estimates flow using neural networks.

**Unsupervised optical flow.** Early work [49] used Boltzmann machines to learn transformations between images. More recently, Yu *et al.* [83] train a neural network to minimize a loss very similar to that of optimization-based approaches. Later work extended this approach by adding edge-aware smoothing [76], hand-crafted features [48], occlusion handling [28,76], learned upsampling [46], and depth and camera pose [82,85]. Another approach learns flow by matching augmented image pairs [41,42,44]. Recently, Jonschkowski *et al.* [32] surveyed previous literature and conducted an exhaustive search to find the best combination of methods and hyperparameters. In contrast to these works, our goal is to learn self-supervised representations for matching, in lieu of hand-crafted features. Moreover, we aim to produce a distribution over motion trajectories suitable for label transfer, rather than motion estimates alone. In very recent work, Stone *et al.* [60] adapted unsupervised flow methods to the RAFT architecture [68] and proposed new augmentation, self-distillation, and multi-frame occlusion inpainting methods. By contrast, we use PWC-net [65], since it is the standard architecture considered in prior work, and since it can obtain strong performance with careful training [64].

**Cycle consistency.** Cycle consistency has long been used to detect occlusions [3,23,36,66,80], and is used to discount the loss of occluded pixels in unsupervised flow [25,28,32,

76]. Zou *et al.* [88] used a cycle consistency loss as part of a system that jointly estimated depth, pose, and flow. Recently, Huang *et al.* [22] combined cycle-consistency with epipolar matching, but their method is weakly supervised with camera pose and assumes egomotion. In contrast, ours is entirely unsupervised and is capable of working *solely* with cycle-consistency and smoothness losses. Without extra constraints, their cycle consistency formulations have trivial solutions (*e.g.*, all-zero flow). Other recent work [69] learns to match images by ensuring that both an image and a warped variation of it match consistently with a second image, and Li *et al.* [37] used random walks with fixed transition matrices to smooth scene flow on point clouds. A random walk formulation of cycle consistency also been used in semi-supervised learning [18], using labels to test consistency.

**Multi-frame matching.** Many methods use a third frame to obtain more local evidence for matching. Classic methods assume approximately constant velocity [30, 63, 70, 72] or acceleration [33, 72] and measure the photo-consistency over the full set of frames. Recently, Janai *et al.* [28] proposed an unsupervised multi-frame flow method that used a photometric loss with a low-acceleration assumption and explicit occlusion handling. In contrast to these approaches, our method can be deployed using two frames at test time. We use subsequent frames as a *training signal*. There have also been a variety of approaches that track over long time horizons, often using sparse (or semi-dense) keypoints [56, 57]. Other work chains together optical flow [7, 66, 73], typically after removing low-texture regions. In contrast, our method also learns “soft” per-pixel tracks, which convey the probability that pairs of pixels match.

**Supervised optical flow.** Early work learned optical flow with probabilistic models, such as graphical models [15]. Other work learns parameters for smoothness and brightness constancy [62] or robust penalties [4, 39]. More recent methods has used neural networks. Fischer *et al.* [12, 24] proposed architectures with a built-in correlation layers. Sun *et al.* [65] introduced a network with built-in coarse-to-fine matching. Recent work [68] iteratively updates a flow with multiscale features, in lieu of coarse-to-fine matching.

### 3. Method

We first show how to learn dense space-time correspondences using mutiscale contrastive random walks, resulting in a model that obtains high quality motion estimates via simple nonparametric matching. We then describe how the learned representation can be combined with regression to handle occlusions and ambiguity, for improved optical flow.

#### 3.1. Multiscale contrastive random walks

We review the single-scale contrastive random walk, then extend the approach to multiscale estimation.

##### 3.1.1 Preliminaries: Contrastive random walks

We build on the contrastive random walk (CRW) formulation of Jabri *et al.* [26]. Given an input video with  $k$  frames, we extract  $n$  patches from each frame and assign each an embedding using a learned encoder  $\phi$ . These patches form the vertices of a graph that connects all patches in temporally adjacent frames. A random walker steps through the graph, forward in time from frames  $1, 2, \dots, k$ , then backward in time from  $k - 1, k - 2, \dots, 1$ . The transition probabilities are determined by the similarity of the learned representations:

$$A_{s,t} = \text{softmax}(X_s X_t^\top / \tau), \quad (1)$$

for a pair of frames  $s$  and  $t$ , where  $X_i \in \mathbb{R}^{n \times d}$  is the matrix of  $d$ -dimensional embedding vectors,  $\tau$  is a small constant, and the softmax is performed along each row. We train the model to maximize the likelihood of *cycle consistency*, *i.e.*, the event that the walker returns to the node it started from:

$$\mathcal{L}_{\text{CRW}} = -\frac{1}{n} \text{tr}(\log(\bar{A}_{t,t+k} \bar{A}_{t+k,t})), \quad (2)$$

where the log is elementwise and  $\bar{A}_{t,t+k}$  are the transition probabilities from frame  $t$  to  $t + k$ :  $\bar{A}_{t,t+k} = \prod_{i=t}^{t+k-1} A_{i,i+1}$ .

##### 3.1.2 Optical flow as a random walk

After training, the transition matrix contains the probability that a pair of patches is in space-time correspondence. We can estimate the optical flow  $\mathbf{f}_{s,t} \in \mathbb{R}^{n \times 2}$  of a patch between frames  $s$  and  $t$  by taking the expected value of the change in spatial position:

$$g_{\text{avg}}(A_{s,t}) = \mathbb{E}_{A_{s,t}}[\mathbf{f}_{s,t}] = A_{s,t} D - D, \quad (3)$$

where  $D \in \mathbb{R}^{n \times 2}$  is the (constant) matrix containing of pixel coordinates for each patch, and  $A_{s,t} D$  is the walker’s expected position in frame  $t$ .

In contrast to widely-used forward-backward cycle consistency formulations [22, 88], which measure the deviation of a predicted motion from a starting point, there is no trivial solution (*e.g.*, all-zero flow). This is because cycle consistency is measured in an embedding space defined solely from the visual content of image regions.

##### 3.1.3 Multiscale random walk

As presented so far, this formulation is expensive to scale to high-resolutions because computing the transition matrix is quadratic in the number of nodes. We overcome this by introducing hierarchy into the search problem. Instead of comparing all pairs, we only attend on a local neighborhood. By integrating local search across scales in a coarse-to-fine manner, the model can efficiently consider a distribution over pixel-level trajectories.

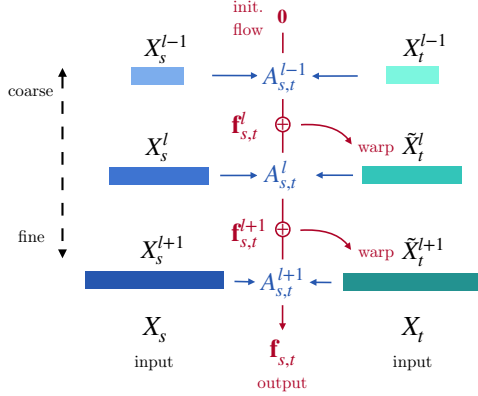


Figure 2. **Coarse-to-fine matching.** Our model performs a contrastive random walk across spatial scales: it computes a transition matrix (Eq. 5), uses it to obtain flow (Eq. 6), and recurses to the next level by using upsampled estimated flow to align the finer scale for matching (Eq. 4). `warp` samples the grid using the flow.

**Coarse-to-fine local attention.** Computing the transition matrix closely resembles cost volume estimation in optical flow [6, 14, 61]. This inspires us to draw on the classic spatial pyramid commonly used for multiscale search in optical flow, by iteratively computing the dense transition matrix, from coarse to fine spatial scales  $l \in [1..L]$ .

For frames  $s$  and  $t$ , we compute feature pyramids  $X_s^l \in \mathbb{R}^{h^{(l)}w^{(l)} \times d}$ , where  $h^{(l)}$  and  $w^{(l)}$  are the width and height of the feature map at scale  $l$ . To match each level efficiently, we `warp` the features of the target frame  $X_t^l$  into the coordinate frame of  $X_s^l$  using the coarse flow from the previous level  $f_{s,t}^{l-1}$ ; we then compute *local* transition probabilities on the warped feature to account for remaining motion. Thus, we estimate the transition matrix and flow in a *coarse-to-fine* manner (Fig. 2), computing levels recurrently:

$$\tilde{X}_t^l = \text{warp}(X_t^l, f_{s,t}^{l-1}) \quad (4)$$

$$A_{s,t}^l = \text{masked\_softmax}(X_s^l \tilde{X}_t^{l\top} / \tau) \quad (5)$$

$$\mathbf{f}_{s,t}^{l+1} = \text{upsample}(g_{\text{avg}}(A_{s,t}^l) + \mathbf{f}_{s,t}^l), \quad (6)$$

where `warp`( $X, \mathbf{f}$ ) samples features  $X$  with flow  $\mathbf{f}$  using bilinear sampling, and  $\mathbf{f}^1 = \mathbf{0}$ . For notational convenience, we write the local transition constraint using `masked_softmax`, which sets values beyond a local spatial window to zero. In practice, we use optimized correlation filtering kernels to compute Eq. 5 efficiently, and represent the transition matrices  $A_{s,t}^l$  as a sparse matrix.

**Loss.** After computing the transition matrices between all pairs of adjacent frames, we sum the contrastive random walk loss over all levels:

$$\mathcal{L}_{\text{msCRW}} = - \sum_{l=1}^L \frac{1}{n_l} \text{tr}(\log(\bar{A}_{t,t+k}^l \bar{A}_{t+k,t}^l)), \quad (7)$$

where  $\bar{A}_{s,t}^l$  is defined as in Eq. 2 and  $n_l$  is the number of nodes in level  $l$ . In our experiments, we use  $L = 5$  scales

and consider  $k \in [2..4]$  length cycles.

### 3.1.4 Smooth random walks

Since natural motions tend to be smooth [54], we follow work in optical flow [61], and incorporate smoothness as an additional desiderata for our random walks. We use the edge-aware loss of Jonschkowski *et al.* [32], which penalizes spatial changes in flow near similarly-colored pixels:

$$\mathcal{L}_{\text{smooth}} = \mathbb{E}_p \sum_{d \in \{x,y\}} \exp(-\lambda_c I_d(p)) \left| \frac{\partial^2 \mathbf{f}_{s,t}(p)}{\partial d^2} \right| \quad (8)$$

where  $I_d(p) = \frac{1}{3} \sum_c \left| \frac{\partial I_c}{\partial d} \right|$  is the spatial derivative averaged over all color channels  $I_c$  in direction  $d$ . The parameter  $\lambda_c$  controls the influence of similarly colored pixels. We apply this loss to each scale of the model.

## 3.2. Handling occlusion

While effective for most image content, nonparametric matching has no mechanism for estimating the motion of pixels that become occluded, since it requires a corresponding patch in the next frame. We propose a variation of the model that combines the multiscale contrastive random walk with a regression module that directly predicts the flow values at each pixel.

The architecture of our regressor closely follows the refinement module of PWC-net [65]. We learn a function  $g_{\text{reg}}(\cdot)$  that regresses the flow at each pixel from the multiscale contrastive random walk cost-volume and convolutional features. These features are obtained from the same shared backbone that is used to compute the embeddings. We provide a diagram in the supplement.

**Regression loss.** We train the regressor using a loss that closely resembles the contrastive random walk objective. Under this loss, pixels that are already well-matched by the nonparametric model (*e.g.*, non-occluded pixels) will be unlikely to change their flow values, while poorly-matched pixels (*e.g.*, occluded pixels) will obtain their flow values using a smoothness prior. We use the same smoothness loss,  $\mathcal{L}_{\text{smooth}}$ , as the nonparametric model (Eq. 8), and penalize the model from deviating from the nonparametric flow estimate (Eq. 3). We also use our learned embeddings as features for a *learned* photometric loss, *i.e.*, incurring loss if the model puts two pixels with dissimilar embedding vectors into correspondence. This results in an additional loss:

$$\mathcal{L}_{\text{feat}} = \|X_s - \text{warp}(X_t, \mathbf{f}_{s,t})\|^2 + \lambda_a \|\mathbf{f}_{s,t} - g_{\text{avg}}(A_{s,t})\|^2, \quad (9)$$

where  $\mathbf{f}_{s,t}$  is the predicted flow and  $\lambda_a$  is a constant. To prevent the regression loss from influencing the learned embeddings that it is based on, we do not propagate gradients from the regressor to the embeddings  $X$  during training. As in Eq. 7, we apply the loss to each scale and sum.

**Augmentation and masking.** We follow [32] and improve the regressor’s handling of occlusions through augmentation, and by discounting the loss of pixels that fail a consistency check [42]. To handle pixels that move off-screen, we compute flow, then randomly crop the input images and compute it again, penalizing deviations between the flows. This results in a new loss,  $\mathcal{L}_{\text{bound}}$ . We also remove the contribution of pixels in the photometric loss (Eq. 9) if they have no correspondence in the backwards-in-time flow estimation from  $t$  to  $s$ . Both are implemented exactly as in [32]. These losses apply only to the regressor, and thus do not directly affect the nonparametric matching.

### 3.3. Training

**Objective.** The pure nonparametric model (Section 3.1) can be trained by simply minimizing the multiscale contrastive random walk loss with a smoothness penalty:

$$\mathcal{L}_{\text{non}} = \mathcal{L}_{\text{msCRW}} + \mathcal{L}_{\text{smooth}}. \quad (10)$$

Adding the regressor results in the following loss:

$$\mathcal{L}_{\text{reg}} = \mathcal{L}_{\text{msCRW}} + \mathcal{L}_{\text{smooth}} + \mathcal{L}_{\text{feat}} + \mathcal{L}_{\text{bound}}. \quad (11)$$

We include weighting factors to control the relative importance of each loss (specified in the supplement).

**Architecture.** To provide a straightforward comparison with unsupervised optical flow methods, we use the PWC-net architecture [65] as our network backbone, after reducing the number of filters [32]. This network uses the feature hierarchy from a convolutional network to provide the features at each scale. We use the cost volume features from this network as the embedding for the random walk,  $X_s^l$ , after performing  $\ell_2$  normalization. We also use its regressor architecture. We provide architecture details in the supplementary material.

**Subcycles.** We follow [26] and include subcycles in our contrastive random walks: when training the model on  $k$ -frame videos, we include losses for walks of length  $k$ ,  $k-1$ , ... 2. These losses can be estimated efficiently by reusing the transition matrices for the full walk.

**Multi-frame training.** When training with  $k > 2$  frames, we use curriculum learning to speed up and stabilize training. We train the model to convergence with 2, 3, ...  $k$  frame cycles in succession.

**Optimization.** To implement the contrastive random walk, we exploit the sparsity of our coarse-to-fine formulation, and represent the transition matrices  $A_{s,t}^l$  as sparse matrices. This significantly improved training times and reduced memory requirements, especially in the finest scales. It takes approximately 3 days to train the full model on one GTX2080 Ti. We train our network with PyTorch [50], using the Adam [34] optimizer with a cyclic learning rate schedule [59] with a base learning rate of  $10^{-4}$  and a max learning rate of  $5 \times 10^{-4}$ . We provide training hyperparameters in supplementary material.

**Avoiding shortcuts.** The contrastive random walk can potentially obtain shortcut solutions when it is trained with a fully convolutional network by exploiting positional information [26]. While recent work has shown this can be solved through augmentation [67], we found that we avoided trivial shortcuts when using reflection padding in our network (for all convolution layers except for in the regressor). This may be because we simultaneously optimize multiple losses and use a limited search window, making the trivial solution harder to find.

## 4. Results

Our model produces two outputs: the optical flow fields and the pixel trajectories (which are captured in the transition matrices). We evaluate these predictions on label transfer and motion estimation tasks. We compare them to space-time correspondence learning methods, and with unsupervised optical flow methods.

### 4.1. Datasets

For simple comparison with other methods, we train on standard optical flow datasets. We note that the training protocols used by unsupervised optical flow literature are not standardized. We therefore follow the evaluation setup of [32]. We pretrain models on unlabeled videos from the **Flying Chairs** dataset [11]. We then train on the **KITTI-2015** [16] multi-view extension and **Sintel** [10]. To evaluate our model’s ability to learn from internet video, we also trained the model on **YouTube-VOS** [81], without pretraining on any other datasets.

We also evaluate our model on standard label transfer tasks. The **JHMDB** benchmark [31] transfers 15 body parts to future frames over long time horizons. The **DAVIS** benchmark [52] transfers object masks.

### 4.2. Label propagation

We evaluate our learned model’s ability to perform video label propagation, a task widely studied in video representation learning work. The goal is to propagate a label map provided in an initial video frame, which might describe keypoints or object segments, to the rest of the video.

We follow Jabri *et al.* [26] and use our model’s probabilistic motion trajectories to guide label propagation. We infer the labels for each *target* frame  $t$  auto-regressively. For each previous *source* frame  $s$ , we have a predicted label map  $L_s \in \mathbb{R}^{n \times c}$ , where  $n$  is the number of pixels and  $c$  is the number classes. As in [26], we compute  $K_{t,s}^l$ , the matrix of weights for attention on the source frames, by keeping the top- $k$  logits in each row of  $\tilde{A}_{t,s}^l$ . We use  $K_{t,s}$  as an attention matrix for each label, *i.e.*,  $L_t = K_{t,s} L_s$ . Using several source frames as context allows for overcoming occlusion.

We use variations of our model that was trained on the unlabeled Sintel and YouTube-VOS datasets, and use the transition matrix and flow fields at the penultimate level of



Figure 3. **Propagating segments and pose along motion trajectories.** We show qualitative results for JHMDB pose (left) and object masks on DAVIS (right). For DAVIS scenes, we show examples of mask propagations (top) and *soft* propagated label distributions (bottom).

Table 1. **Segment and Pose Propagation** with the JHMDB and DAVIS benchmarks, respectively. †: chained flow baseline.

Method	Arch.	Pose (PCK)			Segments
		@0.05	@0.1	@0.2	J&F <sub>m</sub>
UVC [38]	ResNet18	–	58.6	79.6	59.5
CRW [26]	ResNet18	29.1	59.3	80.6	67.6
VFS [79]	ResNet50	–	60.9	80.7	68.8
UFlow [32]†	PWC-Net	24.1	51.3	72.1	42.0
RAFT [68]†	RAFT	30.2	55.6	76.0	46.1
Ours - Sintel	PWC-Net	38.0	63.1	81.4	57.1
Ours - VOS	PWC-Net	38.2	62.6	80.9	57.9

the pyramid, i.e. level 4. Since the transition matrix describes residual motion, we warp (i.e. with  $f_{t,s}^4$ ) each label map before querying. Finally, since the features in level 4 have only 16 channels, we stack features from levels 3 and 4 to obtain hypercolumns [19], before computing attention.

**Evaluation.** We compared our model to recent video representation learning work, including single-scale CRW [26], UVC [38], and the state-of-the-art VFS [79] (Tab. 1). We also report two baselines that chain optical flow: unsupervised UFlow [32] and supervised RAFT [68].

For pose propagation, we evaluate our model on JHMDB [31] and report the PCK metric, which measures the fraction of keypoints within various distances to the ground truth. Our approach outperforms existing self-supervised approaches on this benchmark, especially at the stringent PCK@0.05, which is typically not reported. Note that our approach uses a significantly smaller network. While our model improves on fine-grained matching, it still struggles with occlusions (like other methods), which tend to involve large motions and motion blur (see Fig. 3 left). For object propagation, we evaluate our model on the DAVIS [52] benchmark, and report the mean of  $\mathcal{J}$  and  $\mathcal{F}$  metrics [51], which characterize segment overlap and boundary precision, respectively. Despite our focus on scaling to fine-grained matching, the model achieves competitive performance on DAVIS, outperforming the cycle-consistency method TimeCycle [75]. In attention visualizations of the propagated label distribution, we see that the transition distribution is robust to momentary occlusion (Fig. 3 mid bottom), but can nevertheless suffer from drift (Fig. 3 right bot-

tom). Interestingly, our model significantly outperforms the two optical flow methods, suggesting that the “soft” motion trajectories provided by our model may convey useful information for propagation that is not captured by the flow.

### 4.3. Optical flow

We evaluate our model’s ability to predict optical flow.

#### 4.3.1 Nonparametric motion estimation

Our model is able to estimate motion solely through nonparametric matching, as can be seen qualitatively in Figure 4. Despite the model’s simplicity and the fact that it is based on very different principles than existing flow methods, it obtains strong performance on matching non-occluded pixels (Tab. 2). It outperforms many unsupervised optical flow models, such as SelfFlow [44] (Tab. 4) on KITTI *noc* metric (non-occluded endpoint error). We see that the full regression-based variation of the model obtains better results, particularly on the *all* metric.

To help understand the importance of our multiscale formulation, we compared to Jabri *et al.* [26] on flow benchmarks, using their publicly released model (Tab. 2). This model resembles our nonparametric model, but with the random walk occurring at a single scale, with no smoothness prior. We evaluate the model using dense features, as in their approach. We found that our model significantly outperforms it. To control for training differences, we also tried removing scales from multiscale training (Tab. 6), finding that this significantly reduced performance.

#### 4.3.2 Effects of multi-frame cycles

We asked how the quality of the representation changes as we vary the number of frames  $k$  used to train the random walk. We test all models on 2-frame optical flow. As seen in Table 6, the model obtains better performance on all metrics using 3-frame and 4-frame cycles.

#### 4.3.3 Photometric feature learning

In contrast to other effective unsupervised optical flow approaches, our model does not use hand-crafted features. We evaluate the quality of our learned features when used as a photometric loss, compared to other common designs (Tab. 5). First, we compared our model to a variation of our

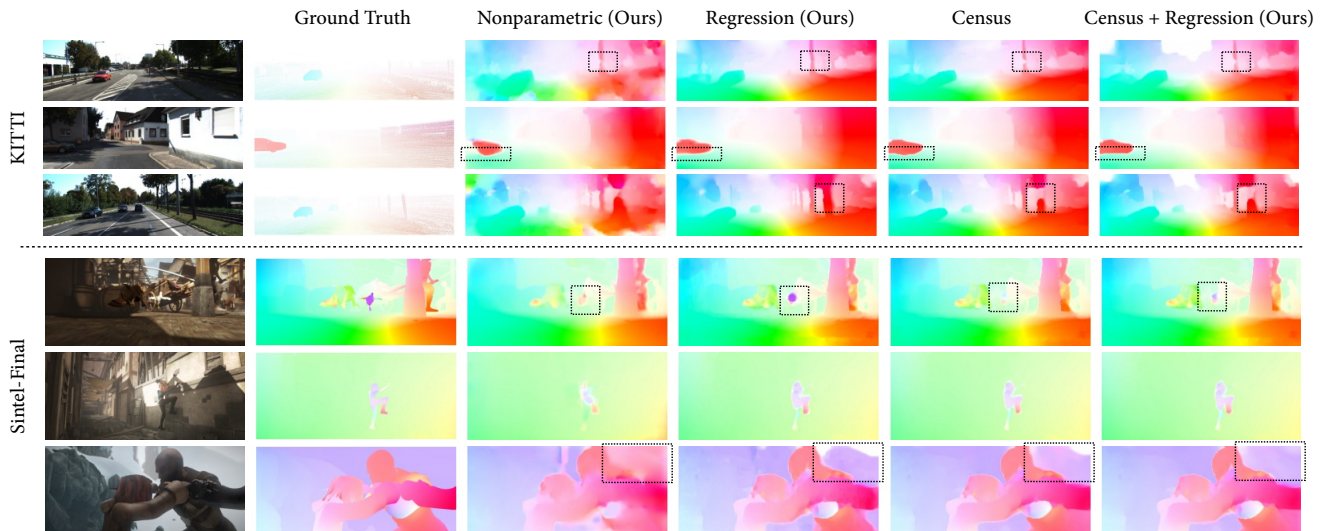


Figure 4. **Optical flow qualitative results.** We show results on images not seen during training, using our nonparametric-only and regression-based models. The highlighted regions show significant differences between the regression-based models. The optical flow vectors are coded by colors.

model that uses raw pixels (rather than hand-crafted features) within a photometric loss, with a robust **Charbonnier** penalty [61]. This baseline model is very similar to the *Charbonnier* variation of UFlow [32], though to control for other differences we use our own network. This amounts to simply replacing  $\mathcal{L}_{\text{feat}}$  with the Charbonnier loss (which disables the contrastive random walk). We found that the resulting model performed significantly worse.

Next, we considered using a state-of-the-art hand-crafted feature, the **Census transform** [47], resulting in a model similar to the *Census* variation of UFlow. We found that our features obtained competitive performance on non-occluded pixels, but that there was a significant advantage to Census features on the *all* metric. This is understandable since the contrastive random walk does not have a way of learning features for occluded pixels. Interestingly, we found that combining the two features together improved performance, and that the gap improves further when multi-frame walks are used, obtaining the overall best results.

Moreover, the combined features show more robustness on image pairs with rapid exposure and hue changes. We evaluated models trained on hue- and brightness-jittered image pairs and found that the model with our learned features performed significantly better, and that the gap increased with the magnitude of the jittering (Tab. 4b). Please see the supplement for details.

Finally, we used our learned features as part of the photometric loss for ARFlow [41], a recent unsupervised flow model. We added the  $\mathcal{L}_{\text{feat}}$  loss to their model and retrained it (while jointly learning the features through a contrastive random walk). The resulting model improves performance on all metrics, with a larger gain on Sintel (Tab. 4).

Table 2. **Nonparametric motion estimation.** Our single-scale, nonparametric, and regression-based methods.

Method	KITTI-15 <i>train</i>		
	noc	all	ER% (occ)
Jabri et al. [26]	12.63	19.41	64.50
Ours - Nonparametric	2.18	9.42	27.98
Ours - With Regression	2.09	3.86	12.45

Table 3. **Model configurations.** All models are 2-frame. Regressor-only finds a shortcut solution, predicting zero flow.

Configuration	KITTI-15 <i>train</i>		
	noc	all	ER%
Full	2.09	3.86	12.45
Nonparametric only	2.18	9.42	27.98
No feature consistency in $\mathcal{L}_{\text{feat}}$	2.20	5.02	17.53
Regressor only	14.21	21.45	41.34
No regressor constraint in $\mathcal{L}_{\text{feat}}$	5.54	10.44	26.43
No $\mathcal{L}_{\text{bound}}$	2.14	4.88	16.54
No $\mathcal{L}_{\text{smooth}}$	10.98	17.43	34.85

#### 4.3.4 Motion estimation ablations


To help understand which properties of our model contribute to its performance, we perform an ablation study on KITTI-15 (Table 3). We asked how the different losses contribute to the performance. We ablated the smoothness loss (Eq. 8), the self-supervision loss (Sec. 3.2), and removing the constraint on the regressor in Eq. 9 by setting  $\lambda_a = 0$ . We see that the smoothness loss significantly improves results. Similarly, we discarded the feature consistency term from Eq. 9, which reduces the quality of the results but outperforms the nonparametric model on the *all* metric.

**Training on internet video.** We found that our model generalized well to benchmark datasets when training solely on YouTube-VOS [81] (Tab. 7). For comparison, we also trained ARFlow [41] on YouTube-VOS. Our model obtained better performance on KITTI, while ARFlow per-

		Sintel				KITTI-15			
		Clean		Final					
Method		train	test	train	test	train	test	train	test
		EPE	EPE	EPE	EPE	all	noc	ER %	ER %
Supervised	FlowNetC [12]	(3.78)	6.85	(5.28)	8.51	-	-	-	-
	FlowNet2 [24]	(1.45)	4.16	(2.01)	5.74	(2.30)	-	(8.61)	11.48
	PWC-Net [65]	(1.70)	3.86	(2.21)	5.13	(2.16)	-	(9.80)	9.60
	RAFT [68]	(0.76)	1.94	(1.22)	3.18	(0.63)	-	(1.50)	5.10
Unsupervised	MFOccFlow [29]*	{3.89}	7.23	{5.52}	8.81	[6.59]	[3.22]	-	22.94
	EPIFlow [85]	3.94	7.00	5.08	8.51	5.56	2.56	-	16.95
	DDFlow [43]	{2.92}	6.18	{3.98}	7.40	[5.72]	[2.73]	-	14.29
	SelfFlow [44]*	[2.88]	[6.56]	{3.87}	{6.57}	[4.84]	[2.40]	-	14.19
	UFlow [32]	{2.50}	5.21	{3.39}	6.50	{2.71}	{1.88}	{9.05}	11.13
	SMURF-PWC [60]	2.63	-	3.66	-	2.73	-	9.33	-
	SMURF-RAFT [60]	{1.71}	3.15	{2.58}	4.18	{2.00}	{1.41}	{6.42}	6.83
Ours	ARFlow [41]	[2.79]	[4.78]	[3.73]	[5.89]	[2.85]	-	-	[11.80]
	Ours + ARFlow	[2.71]	[4.70]	[3.61]	[5.76]	[2.81]	[2.17]	[11.25]	[11.67]
	Ours (2-cycle)	{2.84}	5.68	{3.82}	6.72	{3.86}	{2.09}	{12.45}	13.10

(a) Optical flow benchmarks

Table 4. (a) Our model is trained on the train/test splits of the corresponding datasets. We reprint numbers from [32] and adopt their convention that “{ }” are trained on unlabeled evaluation data, “[ ]” are trained on related data (e.g., full Sintel movie) and “( )” are supervised. Methods that use 3 frames at test time are marked with \*. (b) We evaluate robustness to brightness and hue jittering.



(b) Robustness to jittering

Table 5. **Photometric features.** We evaluate the effectiveness of our features when they are used to define a photometric loss.

Losses	KITTI-15 train		
	noc	all	ER%
Charbonnier	2.28	5.69	19.30
Census	2.05	3.14	11.04
Our feats.	2.09	3.86	12.45
Our feats. + Charbonnier	2.21	4.51	14.25
Our feats. + Census	2.04	3.10	10.89
Our feats. (3-frame) + Census	2.02	3.03	10.67

Table 6. **Contrastive random walk ablations.** We evaluate different model parameters, including cycle length and number of scales in the multiscale random walk (from coarse to fine).

	KITTI-15 train			
	noc	all	ER%	
Cyc. len.	2	2.09	3.86	12.45
	3	2.05	3.46	12.28
	4	2.04	3.39	12.14
# levels	1	4.45	8.98	24.52
	3	2.36	4.55	14.35
	5	2.09	3.86	12.45

Table 7. **Training on internet video.** We train on YouTube-VOS [81] and evaluate on optical flow benchmarks.

	Sintel train		KITTI-15 train		
	Clean	Final	noc	all	ER%
Ours (2-cycle)	3.37	4.65	2.32	5.73	14.69
ARFlow [41]	3.22	4.51	2.65	6.01	16.47

formed better on Sintel.

### 4.3.5 Comparison to recent optical flow methods

To help understand our model’s overall performance, we compare it to recent optical flow methods (Table 4). We include models that use different numbers of frames for the random walk, and a variation of ARFlow [41] that uses our

self-supervised features to augment its photometric loss.

We found that our models outperform many recent unsupervised optical flow methods, including the (3-frame) **MFOccFlow** [29] and **EPIFlow**. In particular, we significantly outperform the recent **SelfFlow** method [44], despite the fact that it takes 3 frames as input at *test* time and uses Census Transform features. By contrast, our model *uses no hand-crafted image features*. The highest performing method is the very recent, highly optimized **SMURF** model [60], which uses the RAFT [68] architecture instead of PWC-net and which extends UFlow [32]. This model uses a variety of additional training signals, such as extensive data augmentation, occlusion inpainting with multi-stage training, self-distillation, and hand-crafted features.

## 5. Discussion

We have proposed a method for learning dense motion estimation using multiscale contrastive random walks. We see our work as a potential step toward unifying self-supervised tracking and optical flow. Moreover, the model can learn from internet video, which suggests that the emergent representations of such a hierarchical tracker may learn interesting part-whole structure at scale.

**Limitations and impact.** Motion analysis has many applications, such as in health monitoring, surveillance, and security. There is also a potential for the technology to be used for harmful purposes if weaponized. The released models are limited in scope to the datasets used in training.

**Acknowledgements.** We thank David Fouhey and Jeff Fessler for the helpful feedback. AO thanks Rick Szeliski for introducing him to multi-frame optical flow. This research was supported in part by Toyota Research Institute, Cisco Systems, and Berkeley Deep Drive.



## References

- [1] Robert Anderson, David Gallup, Jonathan T Barron, Janne Kontkanen, Noah Snavely, Carlos Hernández, Sameer Agarwal, and Steven M Seitz. Jump: virtual reality video. *ACM Transactions on Graphics (TOG)*, 35(6):1–13, 2016. [2](#)
- [2] Simon Baker and Iain Matthews. Lucas-kanade 20 years on: A unifying framework. In *International Journal of Computer Vision*, pages 221–255, February 2004. [2](#)
- [3] Simon Baker, Daniel Scharstein, JP Lewis, Stefan Roth, Michael J Black, and Richard Szeliski. A database and evaluation methodology for optical flow. *IJCV*, 2011. [2](#)
- [4] Michael J Black and Paul Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer vision and image understanding*, 63(1):75–104, 1996. [2](#), [3](#)
- [5] Thomas Brox, Christoph Bregler, and Jitendra Malik. Large displacement optical flow. In *CVPR*, 2009. [2](#)
- [6] Thomas Brox, Andrés Bruhn, Nils Papenberg, and Joachim Weickert. High accuracy optical flow estimation based on a theory for warping. In *ECCV*, 2004. [2](#), [4](#)
- [7] Thomas Brox and Jitendra Malik. Object segmentation by long term analysis of point trajectories. In *European conference on computer vision*, pages 282–295. Springer, 2010. [3](#)
- [8] Andrés Bruhn, Joachim Weickert, and Christoph Schnörr. Lucas/kanade meets horn/schunck: Combining local and global optic flow methods. *International journal of computer vision*, 61(3):211–231, 2005. [2](#)
- [9] Peter J. Burt and Edward H. Adelson. The laplacian pyramid as a compact image code. *IEEE Trans. Commun.*, 31:532–540, 1983. [1](#)
- [10] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In A. Fitzgibbon et al. (Eds.), editor, *European Conf. on Computer Vision (ECCV)*, Part IV, LNCS 7577, pages 611–625. Springer-Verlag, October 2012. [5](#)
- [11] A. Dosovitskiy, P. Fischer, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. v.d. Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2015. [5](#)
- [12] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Häusser, Caner Hazırbaş, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 2015. [3](#), [8](#)
- [13] PF Felzenszwalb and DR Huttenlocher. Efficient belief propagation for early vision. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–I. IEEE, 2004. [2](#)
- [14] Philipp Fischer, Alexey Dosovitskiy, Eddy Ilg, Philip Häusser, Caner Hazırbaş, Vladimir Golkov, Patrick Van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. *arXiv*, 2015. [4](#)
- [15] William T Freeman, Egon C Pasztor, and Owen T Carmichael. Learning low-level vision. *International journal of computer vision*, 40(1):25–47, 2000. [3](#)
- [16] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013. [5](#)
- [17] Daniel Gordon, Kiana Ehsani, Dieter Fox, and Ali Farhadi. Watching the world go by: Representation learning from unlabeled videos, 2020. [2](#)
- [18] Philip Haeusser, Alexander Mordvintsev, and Daniel Cremers. Learning by association—a versatile semi-supervised training method for neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 89–98, 2017. [3](#)
- [19] Bharath Hariharan, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Simultaneous detection and segmentation. In *Eur. Conf. Comput. Vis.*, pages 297–312, 2014. [6](#)
- [20] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. [1](#)
- [21] Berthold Horn and Brain Schunck. Determining optical flow. In *Artificial Intelligence*, pages 185–203, 1981. [2](#)
- [22] Zhaoyang Huang, Xiaokun Pan, Runsen Xu, Yan Xu, Kachun Cheung, Guofeng Zhang, and Hongsheng Li. Life: Lighting invariant flow estimation. *arXiv preprint arXiv:2104.03097*, 2021. [3](#)
- [23] Junhwa Hur and Stefan Roth. Mirrorflow: Exploiting symmetries in joint optical flow and occlusion estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 312–321, 2017. [2](#)
- [24] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017. [3](#), [8](#)
- [25] Eddy Ilg, Tonmoy Saikia, Margret Keuper, and Thomas Brox. Occlusions, motion and depth boundaries with a generic network for disparity, optical flow or scene flow estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 614–630, 2018. [2](#)
- [26] Allan Jabri, Andrew Owens, and Alexei A Efros. Space-time correspondence as a contrastive random walk. *Neural Information Processing Systems (NeurIPS)*, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [27] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks, 2015. [2](#)
- [28] Joel Janai, Fatma Güney, Anurag Ranjan, Michael Black, and Andreas Geiger. Unsupervised learning of multi-frame optical flow with occlusions. In *ECCV*, 2018. [2](#), [3](#)

- [29] Joel Janai, Fatma Güney, Anurag Ranjan, Michael J. Black, and Andreas Geiger. Unsupervised learning of multi-frame optical flow with occlusions. *ECCV*, 2018. 8
- [30] Joel Janai, Fatma Güney, Jonas Wulff, Michael J Black, and Andreas Geiger. Slow flow: Exploiting high-speed cameras for accurate and diverse optical flow reference data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 3
- [31] Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J Black. Towards understanding action recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 3192–3199, 2013. 5, 6
- [32] Rico Jonschkowski, Austin Stone, Jonathan T Barron, Ariel Gordon, Kurt Konolige, and Anelia Angelova. What matters in unsupervised optical flow. *arXiv preprint arXiv:2006.04902*, 2020. 1, 2, 4, 5, 6, 7, 8
- [33] Ryan Kennedy and Camillo J Taylor. Optical flow with geometric occlusion estimation and fusion of multiple frames. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 364–377. Springer, 2015. 3
- [34] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv*, 2014. 5
- [35] Zihang Lai, Erika Lu, and Weidi Xie. Mast: A memory-augmented self-supervised tracker. *arXiv preprint arXiv:2002.07793*, 2020. 1, 2
- [36] Cheng Lei and Yee-Hong Yang. Optical flow estimation on coarse-to-fine region-trees using discrete optimization. In *2009 IEEE 12th International Conference on Computer Vision*, pages 1562–1569. IEEE, 2009. 2
- [37] Ruibo Li, Guosheng Lin, and Lihua Xie. Self-point-flow: Self-supervised scene flow estimation from point clouds with optimal transport and random walk. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15577–15586, 2021. 3
- [38] Xueting Li, Sifei Liu, Shalini De Mello, Xiaolong Wang, Jan Kautz, and Ming-Hsuan Yang. Joint-task self-supervised learning for temporal correspondence. In *Advances in Neural Information Processing Systems*, pages 317–327, 2019. 2, 6
- [39] Yunpeng Li and Daniel P Huttenlocher. Learning for optical flow using stochastic optimization. In *European Conference on Computer Vision*, pages 379–391. Springer, 2008. 3
- [40] Ce Liu et al. *Beyond pixels: exploring new representations and applications for motion analysis*. PhD thesis, Massachusetts Institute of Technology, 2009. 2
- [41] Liang Liu, Jiangning Zhang, Ruifei He, Yong Liu, Yabiao Wang, Ying Tai, Donghao Luo, Chengjie Wang, Jilin Li, and Feiyue Huang. Learning by analogy: Reliable supervision from transformations for unsupervised optical flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6489–6498, 2020. 2, 7, 8
- [42] Pengpeng Liu, Irwin King, Michael R Lyu, and Jia Xu. DdfLOW: Learning optical flow with unlabeled data distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8770–8777, 2019. 2, 5
- [43] Pengpeng Liu, Irwin King, Michael R. Lyu, and Jia Xu. DDFlow: Learning optical flow with unlabeled data distillation. *AAAI*, 2019. 8
- [44] Pengpeng Liu, Michael Lyu, Irwin King, and Jia Xu. Self-low: Self-supervised learning of optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4571–4580, 2019. 2, 6, 8
- [45] Bruce D Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. *IJCAI*, 1981. 1, 2
- [46] Kunming Luo, Chuan Wang, Shuaicheng Liu, Haoqiang Fan, Jue Wang, and Jian Sun. Upflow: Upsampling pyramid for unsupervised optical flow learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1045–1054, 2021. 2
- [47] Simon Meister, Junhwa Hur, and Stefan Roth. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. *AAAI*, 2018. 1, 7
- [48] Simon Meister, Junhwa Hur, and Stefan Roth. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 2
- [49] Roland Memisevic and Geoffrey Hinton. Unsupervised learning of image transformations. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. 2
- [50] 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. *arXiv preprint arXiv:1912.01703*, 2019. 5
- [51] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Computer Vision and Pattern Recognition*, 2016. 6
- [52] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alexander Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv:1704.00675*, 2017. 5, 6
- [53] Jerome Revaud, Philippe Weinzaepfel, Zaid Harchaoui, and Cordelia Schmid. Epicflow: Edge-preserving interpolation of correspondences for optical flow. In *CVPR*, 2015. 2
- [54] Dan Rosenbaum, Daniel Zoran, and Yair Weiss. Learning the local statistics of optical flow. *Advances in Neural Information Processing Systems*, 26:2373–2381, 2013. 4
- [55] Stefan Roth, Victor Lempitsky, and Carsten Rother. Discrete-continuous optimization for optical flow estimation. In *Statistical and Geometrical Approaches to Visual Motion Analysis*, pages 1–22. Springer, 2009. 2

- [56] Michael Rubinstein, Ce Liu, and William T Freeman. Towards longer long-range motion trajectories. 2012. [3](#)
- [57] Peter Sand and Seth Teller. Particle video: Long-range motion estimation using point trajectories. *ICCV*, 2008. [3](#)
- [58] Alexander Shekhovtsov, Ivan Kovtun, and Václav Hlaváč. Efficient mrf deformation model for non-rigid image matching. *Computer Vision and Image Understanding*, 112(1):91–99, 2008. [2](#)
- [59] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, volume 11006, page 1100612. International Society for Optics and Photonics, 2019. [5](#)
- [60] Austin Stone, Daniel Maurer, Alper Ayvaci, Anelia Angelova, and Rico 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*, pages 3887–3896, 2021. [2](#), [8](#)
- [61] Deqing Sun, Stefan Roth, and Michael J Black. Secrets of optical flow estimation and their principles. In *CVPR*, 2010. [1](#), [2](#), [4](#), [7](#)
- [62] Deqing Sun, Stefan Roth, JP Lewis, and Michael J Black. Learning optical flow. In *European Conference on Computer Vision*, pages 83–97. Springer, 2008. [2](#), [3](#)
- [63] Deqing Sun, Erik Sudderth, and Michael Black. Layered image motion with explicit occlusions, temporal consistency, and depth ordering. *Advances in Neural Information Processing Systems*, 23:2226–2234, 2010. [3](#)
- [64] Deqing Sun, Daniel Vlasic, Charles Herrmann, Varun Jampani, Michael Krainin, Huiwen Chang, Ramin Zabih, William T Freeman, and Ce Liu. Autoflow: Learning a better training set for optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10093–10102, 2021. [2](#)
- [65] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume, 2017. [2](#), [3](#), [4](#), [5](#), [8](#)
- [66] Narayanan Sundaram, Thomas Brox, and Kurt Keutzer. Dense point trajectories by gpu-accelerated large displacement optical flow. In *European conference on computer vision*, pages 438–451. Springer, 2010. [2](#), [3](#)
- [67] Yansong Tang, Zhenyu Jiang, Zhenda Xie, Yue Cao, Zheng Zhang, Philip HS Torr, and Han Hu. Breaking shortcut: Exploring fully convolutional cycle-consistency for video correspondence learning. *arXiv preprint arXiv:2105.05838*, 2021. [2](#), [5](#)
- [68] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow, 2020. [2](#), [3](#), [6](#), [8](#)
- [69] Prune Truong, Martin Danelljan, Fisher Yu, and Luc Van Gool. Warp consistency for unsupervised learning of dense correspondences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10346–10356, 2021. [3](#)
- [70] Sebastian Volz, Andres Bruhn, Levi Valgaerts, and Henning Zimmer. Modeling temporal coherence for optical flow. In *2011 International Conference on Computer Vision*, pages 1116–1123. IEEE, 2011. [3](#)
- [71] Carl Vondrick, Abhinav Shrivastava, Alireza Fathi, Sergio Guadarrama, and Kevin Murphy. Tracking emerges by coloring videos. In *ECCV*, 2017. [2](#)
- [72] Chia-Ming Wang, Kuo-Chin Fan, Cheng-Tzu Wang, and Tong-Yee Lee. Estimating optical flow by integrating multi-frame information. *Journal of Information Science & Engineering*, 24(6), 2008. [3](#)
- [73] Heng Wang, Alexander Kläser, Cordelia Schmid, and Cheng-Lin Liu. Dense trajectories and motion boundary descriptors for action recognition. *International journal of computer vision*, 103(1):60–79, 2013. [3](#)
- [74] Ning Wang, Yibing Song, Chao Ma, Wengang Zhou, Wei Liu, and Houqiang Li. Unsupervised deep tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1308–1317, 2019. [2](#)
- [75] Xiaolong Wang, Allan Jabri, and Alexei A Efros. Learning correspondence from the cycle-consistency of time. In *CVPR*, 2019. [1](#), [6](#)
- [76] Yang Wang, Yi Yang, and Wei Xu. Occlusion aware unsupervised learning of optical flow. 2018. [2](#)
- [77] Zhongdao Wang, Hengshuang Zhao, Ya-Li Li, Shengjin Wang, Philip Torr, and Luca Bertinetto. Do different tracking tasks require different appearance models? *NervIPS*, 2021. [1](#)
- [78] Philippe Weinzaepfel, Jerome Revaud, Zaid Harchaoui, and Cordelia Schmid. Deepflow: Large displacement optical flow with deep matching. In *Proceedings of the IEEE international conference on computer vision*, pages 1385–1392, 2013. [2](#)
- [79] Jiarui Xu and Xiaolong Wang. Rethinking self-supervised correspondence learning: A video frame-level similarity perspective. *arXiv preprint arXiv:2103.17263*, 2021. [2](#), [6](#)
- [80] Li Xu, Jianing Chen, and Jiaya Jia. A segmentation based variational model for accurate optical flow estimation. In *European Conference on Computer Vision*, pages 671–684. Springer, 2008. [2](#)
- [81] Ning Xu, Linjie Yang, Yuchen Fan, Dingcheng Yue, Yuchen Liang, Jianchao Yang, and Thomas Huang. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327*, 2018. [5](#), [7](#), [8](#)
- [82] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *CVPR*, 2018. [2](#)
- [83] Jason J. Yu, Adam W. Harley, and Konstantinos G. Derpanis. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In *Computer Vision - ECCV 2016 Workshops, Part 3*, 2016. [1](#), [2](#)
- [84] Jean yves Bouguet. Pyramidal implementation of the lucas kanade feature tracker. *Intel Corporation, Microprocessor Research Labs*, 2000. [1](#)

- [85] Yiran Zhong, Pan Ji, Jianyuan Wang, Yuchao Dai, and Hongdong Li. Unsupervised deep epipolar flow for stationary or dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12095–12104, 2019. [2](#), [8](#)
- [86] Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. Technical report, 2002. [1](#)
- [87] CW Zitnick, Nebojsa Jojic, and Sing Bing Kang. Consistent segmentation for optical flow estimation. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pages 1308–1315. IEEE, 2005. [2](#)
- [88] Yuliang Zou, Zelun Luo, and Jia-Bin Huang. DF-Net: Unsupervised joint learning of depth and flow using cross-task consistency. *ECCV*, 2018. [3](#)