

SMURF: Self-Teaching Multi-Frame Unsupervised RAFT with Full-Image Warping

Austin Stone^{*1} Daniel Maurer^{*2} Alper Ayvaci² Anelia Angelova¹ Rico Jonschkowski¹
¹Robotics at Google ²Waymo
 {austinstone, anelia, rjon}@google.com {maurerd, ayvaci}@waymo.com
<https://github.com/google-research/google-research/tree/master/smurf>

Abstract

We present SMURF, a method for unsupervised learning of optical flow that improves state of the art on all benchmarks by 36% to 40% (over the prior best method UFlow) and even outperforms several supervised approaches such as PWC-Net and FlowNet2. Our method integrates architecture improvements from supervised optical flow, i.e. the RAFT model, with new ideas for unsupervised learning that include a sequence-aware self-supervision loss, a technique for handling out-of-frame motion, and an approach for learning effectively from multi-frame video data while still only requiring two frames for inference.

1. Introduction

Optical flow describes a dense pixel-wise correspondence between two images, specifying for each pixel in the first image, where that pixel is in the second image. The resulting vector field of relative pixel locations represents apparent motion or “flow” between the two images. Estimating this flow field is a fundamental problem in computer vision and any advances in flow estimation benefit many downstream tasks such as visual odometry, multi-view depth estimation, and video object tracking.

Classical methods formulate optical flow estimation as an optimization problem [2, 8]. They infer, for a given image pair, a flow field that maximizes smoothness and the similarity of matched pixels. Recent supervised learning approaches instead train deep neural networks to estimate optical flow from examples of ground-truth annotated image pairs [9, 28, 29, 37]. Since obtaining ground truth flow is extremely difficult for real images, supervised learning is generally limited to synthetic data [5, 20]. While these methods have produced excellent results in the training domain, generalization is challenging when the gap between the target domain and the synthetic training data is too wide.

Unsupervised learning is a promising direction to address this issue as it allows training optical flow models from unlabeled videos of any domain. The unsupervised approach works by combining ideas from classical methods and supervised-learning – training the same neural networks as in supervised approaches but optimizing them with objectives such as smoothness and photometric similarity from classical methods. Unlike those classical methods, unsupervised approaches perform optimization not per image pair but jointly for the entire training set.

Since unsupervised optical flow takes inspiration from classical and supervised learning methods, we can make substantial progress by properly combining novel ideas with insights from these two directions. In this paper, we do exactly that and make the following three contributions:

1. We integrate the current best supervised model, RAFT [29] with unsupervised learning and perform key changes to the loss functions and data augmentation to properly regularize this model for unsupervised learning.
2. We perform unsupervised learning on image crops while using the full image to compute unsupervised losses. This technique, which we refer to as full-image warping, improves flow quality near image boundaries.
3. We leverage a classical method for multi-frame flow refinement [19] to generate better labels for self-supervision from multi-frame input. This technique improves performance especially in occluded regions without requiring more than two frames for inference.

Our method *Self-Teaching Multi-frame Unsupervised RAFT with Full-Image Warping* (SMURF) combines these three contributions and improves the state of the art (SOTA) in unsupervised optical flow in all major benchmarks, i.e. it reduces errors by 40 / 36 / 39 % in the Sintel Clean / Sintel Final / KITTI 2015 benchmarks relative to the prior SOTA set by UFlow [12]. These improvements also reduce the gap to supervised approaches, as SMURF is the first unsupervised optical flow method that outperforms supervised FlowNet2 [9] and PWC-Net [28] on all benchmarks.

* Authors contributed equally.

2. Related Work

Optical flow was first studied in psychology to describe motion in the visual field [6]. Later, the rise of computing in the 1980s led to analytical techniques to estimate optical flow from images [8, 18]. These techniques introduced photometric consistency and smoothness assumptions. These early methods do not perform any learning but instead solve a system of equations to find flow vectors that minimize the objective function for a given image pair. Follow-up work continued to improve flow accuracy, e.g. through better optimization techniques and occlusion reasoning [2, 27].

Machine learning has helped to improve results substantially. First approaches used supervised convolutional neural networks that had relatively little flow-specific structure [5, 9]. Others introduced additional inductive biases from classical approaches such as coarse-to-fine search [23, 28, 37]. The current best network architecture is the Recurrent All-Pairs Field Transforms (RAFT) model. It follows classical work that breaks with the coarse-to-fine assumption [4, 26, 36] and computes the cost volume between all pairs of pixels and uses that information to iteratively refine the flow field [29]. All supervised methods rely heavily on synthetic labeled data for training. While producing excellent results in the supervised setting, RAFT has not previously been used in unsupervised learning.

Unsupervised approaches appeared after supervised methods and showed that even without labels, deep learning can greatly outperform classical flow methods [10, 11, 12, 14, 16, 17, 21, 24, 25, 33, 35, 38, 39, 42, 43]. Besides being more accurate than classical methods, learned methods are also faster at inference because all optimization occurs during training instead of during inference time.¹

A recent study [12] performed an extensive comparison of the many proposed advances in unsupervised flow estimation and amalgamated these different works into a state of the art method called UFlow. We take this work as a starting point and build on the techniques suggested there such as range-map based occlusion reasoning [35], the Census loss for photometric consistency [21, 40], edge aware smoothness [30], and self supervision [16, 17].

We build on and substantially extend this prior work by enabling the RAFT model to work in an unsupervised learning setting through changes in the loss function and data augmentation. We also propose full-image warping to make the photometric loss useful for pixels that leave the (cropped) image plane. And we utilize a flow refinement technique [19] to leverage multi-frame input during training to self-generate improved labels for self-supervision.

¹The reason why learned unsupervised methods outperform classical methods is not obvious as both approaches optimize similar objectives. We speculate that the network structure introduces helpful inductive biases and that optimizing a single network for all training images provides a better regularization than independently optimizing a flow field for each image.

3. SMURF for Unsupervised Optical Flow

This section describes our method *Self-Teaching Multi-frame Unsupervised RAFT with Full-Image Warping* (SMURF) in two parts. The first part covers preliminaries, such as the problem definition and components adopted from prior work. The second part describes the three main improvements in our method.

3.1. Preliminaries on Unsupervised Optical Flow

Given a pair of RGB images, $I_1, I_2 \in \mathbb{R}^{H \times W \times 3}$, we want to estimate the flow field $V_1 \in \mathbb{R}^{H \times W \times 2}$, which for each pixel in I_1 indicates the offset of its corresponding pixel in I_2 . We address this in a learning-based approach where we want to learn a function f_θ with parameters θ that estimates the flow field for any image pair, such that $V_1 = f_\theta(I_1, I_2)$. We learn the parameters θ from data of unlabeled image sequences $D = \{(I_i)_{i=1}^m\}$ by minimizing a loss function \mathcal{L} , $\theta^* = \arg \min \mathcal{L}(D, \theta)$.

We build on the loss function \mathcal{L} from UFlow [12],

$$\mathcal{L}(D, \theta) = \omega_{photo} \mathcal{L}_{photo}(D, \theta) + \omega_{smooth} \mathcal{L}_{smooth}(D, \theta) + \omega_{self} \mathcal{L}_{self}(D, \theta),$$

which is a weighted combination of three terms: occlusion aware photometric consistency \mathcal{L}_{photo} , edge-aware smoothness \mathcal{L}_{smooth} , and self-supervision \mathcal{L}_{self} .

The photometric consistency term is defined as

$$\mathcal{L}_{photo}(D, \theta) = \frac{1}{HW} \sum O_1 \odot \rho(I_1, w(I_2, V_1)),$$

where \odot represents element-wise multiplication and $\frac{1}{HW} \sum$ is shorthand notation for the mean over all pixels; we omit indices for better readability. The function $w(\cdot, \cdot)$ warps an image with a flow field, here estimated as $V_1 = f_\theta(I_1, I_2)$. The function $\rho(\cdot, \cdot)$ measures the photometric difference between two images based on a soft Hamming distance on the Census-transformed images and applies the generalized Charbonnier function [21]. The occlusion mask $O_1 \in \mathbb{R}^{H \times W}$ with entries $\in [0, 1]$ deactivates photometric consistency for occluded locations. This is crucial because occluded pixels, by definition, are not visible in the other image and therefore need not obey appearance consistency – the correct flow vector for an occluded pixel may point to a dissimilar-looking pixel in the other image. Like UFlow, we use a range-map based occlusion estimation [35] with gradient stopping for all datasets except for KITTI, where forward backward consistency [2] seems to yield better results [12].

The k -th order edge-aware smoothness term is defined as

$$\mathcal{L}_{smooth(k)}(D, \theta) = \frac{1}{HW} \sum \left(\exp \left(-\frac{\lambda}{3} \sum_c \left| \frac{\partial I_{1c}}{\partial x} \right| \right) \odot \left| \frac{\partial^k V_1}{\partial x^k} \right| + \exp \left(-\frac{\lambda}{3} \sum_c \left| \frac{\partial I_{1c}}{\partial y} \right| \right) \odot \left| \frac{\partial^k V_1}{\partial y^k} \right| \right).$$

The expression $\frac{1}{HW} \sum$ again computes the mean over all pixels, here for two components that compute the k -th derivative of the flow field V_1 along the x and y axes. These derivatives are each weighted by an exponential of the mean image derivative across color channels c . This weighting provides edge-awareness by ignoring non-smooth transitions in the flow field at visual edges [30]. The sensitivity to visual edges is controlled by λ .

The self-supervision term is defined as

$$\mathcal{L}_{self}(D, \theta) = \frac{1}{HW} \sum \hat{M}_1 \odot (1 - M_1) \odot c(\hat{V}_1, V_1).$$

where a self-generated flow label \hat{V}_1 [16] is compared to the flow prediction V_1 using the generalized Charbonnier function [27] $c(A, B) = ((A - B)^2 + \epsilon^2)^\alpha$ with all operations being applied elementwise. Our experiments use $\epsilon = 0.001$ and $\alpha = 0.5$. Prior work has modulated this comparison by masks \hat{M} and M for the generated labels and the predicted flow respectively, which are computed with a forward-backward consistency check [2], so that self-supervision is only applied to image regions where the label is forward-backward consistent but the prediction is not. However, as we will describe below, our method actually improves when we remove this masking.

The main reason to perform self-supervision is to improve flow estimation near the image boundaries where pixels can easily get occluded by moving outside of the image frame [16]. To address this problem via self-supervision, we apply the model on an image pair to produce a flow label and use this self-generated label to supervise flow predictions from a cropped version of the image pair. This way the model can transfer its own predictions from an easier to a more difficult setting. We build on and extend this technique with additional augmentations as described below.

3.2. SMURF’s Improvements

After having covered the foundation that our method builds on, we will now explain our three major improvements: 1) enabling the RAFT architecture [29] to work with unsupervised learning, 2) performing full-image warping while training on image crops, and 3) introducing a new method for multi-frame self-supervision.

3.2.1 Unsupervised RAFT

As model f_θ for optical flow estimation, we use the Recurrent All-Pairs Field Transforms (RAFT) [29] model, a recurrent architecture that has achieved top performance when trained with supervision but has not previously been used with unsupervised learning. RAFT works by first generating convolutional features for the two input images and then compiling a 4D cost volume $C \in \mathbb{R}^{H \times W \times H \times W}$ that contains feature-similarities for all pixel pairs between both

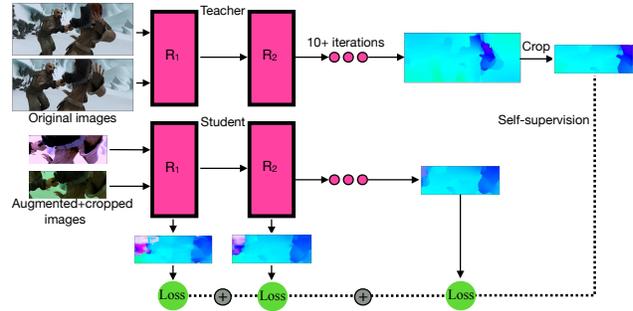


Figure 1. Self-supervision with sequence loss and augmentation. We use a single model as both “student” and “teacher”. As the teacher, we apply the model on full non-augmented images. As student, the model only sees a cropped and augmented version of the same images. The final output of the teacher is then cropped and used to supervise the predictions at all iterations of the student (to which the smoothness and photometric losses are applied as well). The advantages of this self-supervision method are three-fold: (1) the model learns to ignore photometric augmentations (2) the model learns to make better predictions at the borders and in occluded areas of the image, and (3) early iterations of the recurrent model learn from the output at the final iteration.

images. This cost volume is then repeatedly queried and fed into a recurrent network that iteratively builds and refines a flow field prediction. The only architectural modification we make to RAFT is to replace batch normalization with instance normalization [31] to enable training with very small batch sizes. Reducing the batch size was necessary to fit the model and the more involved unsupervised training steps into memory. But more importantly, we found that leveraging RAFT’s potential for unsupervised learning *requires key modifications to the unsupervised learning method*, which we will discuss next.

Sequence Losses In supervised learning of optical flow, it is common to apply losses not only to the final output but also to intermediate flow predictions [29, 28, 37]. In unsupervised learning this is not typically done – presumably because intermediate outputs of other models are often at lower resolutions which might not work well with photometric and other unsupervised losses [12, 16]. However, RAFT produces intermediate predictions at full resolution and does not pass gradients between prediction steps. Thus, we apply the unsupervised losses at the entire sequence of RAFT’s intermediate predictions, which we found to be essential to prevent divergence and to achieve good results. Similarly to supervised methods, we exponentially decay the weight of these losses at earlier iterations [29]. $\mathcal{L}_{sequence} = \sum_{i=1}^n \gamma^{n-i} \mathcal{L}_i$, where n is the number of flow iterations, γ is the decay factor, and \mathcal{L}_i is the loss at iteration i . Our experiments use $n = 12$ and $\gamma = 0.8$.

Improved Self-Supervision For self-supervision, we implement the sequence loss by applying the model to a full

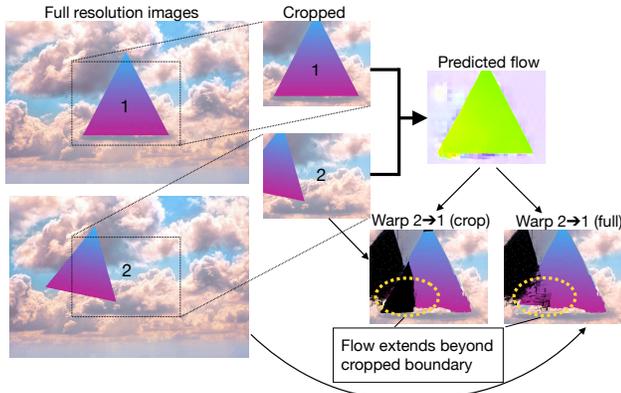


Figure 2. Full-image warping. Images are cropped for flow prediction, but warping of image 2 with the predicted flow (to compute the photometric loss), is done with the full size image. The advantage is shown in the lower right: Compared to warping the cropped image (left), full-image warping reduces occlusions from out-of-frame motion (shown in black) and is able to better reconstruct image 1. When used during training, full-image warping provides a learning signal for pixels that move outside the cropped image boundary. The remaining occlusions in the reconstruction are due to noisy flow predictions for the out-of-frame motion.

image, taking its final output, and using that to supervise all iterations of the model applied to the cropped image (see Figure 1). Applying self-supervision from full to cropped images has been shown to provide a learning signal to pixels near the image border that move out of the image and therefore receive no signal from photometric losses [16]. Performing this self-supervision in a sequence-aware manner allows the earlier stages of the model to learn from the model’s more refined final output. Potentially because of this inherent quality difference between the “student” and the “teacher” flow, we found that simplifying the self-supervision loss by removing masks M and \hat{M} improves performance further. The self-supervision loss without masks is $\mathcal{L}_{self}(D, \theta) = \frac{1}{HW} \sum c(\hat{V}_1, V_1)$. Note that this change does not affect the photometric loss \mathcal{L}_{photo} , where we continue to use O to mask out occlusions.

Extensive Data Augmentation To regularize RAFT, we use the same augmentation as supervised RAFT [29] which is much stronger than what has typically been used in unsupervised optical flow, except for the recent ARFlow [14]. We randomly vary hue, brightness, saturation, stretching, scaling, random cropping, random flipping left/right and up/down, and we apply a random eraser augmentation that removes random parts of each image. All augmentations are applied to the model inputs, but not to the images used to compute the photometric and smoothness losses. The self-generated labels for self-supervision are computed from un-augmented images, which has the benefit of training the model to ignore these augmentations (see Figure 1).

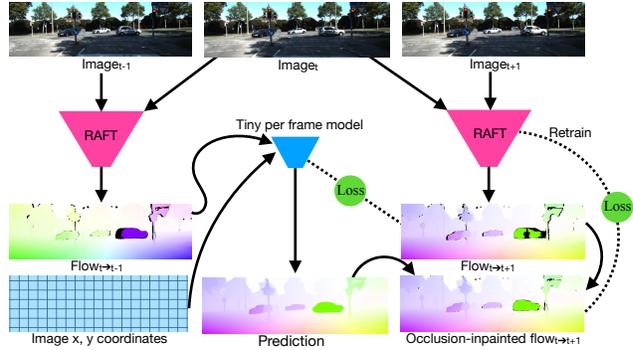


Figure 3. Multi-frame self-supervision. From a sequence of three frames $t - 1$, t , and $t + 1$, we compute the backward flow ($t \rightarrow t - 1$) and the forward flow ($t \rightarrow t + 1$). The backward flow is then inverted via a tiny model (which is trained for this frame pair) and used to inpaint occluded regions in the forward flow. This inpainted flow field is then used to retrain the RAFT model.

3.2.2 Full-Image Warping

The photometric loss, which is essential for unsupervised optical flow estimation, is generally limited to flow vectors that stay inside the image frame because vectors that point outside of the frame have no pixels to compare their photometric appearance to. We address this limitation by computing the flow field from a cropped version of the images I_1 and I_2 while referencing the full, uncropped image I_2 when warping it with the estimated flow V_1 before computing the photometric loss (see Figure 2). As we also no longer mark these flow vectors that move outside the image frame as occluded, they now provide the model with a learning signal. We use full-image warping for all datasets except Flying Chairs, where we found that cropping the already small images hurt performance.

3.2.3 Multi-Frame Self-Supervision

Finally, we propose to leverage multi-frame information for self-supervision to generate better labels in occluded areas, inspired by work that used a similar technique for inference [19]. For multi-frame self-supervision, we take a frame t and compute the forward flow to the next frame ($t \rightarrow t + 1$) and the backward flow to the previous frame ($t \rightarrow t - 1$). We then use the backward flow to predict the forward flow through a tiny learned inversion model and use that prediction to inpaint areas that were occluded in the original forward flow but were not occluded in the backward flow – which is why the estimate from the backward flow is more accurate (see Figure 3). The tiny model for the backward-forward inversion consists of three layers of 3×3 convolutions with [16, 16, 2] channels that are applied on the backward flow and the image coordinates normalized to $[-1, 1]$. The model is re-initialized and trained per frame using the non-occluded forward flow as supervision, after

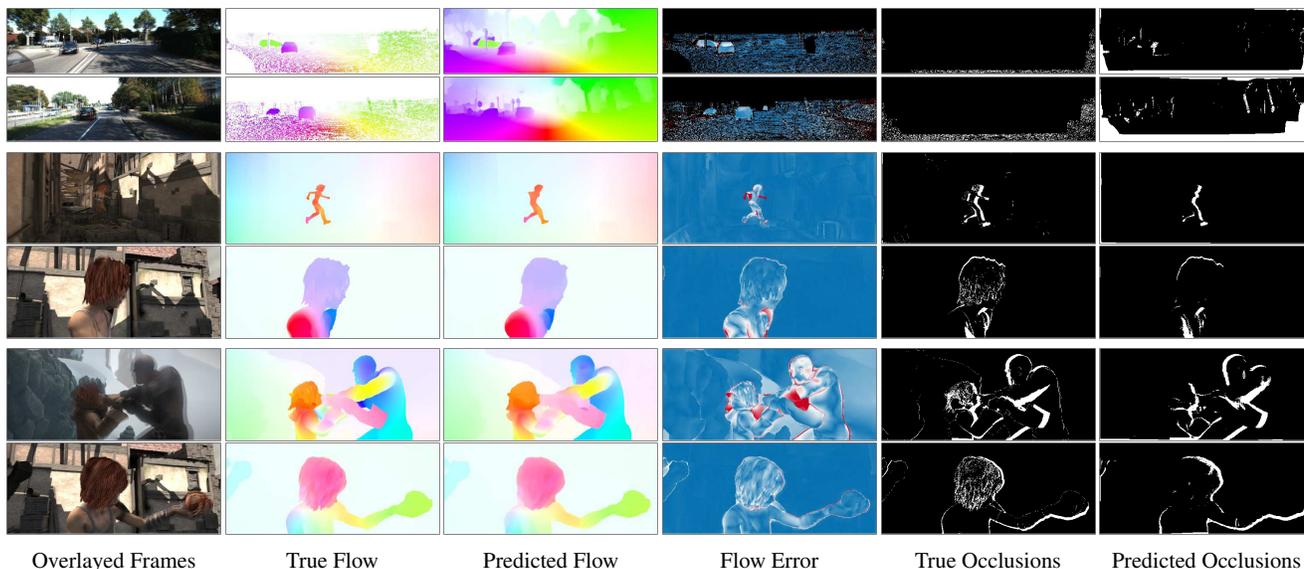


Figure 4. Qualitative results for SMURF on two random examples from KITTI 2015, Sintel Clean, and Sintel Final (top to bottom) not seen during training. These results show the model’s ability to estimate fast motions, relatively fine details, and substantial occlusions.

which the in-painted flow field is stored and used for self-supervision. We apply multi-frame self-supervision only at the final stage of training. Importantly, we use multiple frames only during training and not for inference.

4. Experiments

Our hyperparameters are based on UFlow [12] with slight modifications based on further hyperparameter search. In all experiments, we use weights $\omega_{photo} = 1$ and $\omega_{self} = 0.3$. Regarding smoothness parameters, we set edge sensitivity $\lambda = 150$ and $\omega_{smooth} = 4$ for KITTI and Chairs and $\omega_{smooth} = 2.5$ for Sintel. We use 2nd order smoothness ($k = 2$) for KITTI and 1st order smoothness ($k = 1$) for all other datasets. We train for 75K iterations with batch size 8, except for Sintel where we train for only 15K iterations to avoid overfitting. The self-supervision weight is set to 0 for the first 40% of gradient steps, then linearly increased to 0.3 during the next 10% of steps and then kept constant. All of our training uses Adam [13] ($\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$) with learning rate of 0.0002, which takes about 1 day to converge on 8 GPUs running synchronous SGD. On all datasets, the learning rate is exponentially decayed to $\frac{1}{1000}$ of its original value over the last 20% of steps. On Sintel and on Flying Chairs we train with random crops of 368×496 , and on KITTI we train with random crops of size 296×696 . For all datasets we evaluated our model on the input resolution that had the lowest average end point error on the training set; 488×1144 for KITTI and 480×928 for Sintel. All test images were bilinearly resized to these resolutions during inference, and the resulting flow field was bilinearly resized and rescaled back

to the native image size to compute evaluation metrics.

During the second stage of training when multi-frame self-supervision is applied, we generate labels for all images using the model trained according to the procedure described above. We then continue training the same model with only the self-supervision loss for an additional 30K iterations using the multiframe generated labels. We use the same hyperparameters as in the first stage but exponentially decay the learning rate for the last 5K iterations. For our best performing Sintel model, we train with the KITTI self-supervision labels mixed in at a ratio of 50%.

Datasets We train and evaluate our model according to the conventions in the literature using the following optical flow datasets: Flying Chairs [5], Sintel [3] and KITTI 2015 [22]. We pretrain on Flying Chairs before fine tuning on Sintel or KITTI. Similar to UFlow [12], we did not find a benefit to pretraining on more out-of-domain data, e.g. Flying Things [20]. None of the training techniques in our method uses any ground truth labels. We train on the “training” portion of Flying Chairs, and divide the Sintel dataset according to its standard train / test split. For KITTI, we train on the multi-view extension following the split used in prior work [12, 42]: We train two models, one on the multi-view extension of the training set and one on the extension of the test set, evaluate these models appropriately. For ablations, we report metrics after training on the “test” portion of the dataset (which does not include labels) and evaluating on the training set, and for final benchmark numbers we report results after training on the training portion only. For our benchmark result on Sintel, we train on a 50% mixture of the KITTI and Sintel multi-frame self-supervision labels.

		Sintel Clean [3]		Sintel Final [3]		KITTI 2015 [22]			
Method		EPE		EPE		EPE	EPE (noc)		ER in %
		<i>train</i>	<i>test</i>	<i>train</i>	<i>test</i>	<i>train</i>	<i>train</i>	<i>train</i>	<i>test</i>
Supervised in domain	FlowNet2-ft [9]	(1.45)	4.16	(2.01)	5.74	(2.30)	–	(8.61)	11.48
	PWC-Net-ft [28]	(1.70)	3.86	(2.21)	5.13	(2.16)	–	(9.80)	9.60
	SelfFlow-ft [17] ^(MF)	(1.68)	[3.74]	(1.77)	{4.26}	(1.18)	–	–	8.42
	VCN-ft [37]	(1.66)	2.81	(2.24)	4.40	(1.16)	–	(4.10)	6.30
	RAFT-ft [29]	(0.76)	1.94	(1.22)	3.18	(0.63)	–	(1.5)	5.10
Supervised out of domain	FlowNet2 [9]	2.02	3.96	3.14	6.02	9.84	–	28.20	–
	PWC-Net [28]	2.55	–	3.93	–	10.35	–	33.67	–
	VCN [37]	2.21	–	3.62	–	8.36	–	25.10	–
	RAFT [29]	1.43	–	2.71	–	5.04	–	17.4	–
	Unsupervised	EPiFlow [42]	3.94	7.00	5.08	8.51	5.56	2.56	–
	DDFlow [16]	{2.92}	6.18	{3.98}	7.40	[5.72]	[2.73]	–	14.29
	SelfFlow [17] ^(MF)	[2.88]	[6.56]	{3.87}	{6.57}	[4.84]	[2.40]	–	14.19
	UnsupSimFlow [10]	{2.86}	5.92	{3.57}	6.92	[5.19]	–	–	[13.38]
	ARFlow [14] ^(MF)	{2.73}	{4.49}	{3.69}	{5.67}	[2.85]	–	–	[11.79]
	UFlow [12]	3.01	5.21	4.09	6.50	2.84	1.96	9.39	11.13
	SMURF-test (ours)	1.99	–	2.80	–	2.01	1.42	6.72	–
	SMURF-train (ours)	{1.71}	3.15	{2.58}	4.18	{2.00}	{1.41}	{6.42}	6.83

Table 1. Comparison to state of the art. SMURF-train / test is our model trained on the train / test split of the corresponding dataset. The best results per category are shown in bold – note that supervision “in domain” is often not possible in practice as flow labels for real images are difficult to obtain. Braces indicate results that might have overfit because evaluation data was used for training: “{}” evaluated on the same labeled data as was used for training, “{}” trained on the unlabeled evaluation set, and “[]” trained on data highly related to the evaluation set (e.g., the entire Sintel Movie or < 5 frames away from an evaluation image in KITTI). Methods that use multiple frames at inference are denoted with “MF”; our method uses multiple frames only during training.

For all datasets we report endpoint error (“EPE”), and for KITTI, we additionally report error rates (“ER”), where a prediction is considered erroneous if its EPE is > 3 pixels or > 5% of the length of the true flow vector. We generally compute these metrics for all pixels, except for “EPE (noc)” where only non-occluded pixels are considered.

5. Results

In this section, we compare SMURF to related methods, ablate the proposed improvements, and show limitations.

5.1. Comparison to State of the Art

Qualitative results for SMURF are shown in Figure 4 and a comparison to other methods can be found in Table 1. This comparison shows that our model substantially outperforms all prior published methods on unsupervised optical flow on all benchmarks. Compared to the the previous state of the art method UFlow [12], our method reduces benchmark test errors by 40 / 36 / 39 % for Sintel Clean / Sintel Final / KITTI 2015. When we compare SMURF trained in domain to supervised methods trained out of domain, SMURF outperforms all supervised methods except RAFT [29] on all benchmarks². Compared to supervised RAFT, it performs a bit worse on Sintel (1.99 vs 1.43 on Sintel Clean, 2.80 vs. 2.71 on Sintel Final) but much better on KITTI 2015 (2.01

²This is a fair comparison as obtaining labels for a given domain is extremely difficult while our method trains on readily available video data.

		Chairs	Sintel train		KITTI-15 train	
Method		<i>test</i>	Clean	Final	EPE	ER%
Train on Chairs	DDFlow [16]	2.97	4.83	4.85	17.26	–
	UFlow [12]	2.55	4.36	5.12	15.68	32.69
	SMURF (ours)	1.72	2.19	3.35	7.94	26.51
Train on Sintel	DDFlow [16]	3.46	{2.92}	{3.98}	12.69	–
	UFlow [12]	3.25	3.01	4.09	7.67	17.41
	SMURF (ours)	1.99	1.99	2.80	4.47	12.55
Train on KITTI	DDFlow [16]	6.35	6.20	7.08	[5.72]	–
	UFlow [12]	5.05	6.34	7.01	2.84	9.39
	SMURF (ours)	3.26	3.38	4.47	2.01	6.72

Table 2. Generalization across datasets. These results compare our method to state of the art unsupervised methods in a setting where a model is trained on one dataset and tested on different one.

vs. 5.04 EPE, 6.72% vs. 17.4% ER). SMURF even outperforms some supervised methods when they are finetuned on the test domain, e.g. FlowNet2-ft [9], PWC-Net-ft [28], and SelfFlow-ft [17]. Only VCN-ft [37] and RAFT-ft [29] achieve better performance here. The inference time of our model is the same as the supervised RAFT model, approximately 500ms when using 12 recurrent iterations. A qualitative comparison to supervised RAFT is shown in Figure 5.

In Table 2, we compare generalization across domains to prior unsupervised methods and find substantial improvements for every combination of training and test domain. Even after only training on Flying Chairs, our model already achieves an EPE of 2.19 on Sintel Clean and 3.35 on Sintel Final, which is superior to all prior unsupervised tech-

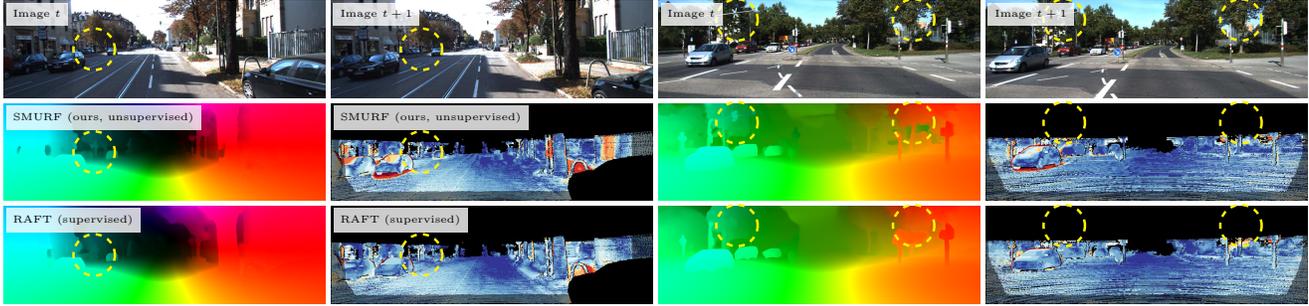


Figure 5. Qualitative comparison of unsupervised SMURF and supervised RAFT. Maybe unsurprisingly supervised training in domain reduces errors especially for challenging and ambiguous cases, e.g. reflections and shadows of moving objects. But interestingly, there are also areas (highlighted) where our unsupervised method works better than the supervised alternative, apparently at small objects in places where labels are sparse and potentially imperfect (left) or non-existent (right).

Method	Trained on		KITTI-15 <i>train</i>		KITTI-15 <i>test</i>	
	Flow	Stereo	EPE	D1(all)%	D1(noc)%	D1(all)%
SGM [7]	–	–	–	–	5.62	6.38
SsSMNet [41]	–	✓	–	–	{3.06}	{3.40}
UnOS [34]	✓	✓	–	5.94	–	6.67
Flow2Stereo [15]	✓	✓	1.34	6.13	6.29	6.61
Reversing-PSMNet [1]	–	✓	1.01	3.85	3.86	4.06
SMURF-train (ours)	✓	–	1.03	4.31	4.51	4.77

Table 3. Stereo depth estimation. Without fine-tuning, our flow model estimates stereo depth “zero-shot” at an accuracy comparable to state of the art unsupervised methods trained for that task.

Method	Model	Training input	Sintel <i>train</i>		KITTI-15 <i>train</i>	
			Clean	Final	EPE	ER%
UFlow [12]	PWC	High res.	3.01	4.09	2.84	9.39
UFlow	PWC	Low res.	3.43	4.33	3.87	13.04
UFlow	RAFT	Low res.	3.36	4.32	4.27 ⁺	14.18 ⁺
SMURF	PWC	Image crop	2.63	3.66	2.73	9.33
SMURF	RAFT	Image crop	1.99	2.80	2.01	6.72

Table 4. Test of whether replacing PWC-Net with RAFT improves the prior best unsupervised method UFlow. Due to memory constraints, training RAFT in this setting requires a lower resolution than UFlow [12] (384×512 for Sintel, 320×704 for KITTI). We use that resolution for a side-by-side comparison to PWC (rows 2-3), which shows that RAFT performs poorly without our proposed modifications. Results marked with ⁺ overfit the dataset and generate worse performance than reported here at the end of training.

niques even if fine tuned on Sintel data (first and third column in Table 1).

We also tested generalization from optical flow to stereo depth estimation. Here we evaluated our exact trained KITTI-15 flow model in the KITTI 2015 stereo depth benchmark [22]. Despite neither tailoring the architecture or losses to stereo depth estimation nor training or fine tuning on any stereo data, our method achieves results that are competitive with the best unsupervised stereo depth estimation methods (see Table 3).

5.2. Ablation Study

To determine which aspects of our model are responsible for its improved performance over prior work, we perform

SQ	AU	FW	MF	Sintel <i>train</i>		KITTI-15 <i>train</i>	
				Clean	Final	EPE	ER%
–	–	–	–	2.92	3.92	8.40	18.57
–	✓	–	–	– <i>diverged</i> –	– <i>diverged</i> –	– <i>diverged</i> –	– <i>diverged</i> –
✓	–	–	–	2.66	3.86	5.01 ⁺	17.50 ⁺
✓	✓	–	–	2.38	3.17	3.64 ⁺	13.05 ⁺
–	✓	✓	–	– <i>diverged</i> –	– <i>diverged</i> –	– <i>diverged</i> –	– <i>diverged</i> –
✓	–	✓	–	2.71	3.46	2.78	8.47
✓	✓	✓	–	2.15	2.99	2.45	7.53
✓	✓	✓	✓	1.99	2.80	2.01	6.72

Table 5. Ablation of proposed improvements. SQ: sequence loss, AU: heavy augmentation, FW: full-image warping, MF: multi-frame self-supervision. All components significantly improve performance and the sequence loss prevents divergence.

an extensive ablation study. In these ablations, we always train one model per domain (on KITTI-2015-test and Sintel-test after pretraining on Flying Chairs), and evaluate those on the corresponding validation split of the same domain.

RAFT Model Our first ablation investigates how much improvement can be obtained by taking the prior state of the art method UFlow and replacing its PWC model with RAFT. As the results in Table 4 show, replacing the model without additional changes to the unsupervised learning method surprisingly does not improve but instead decreases performance. Through extensive experimentation, we identified and added the techniques presented here that enable superior unsupervised learning with RAFT. The gains from these techniques are much smaller with the PWC model, potentially because of the more constrained architecture.

SMURF Components Next, we test different combinations of the novel components in our method. The results in Table 5 show that every component has a significant impact on performance. Sequence losses and heavy augmentations are necessary to achieve good results, especially on Sintel. And when using heavy augmentations, we need to apply sequence losses to prevent divergence of the model. Full-Image warping and multi-frame self-supervision have the strongest effect on KITTI, which makes sense as larger

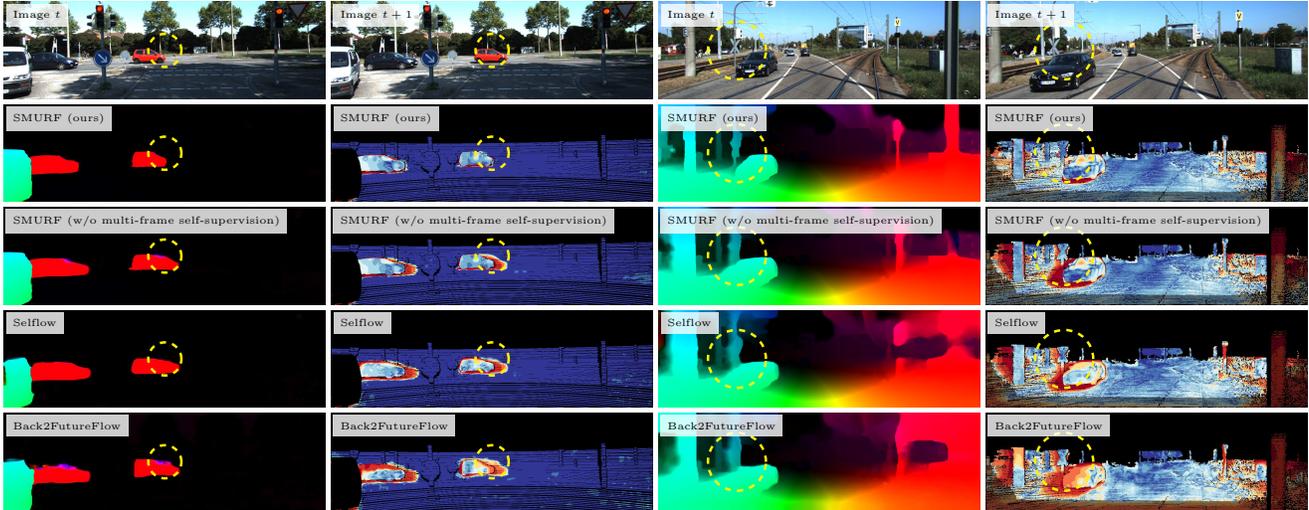


Figure 6. Qualitative ablation and comparison of multi-frame self-supervision. The highlighted areas indicate clear improvements in occluded areas when using multi-frame self-supervision (second vs. third row). These improvements substantially outperform other unsupervised multi-frame methods (fourth and fifth row) although those use additional frames not only during training but also for inference.

Self-sup. variant	Sintel <i>train</i>		KITTI-15 <i>train</i>	
	Clean	Final	EPE	ER%
No self-supervision	3.88	4.50	3.22	8.40
From intermediate predictions	2.51	3.22	2.58	7.54
W/ FB masking	2.53	3.31	2.94	8.14
W/o FB masking (Ours)	2.15	2.99	2.45	7.53

Table 6. Ablation of self-supervision improvements. All results are without multi-frame self-supervision but including all other components. FB masking refers to forward-backward consistency masking that prior work used in the self-supervision loss. Occlusions in the photometric loss are always masked.



Figure 7. Limitations of unsupervised flow. Optimizing photometric consistency can produce incorrect flow at shadows / reflections.

camera motion in that dataset causes more occlusions at the image boundaries which these components help to address.

Self-Supervision Modifications We also ablated our proposed changes to self-supervision (Section 3.2.1). The ablations in Table 6 show that even with full-image warping, self-supervision remains an important component, and that it works best when not masking the loss as in prior work and when we use the final output (not intermediate model predictions) to generate the self-supervision labels.

Multi-Frame Self-Supervision Lastly, we provide a qualitative ablation of multi-frame self-supervision (Section 3.2.3) and a comparison to other unsupervised methods that use multi-frame information [11, 17]. Figure 6

shows that multi-frame self-supervision substantially improves flow accuracy in occluded areas and does this much better than related multi-frame methods while being the only approach that requires multiple frames only during training and not for inference.

5.3. Limitations

A major limitation of unsupervised optical flow is that it estimates apparent visual motion rather than motion of physical objects (see Figure 7). Overcoming this limitation requires some form of supervision, reasoning about the 3-D space as in scene flow [32], reasoning about semantics, or a combination of these. Future work could try to transfer the techniques from our method to such approaches.

6. Conclusion

We have presented SMURF, an effective method for unsupervised learning of optical flow that reduces the gap to supervised approaches and shows excellent generalization across datasets and even to “zero-shot” depth estimation. SMURF brings key improvements, most importantly (1) enabling the RAFT architecture to work in an unsupervised setting via modifications to the unsupervised losses and data augmentation, (2) full-image warping for learning to predict out of frame motion, and (3) multi-frame self-supervision for improved flow estimates in occluded regions. We believe that these contributions are a step towards making unsupervised optical flow truly practical, so that optical flow models trained on unlabeled videos can provide high quality pixel-matching in domains without labeled data.

References

- [1] Filippo Aleotti, Fabio Tosi, Li Zhang, Matteo Poggi, and Stefano Mattoccia. Reversing the cycle: Self-supervised deep stereo through enhanced monocular distillation. In *ECCV*, 2020. 7
- [2] Thomas Brox, Andrés Bruhn, Nils Papenberg, and Joachim Weickert. High accuracy optical flow estimation based on a theory for warping. In *ECCV*, 2004. 1, 2, 3
- [3] Daniel J. Butler, Jonas Wulff, Garrett B. Stanley, and Michael J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, 2012. 5, 6
- [4] Qifeng Chen and Vladlen Koltun. Full flow: Optical flow estimation by global optimization over regular grids. In *CVPR*, 2016. 2
- [5] 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 *ICCV*, 2015. 1, 2, 5
- [6] James J. Gibson. *The Perception of the Visual World*. Houghton Mifflin, 1950. 2
- [7] H. Hirschmüller. Stereo processing by semi-global matching and mutual information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):328–341, February 2008. 7
- [8] Berthold K. P. Horn and Brian G. Schunck. Determining optical flow. *AI*, 1981. 1, 2
- [9] 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 *CVPR*, 2017. 1, 2, 6
- [10] Woobin Im, Tae-Kyun Kim, and Sung-Eui Yoon. Unsupervised learning of optical flow with deep feature similarity. In *ECCV*, 2020. 2, 6
- [11] Joel Janai, Fatma Güney, Anurag Ranjan, Michael J. Black, and Andreas Geiger. Unsupervised learning of multi-frame optical flow with occlusions. In *ECCV*, 2018. 2, 8
- [12] Rico Jonschkowski, Austin Stone, Jonathan T Barron, Ariel Gordon, Kurt Konolige, and Anelia Angelova. What matters in unsupervised optical flow. In *ECCV*, 2020. 1, 2, 3, 5, 6, 7
- [13] Diederick P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [14] 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 *CVPR*, 2020. 2, 4, 6
- [15] Pengpeng Liu, Irwin King, Michael Lyu, and Jia Xu. Flow2Stereo: Effective self-supervised learning of optical flow and stereo matching. In *CVPR*, 2020. 7
- [16] Pengpeng Liu, Irwin King, Michael R. Lyu, and Jia Xu. DDFlow: Learning optical flow with unlabeled data distillation. In *AAAI*, 2019. 2, 3, 4, 6
- [17] Pengpeng Liu, Michael R. Lyu, Irwin King, and Jia Xu. Self-low: Self-supervised learning of optical flow. *CVPR*, 2019. 2, 6, 8
- [18] Bruce D. Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. *DARPA Image Understanding Workshop*, 1981. 2
- [19] Daniel Maurer and Andrés Bruhn. Proflow: Learning to predict optical flow. In *BMVC*, 2018. 1, 2, 4
- [20] Nikolaus Mayer, Eddy Ilg, Philip Häusser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *CVPR*, 2016. 1, 5
- [21] Simon Meister, Junhwa Hur, and Stefan Roth. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In *AAAI*, 2018. 2
- [22] Moritz Menze, Christian Heipke, and Andreas Geiger. Joint 3d estimation of vehicles and scene flow. *ISPRS Workshop on Image Sequence Analysis*, 2015. 5, 6, 7
- [23] Anurag Ranjan and Michael J. Black. Optical flow estimation using a spatial pyramid network. In *CVPR*, 2017. 2
- [24] Anurag Ranjan, Varun Jampani, Lukas Balles, Kihwan Kim, Deqing Sun, Jonas Wulff, and Michael J. Black. Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation. In *CVPR*, 2019. 2
- [25] Zhe Ren, Junchi Yan, Bingbing Ni, Bin Liu, Xiaokang Yang, and Hongyuan Zha. Unsupervised deep learning for optical flow estimation. In *AAAI*, 2017. 2
- [26] Frank Steinbrücker, Thomas Pock, and Daniel Cremers. Large displacement optical flow computation without warping. In *ICCV*, 2009. 2
- [27] Deqing Sun, Stefan Roth, and Michael J. Black. Secrets of optical flow estimation and their principles. In *CVPR*, 2010. 2, 3
- [28] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *CVPR*, 2018. 1, 2, 3, 6
- [29] Zachary Teed and Jia Deng. Raft: Recurrent all pairs field transforms for optical flow. In *ECCV*, 2020. 1, 2, 3, 4, 6
- [30] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *ICCV*, 1998. 2, 3
- [31] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv:1607.08022*, 2016. 3
- [32] Sundar Vedula, Simon Baker, Peter Rander, Robert Collins, and Takeo Kanade. Three-dimensional scene flow. In *CVPR*, 1999. 8
- [33] Yang Wang, Peng Wang, Zhenheng Yang, Chenxu Luo, Yi Yang, and Wei Xu. Unos: Unified unsupervised optical-flow and stereo-depth estimation by watching videos. In *CVPR*, 2019. 2
- [34] Yang Wang, Peng Wang, Zhenheng Yang, Chenxu Luo, Yi Yang, and Wei Xu. Unos: Unified unsupervised optical-flow and stereo-depth estimation by watching videos. In *CVPR*, 2019. 7
- [35] Yang Wang, Yi Yang, Zhenheng Yang, Liang Zhao, Peng Wang, and Wei Xu. Occlusion aware unsupervised learning of optical flow. In *CVPR*, 2018. 2
- [36] Jia Xu, René Ranftl, and Vladlen Koltun. Accurate optical flow via direct cost volume processing. In *CVPR*, 2017. 2

- [37] Gengshan Yang and Deva Ramanan. Volumetric correspondence networks for optical flow. In *NeurIPS*, 2019. 1, 2, 3, 6
- [38] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *CVPR*, 2018. 2
- [39] 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 *ECCV Workshop*, 2016. 2
- [40] Ramin Zabih and John Woodfill. Non-parametric local transforms for computing visual correspondence. In *ECCV*, 1994. 2
- [41] Yiran Zhong, Yuchao Dai, and Hongdong Li. Self-supervised learning for stereo matching with self-improving ability. *arXiv:1709.00930*, 2017. 7
- [42] Yiran Zhong, Pan Ji, Jianyuan Wang, Yuchao Dai, and Hongdong Li. Unsupervised deep epipolar flow for stationary or dynamic scenes. In *CVPR*, 2019. 2, 5, 6
- [43] Yuliang Zou, Zelun Luo, and Jia-Bin Huang. DF-Net: Unsupervised joint learning of depth and flow using cross-task consistency. In *ECCV*, 2018. 2