ATG-PVD: Ticketing Parking Violations on A Drone

Hengli Wang^{1*}, Yuxuan Liu^{1*}, Huaiyang Huang^{1*}, Yuheng Pan^{2*},
Wenbin Yu², Jialin Jiang², Dianbin Lyu², Mohammad J. Bocus³,
Ming Liu¹, Ioannis Pitas⁴, and Rui Fan^{2,5}✉

¹ HKUST Robotics Institute

{hwangdf, yliuhb, hhuangat, eelium}@ust.hk

² ATG Robotics

{panyuheng, yuwenbin, jiangjialin, lvdianbin}@atg-itech.com

³ University of Bristol

junaaid.bocus@bristol.ac.uk

⁴ Aristotle University of Thessaloniki

pitass@csd.auth.gr

⁵ UC San Diego

rui.fan@ieee.org

Abstract. In this paper, we introduce a novel suspect-and-investigate framework, which can be easily embedded in a drone for automated parking violation detection (PVD). Our proposed framework consists of: 1) SwiftFlow, an efficient and accurate convolutional neural network (CNN) for unsupervised optical flow estimation; 2) Flow-RCNN, a flow-guided CNN for car detection and classification; and 3) an illegally parked car (IPC) candidate investigation module developed based on visual SLAM. The proposed framework was successfully embedded in a drone from ATG Robotics. The experimental results demonstrate that, firstly, our proposed SwiftFlow outperforms all other state-of-the-art unsupervised optical flow estimation approaches in terms of both speed and accuracy; secondly, IPC candidates can be effectively and efficiently detected by our proposed Flow-RCNN, with a better performance than our baseline network, Faster-RCNN; finally, the actual IPCs can be successfully verified by our investigation module after drone re-localization.

Dataset and Demo Video:

sites.google.com/view/atg-pvd

1 Introduction

We are currently experiencing an unprecedented crisis due to the ongoing Coronavirus Disease 2019 (COVID-19) pandemic. Its worldwide escalation has taken us by surprise, causing major disruptions to global health, economic and social

* Equal contributions.

systems. Indeed, our lives have changed overnight – businesses and schools are closed, most employees are working from home, and many have found themselves without a job. Millions of people across the globe are confined to their homes, while healthcare workers are at the frontline of the COVID-19 response [1]. With the increase in COVID-19 cases, public transport use has plummeted, as commuters shun buses, trams, and trains in favor of private cars and taxis. For instance, USA Today reported that the transit ridership demand in April 2020 was down by about 75% nationwide, compared to normal, with figures of 85% in San Francisco, 67% in Detroit and 60% in Philadelphia [2].

With the increasing number of vehicles on the roads, parking spaces have become scarce and many vehicles are parked just by the roadside, which in turn results in a significant increase in parking violations. In late March 2020, the Department of Transportation in Los Angeles [3] announced relaxed parking enforcement regulations as part of the emergency response to COVID-19, so that their citizens could practice safe social distancing without being concerned about a ticket. As the Return-to-Work Plan progresses, the relaxed parking enforcement regulations are no longer in force, consequently increasing the workload of the local traffic law enforcement officers. The demand for automated and intelligent parking violation detection (PVD) systems has thus become greater than ever.

The existing automated PVD systems typically recognize illegally parked cars (IPCs) by analyzing the videos acquired by closed-circuit televisions (CCTVs) through 2D/3D object detection algorithms [4] or video surveillance analysis algorithms [5]. However, the efficiency of such methods relies on CCTV camera positions, as IPCs cannot always be detected, especially if they are at a distant location. Deploying more CCTVs can definitely minimize misdetections, but this will also incur a high cost, and/or may not be practical. Therefore, many researchers have turned their focus towards mobile PVD systems, which can be mounted on any vehicle type. For example, the Birmingham City Council in England utilizes surveillance cars to detect IPCs and record their plate numbers [6]. However, such surveillance cars are expensive and typically require drivers. Therefore, autonomous machines, especially drones, have emerged as more efficient and cheaper alternatives.

The cars in the street can be grouped into three categories: 1) moving cars (MCs), 2) legally parked cars (LPCs) and 3) IPCs. MCs can be distinguished from LPCs and IPCs using dynamic object detection techniques, such as optical flow analysis, while IPCs can be distinguished from LPCs using object detection networks, such as Faster-RCNN [7], with the assistance of parking spot information. In this paper, we introduce a novel *suspect-and-investigate* PVD system (see Fig. 1) embedded in a drone. In the suspicion phase, we first employ a novel unsupervised optical flow estimation network, referred to as *SwiftFlow*, to estimate the optical flow F_t between I_t and I_{t+1} . F_t is then incorporated into a novel object detection and classification network, referred to as *Flow-RCNN*, to detect cars and classify them into MCs, LPCs and IPC candidates. A visual simultaneous localization and mapping (VSLAM) module then builds a localizable map

Fig. 1: The framework of our proposed suspect-and-investigate PVD system: the first phase identifies suspected IPC candidates, and the second phase investigates the suspected IPC candidates and issues tickets to the actual IPCs. The frame I_t in the suspicion phase corresponds to the frame I'_t in the investigation phase.

containing the suspected IPC candidates. After a parking grace period (which is typically five minutes) has elapsed, the drone flies back to the same location. The VSLAM module in the investigation phase subsequently detects loop closure and re-localizes the drone in the pre-built map. Finally, the suspected IPC candidates are re-identified, and the actual IPCs are marked in the map. Our main contributions are summarized as follows:

- A novel suspect-and-investigate PVD framework;
- SwiftFlow, a novel unsupervised optical flow estimation network;
- Flow-RCNN, a novel car detection and classification network;
- A large-scale PVD dataset, published for research purposes.

2 Related Work

2.1 Optical Flow Estimation

Traditional approaches generally formulate optical flow estimation as a global energy minimization problem [8,9,10,11]. Recently, convolutional neural net-

works (CNNs) have achieved impressive performance in optical flow estimation. FlowNet [12] was the pioneering work in end-to-end deep optical flow estimation. Its key component is a so-called correlation layer, which can provide explicit matching capabilities. Later methods, PWC-Net [13] and LiteFlowNet [14] introduced the popular coarse-to-fine architecture, which provides a good trade-off between optical flow accuracy and computation efficiency. Meanwhile, IRR-PWCNet [15] demonstrates that occlusion prediction integrated into optical flow estimation can effectively enhance the optical flow estimation accuracy.

Although the aforementioned supervised optical flow estimation methods perform impressively, they generally require a large amount of optical flow ground truth to learn the best solution. Acquiring such ground truth, especially for real-world datasets, is extremely time-consuming and labor-intensive, making these supervised approaches difficult to apply in real-world applications. For these reasons, unsupervised learning has recently become the preferred technique for such applications. For instance, DSTFlow [16] employs a photometric loss and a smooth loss in CNN training, which are similar to the global energy used in traditional methods. Additionally, some methods, such as UnFlow [17], DDFlow [18] and SelFlow [19] integrate occlusion reasoning into unsupervised optical flow estimation frameworks to further improve their accuracy. However, such approaches are typically computationally intensive, and they are difficult to embed in a drone.

2.2 Object Detection

Discovering objects and their locations in images is still a challenging problem in computer vision. Due to their promising results, CNNs have emerged as a powerful tool for object detection. The modern deep object detection algorithms can be grouped into two main types: a) anchor-based and b) anchor-free.

Anchor-based methods predict bounding boxes based on initial guesses. According to the pipelines and primary proposal sources, they can be further categorized as either one-stage or two-stage methods. The former make predictions directly from hand-crafted anchors. For example, RetinaNet [20] employs a feature pyramid network (FPN) to produce dense predictions at multiple scales. On the other hand, the two-stage methods make predictions using the proposals produced by a one-stage detector. For instance, Fast-RCNN [21] and Faster-RCNN [7] perform cropping and resizing on images or feature maps, according to the bounding box proposals. The RCNN branch in Faster-RCNN utilizes a field of view (FOV), that is larger than the bounding box proposals, so as to extract regions of interest (RoIs) directly from the feature maps.

Anchor-free methods usually do not rely on human-designed region proposals to bootstrap the detection process. For example, CornerNet [22] translates the object detection problem into a keypoint detection and matching problem, where specially-designed pooling layers construct biased receptive fields for corner point detection. CenterNet [23], which is based on CornerNet [22], utilizes two customized modules: a) cascade corner pooling and b) center pooling, to

enrich information collected by both the top-left and bottom-right corners. It detects each object as a triplet, rather than a pair, of keypoints.

In recent years, incorporating additional visual information, such as semantic predictions, into object classification is becoming an increasingly ubiquitous part of object detection. Since MCs can be easily distinguished from optical flow images, we incorporate the latter into our framework to improve IPC candidate detection.

2.3 VSLAM

Traditional VSLAM approaches leverage visual features and the geometric relations between multiple views of a 3D scene (typically known as multi-view geometry) to estimate camera poses and construct/update a map of the 3D scene. The state-of-the-art VSLAM approaches are classified as either indirect [24,25,26] or direct [27,28,29]. Both types extract visual features from images and associate them with descriptors. However, the indirect methods sample corners and associate them with higher dimensional descriptors, while the direct methods typically sample pixels with a relatively large local intensity gradient and associate them with a patch of pixels surrounding their sampled location. Furthermore, these two types of methods typically minimize different objective functions: the indirect methods resort to geometric residuals, whereas the direct methods resort to photometric residuals.

In order to combine the advantages of these two types of methods, Froster *et al.* [30] proposed semi-direct visual odometry (SVO), which tracks camera poses via sparse image alignment and utilizes hierarchical bundle adjustment (BA) as the back-end to optimize the geometry structure and camera motion. Furthermore, many researchers have integrated other computer vision tasks, such as 2D object detection [31,32,33], instance segmentation [34,35] and flow/depth prediction [36,37], into their SLAM systems, so as to address the problem of the existence of dynamic objects, by exploiting high-level semantic information. For example, Huang *et al.* [32] proposed ClusterVO, which uses a multi-level probabilistic association scheme to both track low-level visual features and realize high-level object detection. Moreover, Yang *et al.* [31] introduced CubeSLAM, which performs single image 3D cuboid object detection, together with multi-view object SLAM.

3 ATG-PVD Framework

3.1 SwiftFlow

Since our proposed SwiftFlow network is based on the pipeline of PWC-Net [13], we first provide readers with some preliminaries about the latter. In PWC-Net [13], feature maps are first extracted from video frames using a Siamese pyramid network. Then, the feature map x_{t+1}^l of the $(t+1)$ -th video frame at level l is aligned with the feature map x_t^l of the t -th video frame at level l

Fig. 2: The decoder architecture of our proposed SwiftFlow. The pipeline of two adjacent levels in the decoder are displayed for simplicity.

via a warping operation based on the upsampled flow prediction F_t^{l+1} at level $l+1$. A correlation layer is then employed to compute the cost volume, which is subsequently concatenated with x_t^l as well as the upsampled flow prediction F_t^{l+1} at level $l+1$. Finally, the flow residual, predicted by the flow estimation module, is combined with the upsampled flow prediction F_t^{l+1} at level $l+1$ using an element-wise summation to generate the flow prediction F_t^l at level l . We iterate this process and obtain the flow predictions at different scales.

SwiftFlow improves on PWC-Net [13] in terms of computational efficiency, so that it can perform in real time on a drone. The decoder in PWC-Net [13] has too many parameters, so we make three major modifications to the decoder architecture (see Fig. 2) to minimize the model size and improve accuracy. As the decoder in PWC-Net [13] employs a dense connection scheme in each pyramid level, making the network computationally intensive, SwiftFlow establishes connections only between two adjacent levels, which can reduce the number of network parameters by 50%. Furthermore, the optical flow estimation modules at different pyramid levels of PWC-Net [13] have different learnable weights to estimate optical flow residuals. Considering that the optical flow estimation modules at different levels have the same functionality and the optical flow residuals at different levels have similar value ranges, we believe sharing the weights of optical flow estimation modules at all pyramid levels can be a more effective and efficient strategy. We also add an additional convolutional layer before the optical flow estimation module at each level for feature map alignment. Moreover, we notice that the warping operation can induce ambiguity to occluded areas, which breaks correlation layer symmetry. We propose to add an asymmetric layer before the correlation layer to alleviate this problem and improve optical flow estimation accuracy. Therefore, we replace the warping operation with a deformable convolutional layer [38], as shown in Fig. 2.

Fig. 3: Challenging cases for parked car detection and classification.

Referring to the commonly applied unsupervised training strategy, we train SwiftFlow by minimizing the following weighted sum of losses:

$$L = \lambda_{\text{photo}} \cdot L_{\text{photo}} + \lambda_{\text{smooth}} \cdot L_{\text{smooth}} + \lambda_{\text{self}} \cdot L_{\text{self}}, \quad (1)$$

where L_{photo} is the photometric loss that considers an occlusion-aware mask [39], L_{smooth} is the smoothness regularization [40], and L_{self} is the self-supervision Charbonnier loss [18]. Following the instructions in [41], we set $\lambda_{\text{photo}} = 1$ and $\lambda_{\text{smooth}} = 2$ in our experiments. Moreover, we use $\lambda_{\text{self}} = 0$ for the first 50% of training steps, and increase it to 0.3 linearly for the next 10% of training steps, after which it stays at a constant value.

3.2 Flow-RCNN

Given an RGB video frame and its corresponding estimated optical flow, the proposed Flow-RCNN detects cars in the video frame and classifies them into MCs, LPCs, and IPC candidates.

Judging whether a car is legally parked is very challenging. Intuitively, we can resort to the parking spot delimitation lines, which are typically painted in white. However, in real-world environments, methods that rely solely on the parking spot information may fail. For instance, in Fig. 3(a), the white car is not parked entirely within the designated parking spot; in Fig. 3(b), only parts of the white car and parking spot appear; and in Fig. 3(c) and Fig. 3(d), the parking spots are not enclosed. Moreover, parking spots are not always bounded by rectangular line markings, as illustrated in Fig. 3(c). It is challenging to design a rule-guided method to solve for these cases, even with perfectly labeled cars and parking spots. Furthermore, various tall objects, such as light poles and trees, often present salient optical flow estimations. In this case, the methods that rely entirely on optical flow information can wrongly characterize an IPC/LPC as an MC. Therefore, an end-to-end, optical flow-guided, and detect-and-classify architecture for IPC candidate detection provides a better alternative.

The architecture of our proposed Flow-RCNN is illustrated in Fig. 4. It incorporates the optical flow information, obtained by SwiftFlow in Section 3.1, into the conventional Faster-RCNN [7] architecture for IPC candidate detection, and it outputs the position and category (MC, LPC or IPC candidate) of each car in the video frame in an end-to-end manner. The RGB video frame is first passed through a backbone CNN to produce multi-scale feature maps y^i . The features extracted from the optical flow image then dynamically weigh the activation of

Fig. 4: Flow-RCNN architecture. The optical flow image, obtained from Swift-Flow, is fed into multiple convolutional layers. The optical flow features then dynamically weigh each element in the multi-scale feature maps extracted from the RGB image.

each element in the multi-scale feature maps y^i , which enables the detector to focus more on MCs. We then fuse the multi-scale feature maps to produce a feature pyramid for the subsequent region proposal network (RPN) and RCNN heads [7]. Since our dataset is highly imbalanced (see Fig. 7), *i.e.*, most vehicles are regarded as IPC candidates or IPCs, we apply focal loss [20] to mitigate the class imbalance problem in the classification stage.

3.3 Mapping, Re-Localization and Re-Identification

Given RGB images and the corresponding detected IPC candidates, our next target is to build a 3D map, investigate each IPC candidate and mark it in the map. To this end, we develop a mapping, re-localization and re-identification module, as illustrated in Fig. 5, on top of ORB-SLAM2 [42].

Our proposed system applies a suspect-and-investigate scheme to mark IPCs in 3D. In the suspicion phase, we leverage ORB-SLAM2 [42] to build a 3D localizable map and mark the detected IPC candidates in the map. Given an RGB image containing detected IPC candidates, the system first extracts ORB [43] features $\{\mathbf{u}_0, \dots, \mathbf{u}_t\}$ and associates them with 2D bounding boxes $\{\mathcal{B}_0^{2D}, \dots, \mathcal{B}_h^{2D}\}$. We explicitly exclude the ORB features extracted from MCs in the subsequent procedures, *i.e.*, tracking and mapping. The rest of the features are then matched with the 3D keypoints $\{\mathbf{x}_0, \dots, \mathbf{x}_m\}$ in the map. With these 3D-2D correspondences $\mathcal{K} \doteq \{(i_k, j_k)\}_{k=1:N}$, the current camera pose $\mathbf{T} = [\mathbf{R}, \mathbf{t}]$ is estimated in a perspective-n-point (PnP) scheme by minimizing the reprojection error as follows [42]:

$$\mathbf{R}^*, \mathbf{t}^* = \arg \min_{\mathbf{R}, \mathbf{t}} \sum_{(i,j) \in \mathcal{K}} \|\mathbf{u}_i - \pi(\mathbf{R}\mathbf{x}_j - \mathbf{t})\|, \quad (2)$$

where $\pi(\cdot)$ is the camera projection function. After solving the camera pose, the inlier correspondences $\mathcal{K}^* \doteq \{(i_k, j_k)\}_{k=1:N'}$ can be determined via their re-

Fig. 6: Experimental setup.

Fig. 7: Our created ATG-PVD dataset: (a)–(c) the images on the first row are used in the suspicion phase, while the images on the second row are used in the investigation phase; (d) and (e) the statistical analysis of the dataset.

GPU⁹, which has an 8 GB LPDDR4 memory and 256 CUDA cores, for IPC detection. Furthermore, we also equip our drone with an RPLIDAR A2¹⁰, which can perform 360° omnidirectional laser range scanning.

4.2 ATG-PVD Dataset

Using the aforementioned experimental setup, we created a large-scale real-world dataset, named the ATG-PVD dataset, for parking violation detection. Our dataset is publicly available at sites.google.com/view/atg-pvd for research purposes. The ATG-PVD dataset contains seven sequences (resolution: 2160 × 3840 pixels) and the corresponding 2D bounding box annotations for car detection and classification. The ground truth used in the suspicion phase has three classes: a) IPC candidates, b) MCs and c) LPCs, while in the investigation phase, the IPC ground truth is also provided. Examples of the images used in the suspicion and investigation phases are shown in Fig. 7(a)–(c).

In our experiments, we divide our ATG-PVD dataset into a training set and a testing set, which respectively contains 4924 and 4398 images. The statistics for these two sets are shown in Fig. 7(d) and (e), where it can be observed that

⁹ developer.nvidia.com/embedded/jetson-tx2

¹⁰ slamtec.com/en/Lidar/A2

Table 1: Ablation study of our SwiftFlow on the KITTI flow 2015 [44] training dataset. Best results are shown in bold font.

Backbone	Reduce Dense	Shared Weights	Deformable Convolution	F1-all (%)	# Params (M)
PWC-Net	–	–	–	8.37	8.75
	✓	–	–	7.22	5.26
	✓	✓	–	6.95	2.18
	✓	✓	✓	6.51	2.51

Table 2: The evaluation results on the KITTI flow benchmarks, where DDFlow [18], UnFlow [17], Flow2Stereo [45] and SelFlow [19] are the state-of-the-art self-supervised approaches. Best results are shown in bold font.

Approach	KITTI 2012		KITTI 2015		Runtime (s)
	Out-Noc (%)	Rank	F1-all (%)	Rank	
DDFlow [18]	4.57	60	14.29	91	0.06
UnFlow [17]	4.28	53	11.11	66	0.12
Flow2Stereo [45]	4.02	48	11.10	65	0.05
SelFlow [19]	3.32	34	8.42	51	0.09
SwiftFlow (Ours)	2.64	24	7.23	35	0.03

there are more IPC candidates or IPCs than MCs and LPCs. Additionally, most images contain fewer than five cars. Furthermore, our experiments are conducted on downsampled images with a resolution of 540×960 pixels. Sections 4.3, 4.4, and 4.5 respectively discuss the performances of SwiftFlow, Flow-RCNN and our PVD system in terms of both qualitative and quantitative experimental results.

4.3 Evaluation of SwiftFlow

Ablation Study We conduct an ablation study to validate the effectiveness of SwiftFlow. The experimental results are presented in Table 1. We can see that, by removing dense connections between different levels, our approach can reduce many parameters, but still retain a similar optical flow estimation performance, compared with the PWC-Net [13] baseline. Moreover, sharing weights of flow estimation modules can yield a performance improvement with fewer parameters. Furthermore, thanks to deformable convolution, our proposed SwiftFlow achieves the best performance with only a few additional parameters.

Evaluation Since our ATG-PVD dataset does not contain optical flow ground truth, we evaluate our proposed SwiftFlow on the KITTI flow 2012 [46] and 2015 [44] benchmarks. According to the online leaderboard of the KITTI flow benchmarks, as shown in Table 2, our SwiftFlow ranks 24th on the KITTI flow 2012

Fig. 8: Examples from the KITTI flow benchmarks, where rows (a) and (b) on columns (1)–(3) show the optical flow estimations and the corresponding error maps of (1) UnFlow [17], (2) SelfFlow [19] and (3) our SwiftFlow, respectively. Significantly improved regions are highlighted with green dashed boxes.

Fig. 9: Examples of the optical flow estimation results on our ATG-PVD dataset. Our proposed SwiftFlow is compared with DDFlow [18] and UnFlow [17].

benchmark¹¹ and 35th on the KITTI flow 2015 benchmark¹², outperforming all other state-of-the-art unsupervised optical flow estimation approaches, with a faster running speed (in real time) achieved in the mean time. Fig. 8 presents examples from the KITTI flow benchmarks, where we can see that SwiftFlow yields more robust results than others. Furthermore, Fig. 9 shows optical flow estimation results on our ATG-PVD dataset, indicating that our proposed SwiftFlow performs much more accurately than both DDFlow [18] and UnFlow [17], another two well-known unsupervised optical flow estimation approaches, especially on the boundary of the MCs.

4.4 Evaluation of Flow-RCNN

In our experiments, we compute the mean average precision (mAP) over ten IoU thresholds between 0.50 and 0.95 (refer to [47] for more details) to quantitatively evaluate the performance of our proposed Flow-RCNN. It should be noted that

¹¹ cvlibs.net/datasets/kitti/eval_stereo_flow.php?benchmark=flow

¹² cvlibs.net/datasets/kitti/eval_scene_flow.php?benchmark=flow

Table 3: Car detection mAP, where the best results are shown in bold font.

Method	Total	IPC candidate	MC	LPC
Faster-RCNN [7]	0.770	0.844	0.672	0.789
Flow-RCNN (ours)	0.789	0.845	0.733	0.796

(a) Occlusion

(b) Motion blur

(c) Complex environment

Fig. 10: Examples of our Flow-RCNN results.

IPCs are regarded as IPC candidates in both training and testing experiments, due to the fact that IPCs are re-identified as IPC candidates.

We compute mAP for all three categories (IPC candidate, MC and LPC) so as to comprehensively evaluate the performance of our proposed Flow-RCNN. The quantitative results are provided in Table 3, where it can be observed that Flow-RCNN outperforms the baseline network Faster-RCNN [7] (especially for MC detection) in terms of both car detection and classification. It is rather astonishing that Faster-RCNN can still successfully detect many MCs from only RGB images, even without using optical flow information. We speculate that the baseline network might also consider the road textures around a car when inferring its category. For instance, an MC is typically at the center of a lane, and the road textures around it are similar, which can weaken the influence caused by motion blur problem.

Experimental results of our Flow-RCNN are given in Fig. 10, showing the robustness of our proposed approach. For example, in Fig. 10(a), the light pole, that occludes part of an IPC candidate, can produce a similar optical flow estimation to an MC. Fortunately, our Flow-RCNN which fuses both RGB and flow information can still detect the IPC candidate correctly. Furthermore, although it is hard to extract features from a blurred car image, it can be seen in Fig. 10(b) that our proposed approach can avoid such misdetections by leveraging additional optical flow information. Moreover, in complex environments, such as the case shown in Fig. 10(c), car with different categories can still be successfully detected and classified.

4.5 Evaluation of Parking Violation Detection

We also comprehensively evaluate the performance of the entire system for parking violation detection using our ATG-PVD dataset, and a precision of 91.7%, a recall of 94.9% and an F1-Score of 93.3% are achieved. An example of the detected IPCs in the map is illustrated in Fig. 11, where readers can observe

Fig. 11: An example of the detected IPCs in the map.

that our proposed suspect-and-investigate system can detect parking violations effectively and efficiently.

5 Conclusion

In this paper, we proposed a novel, robust and cost-effective parking violation detection system embedded in an ATG-R680 drone equipped with a TX2 GPU. Our system utilizes a so-called suspect-and-investigate framework, which consists of: 1) an unsupervised optical flow estimation network named SwiftFlow, 2) a novel flow-guided object detection network named Flow-RCNN, and 3) a drone re-localization and IPC re-identification module based on VSLAM. On the KITTI flow 2012 and 2015 benchmarks, our proposed SwiftFlow outperforms all other state-of-the-art unsupervised optical flow estimation approaches in terms of both speed (real-time performance was achieved) and accuracy. By incorporating the inferred optical flow information into our object detection framework, IPC candidates, MCs and LPCs can be effectively detected and classified, even in many challenging cases. In the investigation phase, our VSLAM module detects loop closure to re-localize the drone in the pre-built map. After the drone is successfully re-localized, we further re-identify whether an existing IPC candidate is an actual IPC. The experimental results both qualitatively and quantitatively demonstrate the effectiveness and robustness of our proposed parking violation detection system.

Acknowledgements

This work was supported by the National Natural Science Foundation of China, under grant No. U1713211, Collaborative Research Fund by the Research Grants Council Hong Kong, under Project No. C4063-18G, and the Research Grant Council of Hong Kong SAR Government, China, under Project No. 11210017, awarded to Prof. Ming Liu.

References

1. McKee, M., Stuckler, D.: If the world fails to protect the economy, covid-19 will damage health not just now but also in the future. *Nature Medicine* **26**(5) (2020) 640–642
2. Hughes, T.: Poor, essential and on the bus: Coronavirus is putting public transportation riders at risk (April 2020)
3. Garcetti, E.: Mayor garcetti relaxes parking enforcement (March 2020)
4. Zhou, Y., Liu, L., Shao, L., Mellor, M.: Dave: A unified framework for fast vehicle detection and annotation. In: *European Conference on Computer Vision*, Springer (2016) 278–293
5. Regazzoni, C.S., Cavallaro, A., Wu, Y., Konrad, J., Hampapur, A.: Video analytics for surveillance: Theory and practice [from the guest editors]. *IEEE Signal Processing Magazine* **27**(5) (2010) 16–17
6. Council, B.C.: Codes of practice for operation of CCTV Enforcement Cameras. Birmingham City Council (2013)
7. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **39**(6) (2017) 1137–1149
8. Horn, B.K., Schunck, B.G.: Determining optical flow. In: *Techniques and Applications of Image Understanding*. Volume 281., International Society for Optics and Photonics (1981) 319–331
9. Mémin, E., Pérez, P.: Dense estimation and object-based segmentation of the optical flow with robust techniques. *IEEE Transactions on Image Processing* **7**(5) (1998) 703–719
10. Brox, T., Bruhn, A., Papenberger, N., Weickert, J.: High accuracy optical flow estimation based on a theory for warping. In: *European conference on computer vision*, Springer (2004) 25–36
11. Zach, C., Pock, T., Bischof, H.: A duality based approach for realtime tv-l 1 optical flow. In: *Joint pattern recognition symposium*, Springer (2007) 214–223
12. Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., Brox, T.: FlowNet: Learning optical flow with convolutional networks. In: *Proceedings of the IEEE international conference on computer vision*. (2015) 2758–2766
13. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2018) 8934–8943
14. Hui, T.W., Tang, X., Change Loy, C.: LiteflowNet: A lightweight convolutional neural network for optical flow estimation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2018) 8981–8989
15. Hur, J., Roth, S.: Iterative residual refinement for joint optical flow and occlusion estimation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019) 5754–5763
16. Ren, Z., Yan, J., Ni, B., Liu, B., Yang, X., Zha, H.: Unsupervised deep learning for optical flow estimation. In: *Thirty-First AAAI Conference on Artificial Intelligence*. (2017)
17. Meister, S., Hur, J., Roth, S.: Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In: *Thirty-Second AAAI Conference on Artificial Intelligence*. (2018)

18. Liu, P., King, I., Lyu, M.R., Xu, J.: Ddflow: Learning optical flow with unlabeled data distillation. In: Proceedings of the AAAI Conference on Artificial Intelligence. Volume 33. (2019) 8770–8777
19. Liu, P., Lyu, M., King, I., Xu, J.: Selfflow: Self-supervised learning of optical flow. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2019) 4571–4580
20. Lin, T., Goyal, P., Girshick, R., He, K., Dollr, P.: Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **42**(2) (2020) 318–327
21. Girshick, R.: Fast r-cnn. In: International Conference on Computer Vision (ICCV). (2015)
22. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints. In: The European Conference on Computer Vision (ECCV). (September 2018)
23. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. In: arXiv preprint arXiv:1904.07850. (2019)
24. Klein, G., Murray, D.: Parallel tracking and mapping for small ar engel2018dso. In: Proceedings of the 2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality, IEEE Computer Society (2007) 1–10
25. Strasdat, H., Davison, A.J., Montiel, J.M., Konolige, K.: Double window optimisation for constant time visual slam. In: 2011 International Conference on Computer Vision, IEEE (2011) 2352–2359
26. Mur-Artal, R., Montiel, J.M.M., Tardos, J.D.: ORB-SLAM: a versatile and accurate monocular slam system. *IEEE Transactions on Robotics* **31**(5) (2015) 1147–1163
27. Newcombe, R.A., Lovegrove, S.J., Davison, A.J.: Dtam: Dense tracking and mapping in real-time. In: 2011 international conference on computer vision, IEEE (2011) 2320–2327
28. Engel, J., Schöps, T., Cremers, D.: Lsd-slam: Large-scale direct monocular slam. In: European conference on computer vision, Springer (2014) 834–849
29. Engel, J., Koltun, V., Cremers, D.: Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence* **40**(3) (2018) 611–625
30. Forster, C., Pizzoli, M., Scaramuzza, D.: Svo: Fast semi-direct monocular visual odometry. In: 2014 IEEE international conference on robotics and automation (ICRA), IEEE (2014) 15–22
31. Yang, S., Scherer, S.: Cubeslam: Monocular 3-d object slam. *IEEE Transactions on Robotics* **35**(4) (2019) 925–938
32. Huang, J., Yang, S., Mu, T.J., Hu, S.M.: Clustervo: Clustering moving instances and estimating visual odometry for self and surroundings. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2020) 2168–2177
33. Nicholson, L., Milford, M., Sünderhauf, N.: Quadricslam: Dual quadrics from object detections as landmarks in object-oriented slam. *IEEE Robotics and Automation Letters* **4**(1) (2018) 1–8
34. Runz, M., Buffier, M., Agapito, L.: Maskfusion: Real-time recognition, tracking and reconstruction of multiple moving objects. In: 2018 IEEE International Symposium on Mixed and Augmented Reality (ISMAR), IEEE (2018) 10–20
35. McCormac, J., Clark, R., Bloesch, M., Davison, A., Leutenegger, S.: Fusion++: Volumetric object-level slam. In: 2018 international conference on 3D vision (3DV), IEEE (2018) 32–41
36. Zhang, T., Zhang, H., Li, Y., Nakamura, Y., Zhang, L.: Flowfusion: Dynamic dense rgb-d slam based on optical flow. arXiv preprint arXiv:2003.05102 (2020)

37. Tateno, K., Tombari, F., Laina, I., Navab, N.: Cnn-slam: Real-time dense monocular slam with learned depth prediction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2017) 6243–6252
38. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: Proceedings of the IEEE international conference on computer vision. (2017) 764–773
39. Wang, Y., Yang, Y., Yang, Z., Zhao, L., Wang, P., Xu, W.: Occlusion aware unsupervised learning of optical flow. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2018) 4884–4893
40. Tomasi, C., Manduchi, R.: Bilateral filtering for gray and color images. In: Sixth international conference on computer vision (IEEE Cat. No. 98CH36271), IEEE (1998) 839–846
41. Jonschkowski, R., Stone, A., Barron, J.T., Gordon, A., Konolige, K., Angelova, A.: What matters in unsupervised optical flow. arXiv preprint arXiv:2006.04902 (2020)
42. Mur-Artal, R., Tardós, J.D.: ORB-SLAM2: An open-source slam system for monocular, stereo, and rgb-d cameras. IEEE Transactions on Robotics **33**(5) (2017) 1255–1262
43. Rublee, E., Rabaud, V., Konolige, K., Bradski, G.: Orb: An efficient alternative to sift or surf. In: 2011 International conference on computer vision, Ieee (2011) 2564–2571
44. Menze, M., Heipke, C., Geiger, A.: Joint 3d estimation of vehicles and scene flow. In: ISPRS Workshop on Image Sequence Analysis (ISA). (2015)
45. Liu, P., King, I., Lyu, M.R., Xu, J.: Flow2stereo: Effective self-supervised learning of optical flow and stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2020) 6648–6657
46. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: Conference on Computer Vision and Pattern Recognition (CVPR). (2012)
47. Lin, T., Maire, M., Belongie, S.J., Bourdev, L.D., Girshick, R.B., Hays, J., Perona, P., Ramanan, D., Doll’ar, P., Zitnick, C.L.: Microsoft COCO: common objects in context. CoRR **abs/1405.0312** (2014)