

Optical Flow for Autonomous Driving: Applications, Challenges and Improvements

Shihao Shen¹, Louis Kerofsky² and Senthil Yogamani³

¹Carnegie Mellon University, Pittsburgh, Pennsylvania, U.S.

²Qualcomm Technologies, Inc., San Diego, California, U.S.

³Automated Driving, QT Technologies Ireland Limited.

ABSTRACT

Optical flow estimation is a well-studied topic for automated driving applications. Many outstanding optical flow estimation methods have been proposed, but they become erroneous when tested in challenging scenarios that are commonly encountered. Despite the increasing use of fisheye cameras for near-field sensing in automated driving, there is very limited literature on optical flow estimation with strong lens distortion. Thus we propose and evaluate training strategies to improve a learning-based optical flow algorithm by leveraging the only existing fisheye dataset with optical flow ground truth. While trained with synthetic data, the model demonstrates strong capabilities to generalize to real world fisheye data. The other challenge neglected by existing state-of-the-art algorithms is low light. We propose a novel, generic semi-supervised framework that significantly boosts performances of existing methods in such conditions. To the best of our knowledge, this is the first approach that explicitly handles optical flow estimation in low light.

I INTRODUCTION

Advancement in the field of computer vision has enabled the rapid development of perception systems for autonomous vehicles (AV) in recent years. Optical flow estimation, known as the study of how to estimate per-pixel 2D motion between two temporally consecutive frames, is one of the fundamental problems in computer vision that are widely used in autonomous driving. Specifically, optical flow estimation helps vehicles perceive the temporal continuity of the surrounding environment and hence it plays significant roles in time-series-based tasks such as object tracking [1, 2], visual odometry [3], and SLAM systems [4], to point out a few. Horn and Schunck [5] introduce the first method to compute optical flow through energy minimization and many excellent methods obtain better results based on it. However, the optimizing problem of a complex objective is usually computationally expensive in terms of real-time applications such as AV. To achieve faster and more reliable performance, end-to-end neural networks are proposed [6, 7, 8, 9]. These data-driven learning-based methods are more efficient and robust against challenges, such as occlusions, large displacement and motion blur, that break the brightness constancy and small motion assumptions traditional methods are built upon. Nevertheless, there are still a few unique challenges in AV applications that have been neglected by existing state-of-the-art methods. In this paper, we investigate two commonly encountered challenges among them and propose the

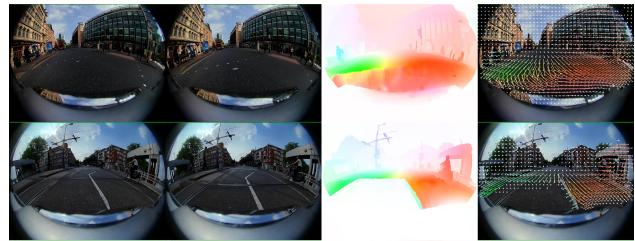


Figure 1: Erroneous optical flow estimation by feeding fisheye images into off-the-shelf RAFT [9]. From left to right in each row: current frame, next frame, color coded result, sparse vector overlay plots for better visualization. Note how the estimated flow vectors on the ground are either missing or inconsistent with the vehicle motion.

solutions respectively: lens distortion and low-light scenes.

Near-field sensing is a ubiquitous topic for automated driving. Some primary use cases are automated parking systems and traffic jam assistance systems. Near-field sensing is usually achieved by building a surround-view system with a number of wide-angle cameras that come with strong radial distortion. For example, fisheye cameras offer a significantly wider field-of-view (FoV) than standard pinhole cameras, and in practice four fisheye cameras located at the front, rear, and on each wing mirror are sufficient to build a surround-view system for a full-size vehicle [10]. Although such fisheye systems are widely deployed, to the best of our knowledge, there is no previous work explicitly handling optical flow estimation on images with strong lens distortion, such as fisheye imagery. As shown in Figure 1, one of the current state-of-the-art methods [9] shows erroneous results when taking in fisheye images from WoodScape [11] due to its focus on narrow field-of-view cameras with mild radial distortion only. An intuitive way to solve this is to correct the distortion in the input images as a preprocessing step before passing through the neural network. However, this inevitably leads to reduced field-of-view and resampling distortion artifacts at the periphery [12]. Without rectification, building an automotive dataset is the major bottleneck in optical flow estimation on fisheye imagery. Very few synthetic datasets provide optical flow ground truth associated with fisheye images [13], whereas no real-world dataset exists with optical flow ground truth. This is due to the fact that per-pixel motion between every two consecutive frames is extremely difficult to be manually labelled. Simulators [14, 15] can readily generate background motion but dynamic foreground objects need to be explicitly taken care of. In this paper, we investigate and boost the performance of RAFT on strongly distorted inputs by mak-

ing use of the only existing dataset with optical flow groundtruth, SynWoodScape [13].

Most AV applications are expected to operate not only during the day but also at night. Cameras become unreliable and camera-based computations are prone to failure under low-light conditions due to its susceptibility to noise and inconsistent exposure. Alternatively, LiDAR sensors can perform robustly in low-light autonomous driving [16] because active sensors that measure the time-of-flight of the emitted lasers are independent of illumination. However, LiDAR is bulky, costly, and requires much more computation as well as memory resources to process the output, which makes it inferior to cameras if the latter can provide equivalently reliable results in low light. Current optical flow methods show poor capabilities of dealing with low-light data because low light is a complex scenario coming with low signal-to-noise ratio, motion blur and local illumination changes brought by multiple light sources. In addition, current optical flow datasets [17, 18, 19] are dominated by daytime images. In this paper, we propose a novel, generic architecture that facilitates learning nighttime-robust representations in a semi-supervised manner, without the help of any extra data or sacrificing the daytime performance. To the best of our knowledge, this is the first learning-based method that explicitly handles optical flow estimation in low light. The main contributions of this paper are:

1. Introduction and investigation of two challenges in optical flow estimation for AV applications: strong lens distortion and low-light scenes.
2. Implementation and improvement of a baseline optical flow algorithm on fisheye inputs and experimental evaluation.
3. Implementation of an effective but also generic framework of novel strategies to learn nighttime-robust representations for learning-based optical flow algorithms.

The paper is organized as follows. Section II discusses related work on optical flow estimation in the automotive industry and existing attempts to solve the two aforementioned challenges. Section III describes the implementation of our proposed flow estimation algorithms for fisheye and low-light inputs respectively, as well as presents the experimental evaluation and results analysis. Finally, Section IV discusses the remaining challenges for flow estimation in AV applications and concludes the paper.

II RELATED WORK

Optical Flow Estimation: Traditional solutions have been studied and adapted for decades [5, 20]. In order to be robust against more challenging open world problems including lack of features, motions in different scales, and occlusions, recent learning-based methods outperform traditional ones. Dosovitskiy et al. [6] propose FlowNetS and FlowNetC, which is a pioneer work in showing the feasibility of directly estimating optical flow given images. Sun et al. [7] design PWC-Net, a much more efficient solution based on pyramidal processing, warping and the use of a cost volume. RAFT [9], proposed by Teed and Deng, demonstrates notable improvement by building multi-scale 4D correlation volumes for all pairs of pixels and iteratively updating flow estimates through refinement module based on gated recurrent units (GRU). All these methods are fully supervised and trained using imagery from a standard pinhole camera. The training data are also collected during the day with sufficient brightness. None of them

pays attention to the performance of optical flow in more challenging AV applications such as strong lens distortion and driving at night, which leads to errors and even catastrophic failures.

Strong Lens Distortion: There is very limited work on perception tasks for strongly distorted images such as fisheye images. Popular approaches include rectifying the radial distortion before passing images into any regular perception pipeline. However, this will inevitably bring reduced field-of-view and resampling distortion artifacts especially at the image borders [12]. Spatially variant distortion that makes closer objects appear larger also poses scaling problems and complexity to geometric perception tasks. Additionally, Rashed et al. [21] show that the common use of bounding boxes for object detection no longer fit well for rectangular objects in distorted images. More sophisticated representations for detected objects, such as a curved bounding box exploiting the known radial distortion, are explored in [22]. Although there is some literature using distorted images without rectification on other perception tasks, such as depth estimation [23, 24], soiling [25], visual odometry [26] and motion segmentation [27], there is no previous work estimating optical flow due to the difficulty in labeling ground truth. WoodScape [11], KITTI 360 [28] and Oxford RobotCar [29] are some well-known autonomous driving datasets containing strongly distorted images such as fisheye images, but none of them has optical flow ground truth. In this paper, we take advantage of the synthetic fisheye dataset published recently, SynWoodScape [13], which is the first dataset providing optical flow for both foreground and background motions by computing it analytically using other data modalities extracted from the simulator. We train our network using synthetic data from SynWoodScape and evaluate it on real-world fisheye data from WoodScape.

Low-Light Scenes: Similar to optical flow estimation on strongly distorted images, there is some work handling low light in a few perception tasks [16, 30, 31] but none of them has proposed an optical flow estimation algorithm that is robust against low-light scenes. Very related to ours, Zheng et al. [32] propose a method to synthesize low-light optical flow data by simulating the noise model on dark raw images, which is then used to finetune an off-the-shelf network. However, their method is not able to synthesize more realistic characteristics of real-world low-light scenes one would observe in AV applications, such as the motion blur and local illumination changes brought by multiple light sources. Their improvement is also very limited due to the off-the-shelf network is not designed nor trained to learn nighttime-robust representations. In addition, a variety of techniques have been developed for low-light image enhancement [33, 34] and image-to-image translation [35, 36]. The former can preprocess inputs to a flow estimation network during inference by brightening up a given low-light image, whereas the latter can translate a daytime image into its nighttime counterpart so as to complement the lack of optical flow datasets in low light [16]. But neither approach facilitates the network training in that the processed data bring in extra complexities such as additional artificial noise, overexposure, or inconsistent image translation across frames. Finally, semi-supervised learning is a common approach to tackling the lack of optical flow data in particular scenarios, where a set of predefined transformations are applied to the original labeled data and the output of the perturbed data are enforced to agree with the out-

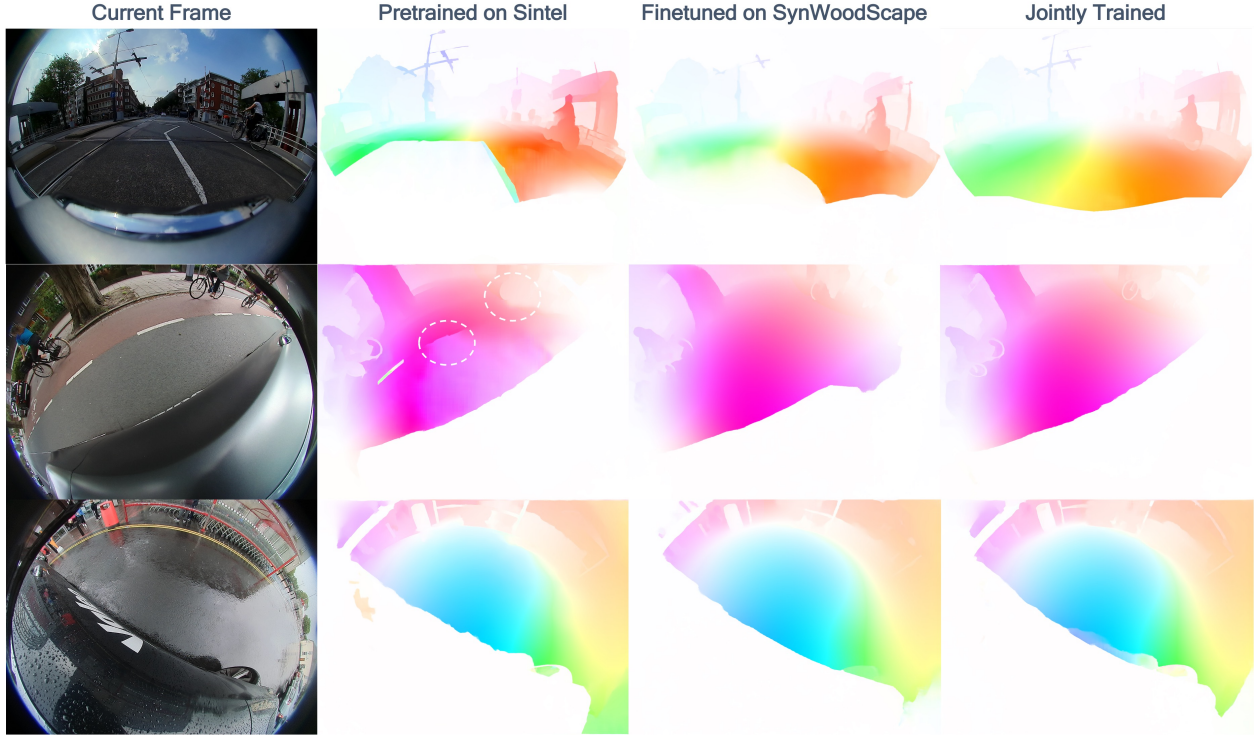


Figure 2: Optical flow estimation (color coded) on real-world automotive data from WoodScape [11]. Input frames are from the fisheye cameras of front view, right-side view, and left-side view respectively.

puts of the original data [37]. For example, Jeong et al. [38] use a semi-supervised setup to impose translation and rotation consistency equivariance for optical flow estimation. Yan et al. [39] synthesize foggy images from clean and labelled images in order to avoid flow estimation errors caused in dense foggy scenes. Similar to these semi-supervised methods, we incorporate low-light consistency that facilitates learning explicit nighttime-robust representations without additional labeling.

III PROPOSED ALGORITHMS AND RESULTS

In this section, we describe the two proposed optical flow estimation algorithms for strongly distorted inputs and low-light inputs respectively. We also present the corresponding experimental evaluation and results analysis.

III. A Strong Lens Distortion

The limited availability of datasets with strong lens distortion is the bottleneck that prevents recent methods from generalizing to more distorted inputs. With the help of SynWoodScape [13], the first fisheye dataset providing optical flow ground truth for both foreground and background motions, we are able to train an optical flow model, using RAFT [9] as the backbone, that generalizes well on strongly distorted lenses without sacrificing its original performance on pinhole cameras.

If we run the off-the-shelf RAFT on real-world fisheye automotive dataset, e.g. WoodScape [11], we will find sharp and inconsistent optical flow estimation, which is especially illustrated on the ground plane in Figure 1. To solve this, we provide

two baselines and their qualitative as well as quantitative evaluation. One is to finetune the pretrained RAFT using SynWoodScape, following the training schedule in Table 1a. The other is to jointly train RAFT on both SynWoodScape and images from pin-hole camera that are regularly used in learning-based optical flow methods [6, 17, 18, 19, 40]. The jointly training baseline follows the training schedule in Table 1b.

Table 1: Details of the training schedule. Column header abbreviations: LR: learning rate, BS: batch size, WD: weight decay, CS: crop size. Training dataset abbreviations: C: FlyingChairs, W: SynWoodScape, S: Sintel, T: FlyingThings3D, K: KITTI-2015, H: HD1K.

(a) Finetuning baseline. During the Sintel stage, the dataset distribution is S(.67), T(.12), K(.13), H(.08).

Stage	Weights	Dataset	LR	BS	WD	CS
Chairs	-	C	4e-4	6	1e-4	[368, 496]
Things	Chairs	T	1.2e-4	3	1e-4	[400, 720]
Sintel	Things	S+T+K+H	1.2e-4	3	1e-5	[368, 768]
Finetune	Sintel	W	1e-4	3	1e-5	[600, 800]

(b) Jointly training baseline. During the Joint stage, the dataset distribution is W(.65), S(.17), T(.13), K(.03), H(.02).

Stage	Weights	Dataset	LR	BS	WD	CS
Chairs	-	C	4e-4	6	1e-4	[368, 496]
Things	Chairs	T	1.2e-4	3	1e-4	[400, 720]
Joint	Things	W+S+T+K+H	1e-4	3	1e-5	[368, 768]

We then show the quantitative results in Table 2. We use the endpoint error (EPE) as the metric, which is the standard error measure for optical flow estimation. It is the Euclidean dis-

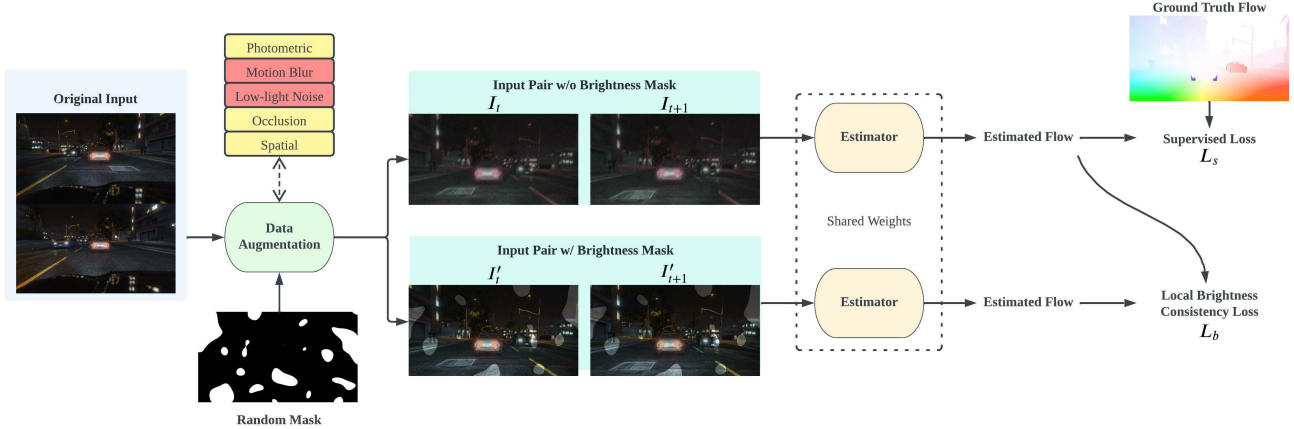


Figure 3: Overview of our proposed framework. During training, the framework takes two consecutive frames as input and passes them through a set of low-light-specific data augmentations as well as applies a random illumination mask. Then the optical flow estimator estimates flow on two pairs of augmented frames in parallel. The network is supervised by two losses: the conventional optical flow loss and the novel brightness consistency loss. During inference, the input frames are directly passed into the estimator which outputs optical flow, as is the standard way in the existing state of the art.

tance between the estimated flow vector and the ground truth, averaged over all pixels. We evaluate the two baselines (Stages "Finetune" and "Joint") described above along with the pretrained model (Stage "Sintel") provided by the author on four hold-out test sets from SynWoodScape, Sintel (clean and final passes), and KITTI. SynWoodScape is the only test set of strongly distorted inputs, while the other three assume a pinhole camera model with very little distortion. Although the pretrained model gives outstanding performance on pinhole cameras, its performance significantly drops on fisheye inputs. Our first baseline, the one finetuned on fisheye images, gives the best result on SynWoodScape but has very poor performance on the others. This matches our expectation because both the pretrained and the finetuned models are trained to the best for pinhole camera and fisheye camera respectively, without taking generalization into account. On the other hand, our second baseline, the jointly trained model, keeps the second best while being very close to the best score on all four datasets. Therefore, jointly training provides a straightforward yet strong baseline that generalizes well over lenses with distinct distortions.

Table 2: Endpoint-error results on datasets with various distortion.

Stage	SynWoodScape	Sintel - Clean	Sintel - Final	KITTI
Sintel	5.12	1.94	3.18	5.10
Finetune	1.40	5.44	10.32	14.34
Joint	1.48	2.44	4.14	7.31

In Figure 2, we further show their qualitative results on WoodScape that support the improvements we obtain by jointly training RAFT on a mixture of lens distortions. In the front view case, note how the jointly trained model is able to consistently estimate the flow on the ground as is the major failure of recent methods shown in Figure 1. The results on side-view cameras also show the jointly trained model captures finer details than its finetuned counterpart. For example in the right-side view, not only the inconsistency on the ground is solved, but optical flow associated with the bicycle wheel in the upper right corner is also clearly estimated. In the left-side view, the finetuned model misses the flow associated with the vehicle's front wheel, which is captured by the pretrained model, but the jointly trained model "regains" such detailed estimations. In other words, the finetuned model es-

timates more consistent optical flow, which poses a challenge to the pretrained model due to markedly different projection geometries between fisheye and pinhole cameras, but in return, it loses some details observed by the pretrained model because interesting local features become much less significant given the strong lens distortion. However, the jointly trained model achieves a great trade-off among the previous two: it re-captures the details locally while maintaining good performance globally across different camera views.

III. B Low-Light Scenes

We propose a novel and generic semi-supervised framework that significantly boosts performances of existing state-of-the-art methods in low light conditions. Figure 3 shows the architecture of the framework. The benefit of our framework is three-fold. First, it is independent from the design of the existing methods, so one can apply it generically to an estimator of his choice (e.g. [6, 7, 9, 41]) and augment its nighttime performance out of the box. Second, semi-supervised learning does not require any extra data as the labeling cost for nighttime optical flow datasets is immense. Lastly, it maintains the estimator's competitive performance on the original daytime data without making any trade-off.

We first break down the root causes of failures in optical flow estimation under low light and then describe our proposed strategies in the framework that address these root causes accordingly:

1. the complex noise model of images captured at night,
2. severe motion blur caused by longer exposure time,
3. inconsistent local brightness brought by multiple independent light sources in the scene.

Images captured in low light tend to have more complex noises than those captured with sufficient ambient light. Such noises are never synthesized in the data augmentation step by existing methods, which is the first reason why their optical flow estimators fail in low light. Similar to [32], we decompose the noise model in low light as an aggregate of the photon shot noise and thermal noise. The former is due to the changing amount of photons hitting the sensor with different exposure levels and pixel locations. The photon shot noise is approximated by a Pois-



Figure 4: Effects of low-light noise augmentation and motion blur augmentation.

son distribution. Thermal noise refers to the noise in readout circuitry in the sensor and is approximated by a Gaussian distribution. Therefore, we synthesize the low-light noise onto the input frames as one extra data transform in the data augmentation step. Specifically, we sample Poisson and Gaussian parameters, (a, b) , from ranges observed in real-world low-light images, formulate it into a single heteroscedastic Gaussian (Equation 1), and apply it to an input frame I . With probability 0.5, the low-light noise augmentation is performed on each pair of consecutive frames.

$$I(x) = \mathcal{N}(\mu = x, \sigma^2 = ax + b) \quad (1)$$

Motion blur is another root cause we need to address when estimating optical flow in low light. In order to mimic the blurring effects caused by longer exposure length, we generate authentic motion blur kernels using Point Spread Functions (PSF) at different kernel sizes and intensities. The intensity determines how non-linear and shaken the motion blur looks. Similar to low-light noise, we apply the authentic blurring to a pair of input frames as one extra data augmentation, with probability 0.6. An illustration of the two introduced data augmentation strategies is shown in Figure 4.

Inconsistent local brightness is the last but not least root cause. This is due to multiple independent lighting sources existing in a low-light scene (street light, headlight, moonlight, etc.), which leads to uneven bright areas in an image. For example, the ground plane in the original input in Figure 3 is illuminated only in front of vehicles' headlights but remains dark elsewhere. Unlike in the daytime where sun is the dominant lighting source, images captured at night have inconsistent local brightness even on the same object. Because optical flow is estimated by matching pixels across two images, such inconsistencies cause existing methods to fail easily. For instance, in Figure 5, the first row shows the catastrophic failure of RAFT when a pedestrian walks from the dark into the vehicle headlight and his illumination changes drastically across frames. In order to resolve this, we resort to semi-supervised learning. Similar to [38], we also adopt the cow-mask [42] to create sufficiently random yet locally connected illumination patterns as the inconsistent local brightness occurs in any size, shape and position in images while exhibiting locally explainable structures, depending on the driving environment and the time. We apply the same binary mask to the original pair of input frames and randomly adjust the brightness of pixels according to the mask. The true area of the binary mask is

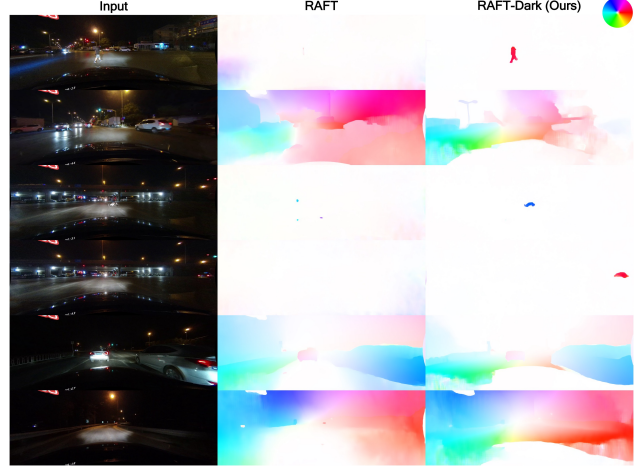


Figure 5: Optical flow estimation on low-speed sequences from CULane [43].

uniformly sampled from 40% to 70% of the image. With a probability of 0.5, we increase the absolute brightness of the true area, whereas in the remaining time we increase the brightness of the false area. Finally, we introduce the local brightness consistency regularization. We use (I'_t, I'_{t+1}) and (I_t, I_{t+1}) to denote the input pair after data augmentation with and without applying a random brightness mask. Both passes in Figure 3 are independent except that the spatial transform is shared in order to keep the same cropped areas for consistency loss calculation. The local brightness consistency loss is calculated as follows

$$L_b = \|\text{Estimator}(I_t, I_{t+1}) - \text{Estimator}(I'_t, I'_{t+1})\|_2^2. \quad (2)$$

This regularization explicitly constrains the network to output consistent optical flow on (I'_t, I'_{t+1}) as on (I_t, I_{t+1}) , which enforces illumination invariance between the estimated optical flow for the original pair the estimated optical flow for the randomly transformed pair. Note how this semi-supervised approach is different from simply adjusting brightness randomly as another data augmentation scheme, which expand training samples without imposition of a sophisticated consistency loss during training.

We choose RAFT [9] as the estimator and we supervise our network on the aggregated loss $L = L_s + L_b$. L_s is the l_1 distance between the predicted flow $\tilde{f}^i(I_t, I_{t+1})$ and ground truth flow f_t over all iterations i , as in [9]:

$$L_s = \sum_{i=1}^N \gamma^{N-i} \|\tilde{f}^i(I_t, I_{t+1}) - f_t\|_1. \quad (3)$$

Due to the lack of nighttime data with optical flow ground truth, we are restricted to qualitatively evaluating our approach, which we call RAFT-Dark for short. We use CULane [43], a large automotive dataset containing a lot of challenging real-world low light sequences. In Figure 5, we show the comparison between vanilla RAFT and RAFT-Dark on some low-speed sequences. RAFT-Dark demonstrates superior performance to RAFT. In the first, third and fourth rows, RAFT-Dark is able to detect motions associated with pedestrians and vehicles that either experience some drastic illumination change or appear to be too dark and noisy. In the other cases, note how RAFT-Dark gives a signif-

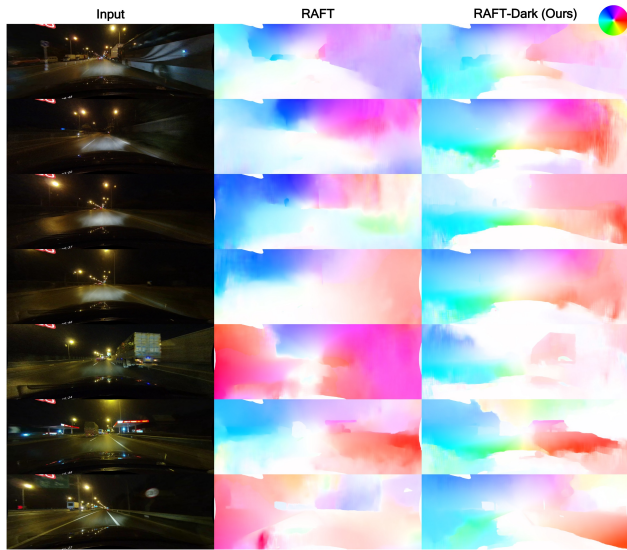


Figure 6: Optical flow estimation on high-speed sequences from CULane [43].

icantly better estimation on the ground plane as well as the directions and magnitudes that are consistent with the ego vehicle’s motion. For convenience, a color coding wheel to visualize per-pixel optical flow vectors is attached to the top right corner: color denotes direction of the flow vector while intensity denotes length of the displacement. Since the ego vehicle always heads forward, the ground truth optical flow vectors in the front camera’s image should intuitively point to the image boundaries and away from the image center. And due to the motion parallax, one should expect larger magnitudes of flow vectors toward the image boundaries and small magnitudes around the image center. In other words, although we have no access to numerical ground truth flow, we know the color coded ground truth should exhibit the same pattern as the color wheel: bluish or greenish on the left side while reddish or yellowish on the right side of the image. With this in mind, RAFT fails to estimate correct optical flow consistent with the vehicle’s motion, especially in background areas such as the ground plane. On the other hand, RAFT-Dark not only performs well on these areas but also learns to separate the dark sky in some cases and to capture details such as the street light in the second row. Such improvements are further illustrated in high-speed sequences from CULane in Figure 6.

Our framework of learning strategies enables RAFT to improve estimation accuracy by more than 50% on average (based on visual observations), and even solves some catastrophic failures. Although we show our results based on RAFT as the estimator, our framework is generic and one can replace RAFT with any existing state-of-the-art method of one’s choice.

III. C Discussion

The goal of this work is to emphasize the importance of addressing optical flow challenges which are not well explored in automated driving. We investigate two of them and propose our solutions accordingly, but the others require further research. Lack of data tends to be the major bottleneck for most data-driven optical flow algorithms. We are able to leverage synthetic data to

improve existing methods’ adaptation of various lens distortions but the sim-to-real gap still exists when these methods are evaluated on real world fisheye data. Optical flow in low light cannot be addressed in the same way without any synthetic data available. We experiment image enhancement prior to the network inference, but it leads to even worse results because enhancement happens per frame rather than per pair of frames and temporal consistency is easily broken. Without any extra data, our approach takes full advantage of publicly available data and simulates three root causes through novel data augmentation schemes and semi-supervised learning. However, low light is merely one of many scenarios that make optical flow estimation harder. Others include foggy, rainy or snowy weather [44]. A unified and robust approach aiming for all these cases is encouraged and we see it also as an opportunity for further investigation by the community.

IV Conclusion

Both lens distortion and low light are important problems for higher levels of automated driving, but they are not explored in detail in the optical flow community as there is no public dataset available. Thus we propose our approaches to these two respectively. We implement and improve a state-of-the-art optical flow algorithm by training it on synthetic fisheye data and demonstrating its adaptation to real-world distorted images as well as generalizability over various lens distortions. We implement a novel, generic framework that facilitates learning nighttime-robust representations in a semi-supervised manner, which shows superior performance to the existing state of the art. In future work, we plan to integrate our current solutions into higher-level pipelines as well as explore other unique challenges of optical flow estimation in the context of automated driving.

REFERENCES

- [1] K. Kale, S. Pawar, and P. Dhulekar, “Moving object tracking using optical flow and motion vector estimation,” in *2015 4th international conference on reliability, infocom technologies and optimization (ICRITO)(trends and future directions)*, pp. 1–6, IEEE, 2015.
- [2] H. Zhou, B. Ummenhofer, and T. Brox, “Deeptram: Deep tracking and mapping,” in *Proceedings of the European conference on computer vision (ECCV)*, pp. 822–838, 2018.
- [3] S. Wang, R. Clark, H. Wen, and N. Trigoni, “Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks,” in *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 2043–2050, IEEE, 2017.
- [4] Z. Teed and J. Deng, “Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 16558–16569, 2021.
- [5] B. K. Horn and B. G. Schunck, “Determining optical flow,” *Artificial intelligence*, vol. 17, no. 1-3, pp. 185–203, 1981.
- [6] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. Van Der Smagt, D. Cremers, and T. Brox, “Flownet: Learning optical flow with convolutional networks,” in *Proceedings of the IEEE international conference on computer vision*, pp. 2758–2766, 2015.
- [7] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, “Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8934–8943, 2018.

- [8] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, "Models matter, so does training: An empirical study of cnns for optical flow estimation," *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 6, pp. 1408–1423, 2019.
- [9] Z. Teed and J. Deng, "Raft: Recurrent all-pairs field transforms for optical flow," in *European conference on computer vision*, pp. 402–419, Springer, 2020.
- [10] V. R. Kumar, C. Eising, C. Witt, and S. Yogamani, "Surround-view fisheye camera perception for automated driving: Overview, survey and challenges," *arXiv preprint arXiv:2205.13281*, 2022.
- [11] S. Yogamani, C. Hughes, J. Horgan, G. Sistu, P. Varley, D. O'Dea, M. Uricár, S. Milz, M. Simon, K. Amende, *et al.*, "Woodscape: A multi-task, multi-camera fisheye dataset for autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9308–9318, 2019.
- [12] V. R. Kumar, M. Klingner, S. Yogamani, M. Bach, S. Milz, T. Fingscheidt, and P. Mäder, "Svdistnet: Self-supervised near-field distance estimation on surround view fisheye cameras," *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [13] A. R. Sekkat, Y. Dupuis, V. R. Kumar, H. Rashed, S. Yogamani, P. Vasseur, and P. Honeine, "Synwoodscape: Synthetic surround-view fisheye camera dataset for autonomous driving," *IEEE Robotics Automation Letters*, 2022.
- [14] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*, pp. 1–16, PMLR, 2017.
- [15] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "Airsim: High-fidelity visual and physical simulation for autonomous vehicles," in *Field and service robotics*, pp. 621–635, Springer, 2018.
- [16] H. Rashed, M. Ramzy, V. Vaquero, A. El Sallab, G. Sistu, and S. Yogamani, "Fusmodnet: Real-time camera and lidar based moving object detection for robust low-light autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pp. 0–0, 2019.
- [17] M. Menze and A. Geiger, "Object scene flow for autonomous vehicles," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3061–3070, 2015.
- [18] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox, "A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4040–4048, 2016.
- [19] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black, "A naturalistic open source movie for optical flow evaluation," in *European conference on computer vision*, pp. 611–625, Springer, 2012.
- [20] T. Brox, A. Bruhn, N. Papenberger, and J. Weickert, "High accuracy optical flow estimation based on a theory for warping," in *European conference on computer vision*, pp. 25–36, Springer, 2004.
- [21] H. Rashed, E. Mohamed, G. Sistu, V. R. Kumar, C. Eising, A. El-Sallab, and S. Yogamani, "Generalized object detection on fisheye cameras for autonomous driving: Dataset, representations and baseline," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 2272–2280, 2021.
- [22] H. Rashed, E. Mohamed, G. Sistu, C. E. Ganesh, A. El-Sallab, and S. Yogamani, "Fisheyeyolo: Object detection on fisheye cameras for autonomous driving," in *Machine Learning for Autonomous Driving NeurIPS 2020 Virtual Workshop*, vol. 8, 2020.
- [23] V. R. Kumar, S. Milz, C. Witt, M. Simon, K. Amende, J. Petzold, S. Yogamani, and T. Pech, "Near-field depth estimation using monocular fisheye camera: A semi-supervised learning approach using sparse lidar data," in *CVPR Workshop*, vol. 7, p. 2, 2018.
- [24] V. R. Kumar, S. A. Hiremath, M. Bach, S. Milz, C. Witt, C. Pinard, S. Yogamani, and P. Mäder, "Fisheye distance net: Self-supervised scale-aware distance estimation using monocular fisheye camera for autonomous driving," in *2020 IEEE international conference on robotics and automation (ICRA)*, pp. 574–581, IEEE, 2020.
- [25] M. Uricár, J. Ulicny, G. Sistu, H. Rashed, P. Krizek, D. Hurych, A. Vobecky, and S. Yogamani, "Desoiling dataset: Restoring soiled areas on automotive fisheye cameras," in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pp. 0–0, 2019.
- [26] P. Liu, L. Heng, T. Sattler, A. Geiger, and M. Pollefeys, "Direct visual odometry for a fisheye-stereo camera," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1746–1752, IEEE, 2017.
- [27] V. R. Kumar, S. Yogamani, H. Rashed, G. Sistu, C. Witt, I. Leang, S. Milz, and P. Mäder, "Omnidet: Surround view cameras based multi-task visual perception network for autonomous driving," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2830–2837, 2021.
- [28] Y. Liao, J. Xie, and A. Geiger, "Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [29] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The oxford robotcar dataset," *The International Journal of Robotics Research*, vol. 36, no. 1, pp. 3–15, 2017.
- [30] H. Alismail, M. Kaess, B. Browning, and S. Lucey, "Direct visual odometry in low light using binary descriptors," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 444–451, 2016.
- [31] J. Ge, Y. Luo, and G. Tei, "Real-time pedestrian detection and tracking at nighttime for driver-assistance systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 10, no. 2, pp. 283–298, 2009.
- [32] Y. Zheng, M. Zhang, and F. Lu, "Optical flow in the dark," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6749–6757, 2020.
- [33] C. Chen, Q. Chen, J. Xu, and V. Koltun, "Learning to see in the dark," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3291–3300, 2018.
- [34] F. Lv, F. Lu, J. Wu, and C. Lim, "Mbllen: Low-light image/video enhancement using cnns," in *British Machine Vision Conference (BMVC)*, 2018.
- [35] M.-Y. Liu, T. Breuel, and J. Kautz, "Unsupervised image-to-image translation networks," *Advances in neural information processing systems*, vol. 30, 2017.
- [36] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz, "Multimodal unsupervised image-to-image translation," in *Proceedings of the European conference on computer vision (ECCV)*, pp. 172–189, 2018.
- [37] S. Laine and T. Aila, "Temporal ensembling for semi-supervised learning," *arXiv preprint arXiv:1610.02242*, 2016.
- [38] J. Jeong, J. M. Lin, F. Porikli, and N. Kwak, "Imposing consistency for optical flow estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3181–3191, 2022.
- [39] W. Yan, A. Sharma, and R. T. Tan, "Optical flow in dense foggy scenes using semi-supervised learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13259–13268, 2020.

- [40] D. Kondermann, R. Nair, K. Honauer, K. Krispin, J. Andrulis, A. Brock, B. Gussfeld, M. Rahimimoghaddam, S. Hofmann, C. Brenner, *et al.*, “The hci benchmark suite: Stereo and flow ground truth with uncertainties for urban autonomous driving,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 19–28, 2016.
- [41] H. Xu, J. Zhang, J. Cai, H. Rezatofighi, and D. Tao, “Gmflow: Learning optical flow via global matching,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8121–8130, 2022.
- [42] G. French, A. Oliver, and T. Salimans, “Milking cowmask for semi-supervised image classification,” *arXiv preprint arXiv:2003.12022*, 2020.
- [43] X. Pan, J. Shi, P. Luo, X. Wang, and X. Tang, “Spatial as deep: Spatial cnn for traffic scene understanding,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, 2018.
- [44] M. M. Dhananjaya, V. R. Kumar, and S. Yogamani, “Weather and light level classification for autonomous driving: Dataset, baseline and active learning,” in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pp. 2816–2821, IEEE, 2021.

AUTHORS BIOGRAPHY

Shihao Shen is a second-year graduate student in the Robotics Institute at Carnegie Mellon University and expects to receive his M.Sc. in Robotic Systems Development in 2023. He worked as an Interim Engineering Intern in the Multimedia Research and Development department at Qualcomm in summer 2022 and this

is his work done during his internship. His main research focus is machine learning with applications in computer vision as well as simultaneous localization and mapping (SLAM).

Louis Kerofsky is researcher in video compression, video processing and display. He received M.S. and Ph.D. degrees in Mathematics from the University of Illinois, Urbana-Champaign (UIUC). He has over 20 years of experience in research and algorithm development and standardization of video compression. He has served as an expert in the ITU and ISO video compression standards committees. He is an author of over 40 publications which have over 5000 citations. He is an inventor on over 130 issued US patents. He is a senior member of IEEE, member of Society for Information Display.

Senthil Yogamani is an artificial intelligence architect for autonomous driving and holds a principal engineer position at Qualcomm. He leads the research and design of AI algorithms for various modules of autonomous driving systems. He has over 17 years of experience in computer vision and machine learning including 14 years of experience in industrial automotive systems. He is an author of 110+ publications which have 4000+ citations and 100+ inventions with 75 filed patent families. He serves on the editorial board of various leading IEEE automotive conferences including ITSC and IV and advisory board of various industry consortia including Khronos, Cognitive Vehicles and IS Auto. He is a recipient of the best associate editor award at ITSC 2015 and best paper award at ITST 2012.