Multimed Tools Appl (2015) 74:3121–3135
DOI 10.1007/s11042-013-1771-7

Synthetic sequences and ground-truth flow field generation for algorithm validation

Naveen Onkarappa · Angel D. Sappa

Published online: 19 November 2013

Abstract Research in computer vision is advancing by the availability of good datasets that help to improve algorithms, validate results and obtain comparative analysis. The datasets can be real or synthetic. For some of the computer vision problems such as optical flow it is not possible to obtain ground-truth optical flow with high accuracy in natural outdoor real scenarios directly by any sensor, although it is possible to obtain ground-truth data of real scenarios in a laboratory setup with limited motion. In this difficult situation computer graphics offers a viable option for creating realistic virtual scenarios. In the current work we present a framework to design virtual scenes and generate sequences as well as ground-truth flow fields. Particularly, we generate a dataset containing sequences of driving scenarios. The sequences in the dataset vary in different speeds of the on-board vision system, different road textures, complex motion of vehicle and independent moving vehicles in the scene. This dataset enables analyzing and adaptation of existing optical flow methods, and leads to invention of new approaches particularly for driver assistance systems.

Keywords Ground-truth optical flow · Synthetic sequence · Algorithm validation

1 Introduction

Computer Vision has got applications in different fields of life. The research in computer vision is always motivated, as well as supported by the benchmarking dataset with ground-truth information. The availability of ground-truth information makes the dataset very useful for the evaluation of different methods. Generally, ground-truth can be obtained by manual labelling and/or sophisticated equipments such as for detection, recognition and segmentation tasks. Optical flow technique is an important approach in motion estimation that is

N. Onkarappa (✉) · A. D. Sappa
Computer Vision Center, Campus UAB, 08193, Bellaterra, Barcelona, Spain
e-mail: naveen@cvc.uab.es

A. D. Sappa
e-mail: asappa@cvc.uab.es

useful in many fields such as action recognition, surveillance, image compression, robot navigation, driver assistance systems to mention a few. In spite of the fact that the research on optical flow started more than three decades ago, the seminal methods were just proposed in 1981 [10, 12]. There are many advances in this field since then and it has got momentum since several years from now. This is mainly due to the availability of datasets with ground-truth flow information as well as due to the large increase in computational power. The availability of dataset challenges the existing state of the art methods and promotes research to propose newer methods. Also it allows the evaluation of existing methods. Another benefit is that the ground-truth data satisfies the need of the learning based approaches. Particular to the optical flow, there are issues in obtaining ground-truth optical flow of real scenarios. There is no such direct sensor to obtain ground-truth information with a good accuracy in real scenarios unless performed in a restricted laboratory environment. It can be possible in a laboratory environment for limited motion using hidden texture [6], but for natural outdoor scenes, there is no way at present to have ground-truth optical flow with good accuracy. In such a scenario, the alternate viable solution is to have synthetic datasets. Hence, the developments in the area of computer graphics have given the advantage of creating such synthetic datasets facilitating the validation of computer vision algorithms.

The safety is one of the top priority while driving. Although there are improvements in vehicular technology and infrastructures that increases the safety of human and vehicles, the main cause of accidents are the human errors. Advanced driver assistance systems (ADAS) is an upcoming area where many computer vision techniques have potential to tackle challenging situations. Motion is an important input for many of the ADAS. The optical flow as a motion estimation technique has an important role in ADAS. There are several variations in a driving scenario affected by environments such as urban, highway, countryside where the vehicle is being driven, types of road (well structured, differently textured), weather conditions, and daylight conditions. There exist no sensor to directly acquire ground-truth flow field along with the image sequences and it is not possible to obtain accurate ground-truth by other ways such as by using depth sensors. The alternative way for this situation is to create virtual scenarios using 3D designing tools. The advantage to go with synthetic sequences is that one can create all the possible different environments and scenarios as mentioned before in the case of driving. Although there exist a question that how realistic are these synthetic ones compared to real ones, one can thrive to integrate realism into virtual scenes with the latest advances in computer graphics. Actually, there are some work in this direction, for instance recently in [15] the authors did an attempt to create realistic synthetic scenarios. They show that it is possible to create more realism by varied lighting conditions, varied material properties and by exploiting state of the art in computer graphics.

As discussed above, driving scenarios involve varied complexities due to motion of on-board camera, dynamic scene with independently moving vehicles and other additional factors in the environment. There are several of such characteristics that need to be analyzed to develop a good optical flow method in driving scenarios. For example, change in road texture influences the optical flow accuracy. To do such a study one needs to have image sequences of the same structural scene but with different road textures. In reality, it is very difficult to create such one and impossible to generate ground-truth optical flow of good accuracy. The other best solution is to have synthetic sequences generated and then obtaining ground-truth will also be easier. The existing datasets do not provide any of such sequences. In this work we are proposing a framework to generate such sequences along with the ground-truth flow fields. The most important contribution is the generated dataset of sequences with different speeds, different textures with added complexity. This is the first work on this kind of dataset in the ADAS domain. Moreover we are proposing a framework that can be used to generate any kind of sequences of complex scenarios and being able to generate ground-truth optical flow. The research community can improve and develop new datasets according to their own requirements.

The paper is organized as follows. Next section describes existing optical flow datasets, their applicability and drawbacks. The proposed framework is presented in Section 3 followed by the technique of generating ground-truth optical flow in Section 4 and dataset generated in Section 5. The generated ground-truth optical flow is validated by back-projecting images. Also our dataset is compared with state of the art synthetic driving sequences. Further a simple analysis of effectiveness of different optical flow algorithms on our dataset is also performed. This analysis and validation is provided in Section 6. Then, the paper is concluded with Section 7.

2 Related work

There are several datasets for optical flow (e.g., [6–9, 11, 18]) available to the research community for algorithm validation, evaluation and comparison. One of the most well known dataset is Middlebury [6], which contains both real and synthetic sequences. The ground-truth for real sequences are generated with hidden texture in a controlled laboratory environment. Most of the contributions in optical flow are evaluated on this dataset [1], which contains limited scenarios and image pairs have small displacements. One of the first performance evaluation work has been done in [7] while presenting a few synthetic sequences with an evaluation methodology. The sequences are very simplistic. Later McCane et al. [14] introduce several complex synthetic sequences and also compared several optical flow methods. A real sequence is provided by Liu et al. [11]. In this work the authors also present dense ground-truth data. They annotate an image into different layered segments and compute optical flow using existing methods for each layers seperately. Obviously this process adds several errors into the ground-truth. This dataset do not involve much realistic characteristics that one can expect in driving scenarios. Few real sequences were proposed by Otte and Nagel [16]. These sequences are also simplistic with geometry, texture and small displacements.

A big challenge when a real dataset with realistic scenarios need to be obtained lies on the difficulty in obtaining ground-truth optical flow. Recently, Geiger et al. [9] have proposed a new real dataset of driving scenarios containing large displacements, specularity, shadows and different illuminations. They have also provided sparse ground-truth flow field with a density of around 50 %. This dataset is referred to as KITTI. Although this dataset is obtained from real data, the main limitation lies in the sparseness of the data. It does not provide dense ground-truth and the ground-truth is not accurate due to errors in the registration of laser scanner point clouds to the image plane. One can observe that the camera has a wide angle, but the images are not well focussed in all the regions. The work in [18] introduces few synthetic sequences of driving scenarios with ground-truth optical flow. This dataset is referred to as EISATS. The ground-truth flow fields in these sequences do not show the occlusion areas when there are moving vehicles. Also a set of simple sequences are provided by [13]. In this work the authors attempt to find the best suitable optical flow algorithm based on the flow confidence at every pixel.

Considering the drawbacks of the existing datasets, recently Butler et al. [8] presented a large synthetic dataset from the animated short film Sintel. This dataset is referred to as MPI-Sintel. They have incorporated several complexities such as motion blur, defocus blur, shading and atmospheric effects. The dataset contains the same image sequences with three

Fig. 1 Snapshot of a 3D design in Maya

levels of complexities. This dataset is expected to be a new benchmark for optical flow research [2]. The datasets [6, 8, 9] have separate training and evaluation sets. The evaluation set does not provide ground-truth flow data to the public for the purpose of evaluation. One can think that the state of the art methods that give the best results on Middlebury dataset can also perform similarly on KITTI dataset and MPI-Sintel. However, by analyzing the KITTI flow evaluation [3] and Sintel evaluation [2] one can realize that such a statement is wrong due to the difficulties of the particular datasets. This proves that a diverse collection of datasets will take forward research to new levels. Among the existing datasets, the sequences from [9, 18] are intended for ADAS applications.

3 Synthetic sequence generation framework

The objective of the current work is to generate synthetic sequences along with ground-truth flow fields. We present a framework to generate sequences in a driving scenario considering three particular cases: i) on-board vision system in a vehicle with different speeds; ii) roads with different textures; iii) scenarios with independently moving vehicles. For analyzing the influence of speed on optical flow accuracy, we need to have image sequences of the same scene, but the on-board vision system vehicle moving with different speeds. Similarly, for analyzing the impact of texture, we need image sequences of the same scene (i.e., surrounding scene structure) but with just different textures. In reality, it is impossible to have such scenarios by default and also to generate ground-truth optical flow in real life. Hence, in the current work we propose a framework similar to the one presented in [13]. We use Maya[1] to develop a 3D scene. We have built a synthetic 3D urban scenario that consists of a straight road and buildings around it with appropriate textures. A camera assumed to be fixed in a vehicle referred to as on-board camera moves along the road in the model. The images are rendered using in-built Maya software with production quality. All the images are rendered with a resolution of 640×480 pixels. Figure 1 shows a snapshot of the 3D urban scenario designed in Maya.

[1] www.autodesk.com/maya

For case i) the on-board camera moves at different speeds along the same path straight along the road and the images are rendered. Some of the rendered images are shown in Fig. 2. The *top-left* is the first image which is common to all the sequences. Second *row-left* is the second image in a sequence and *bottom-left* is the second image of another sequence of higher speed. In ADAS scenarios, the road surface covers a major part in the images taken through vehicle's camera. The flow vectors computed from this surface are more reliable as there could be more inaccuracies in other areas of the image due to occlusions, specularities etc. For case ii) to analyze texture influence particularly, only road texture can be changed very easily without disturbing the 3D design. Hence, for a given speed, several sequences

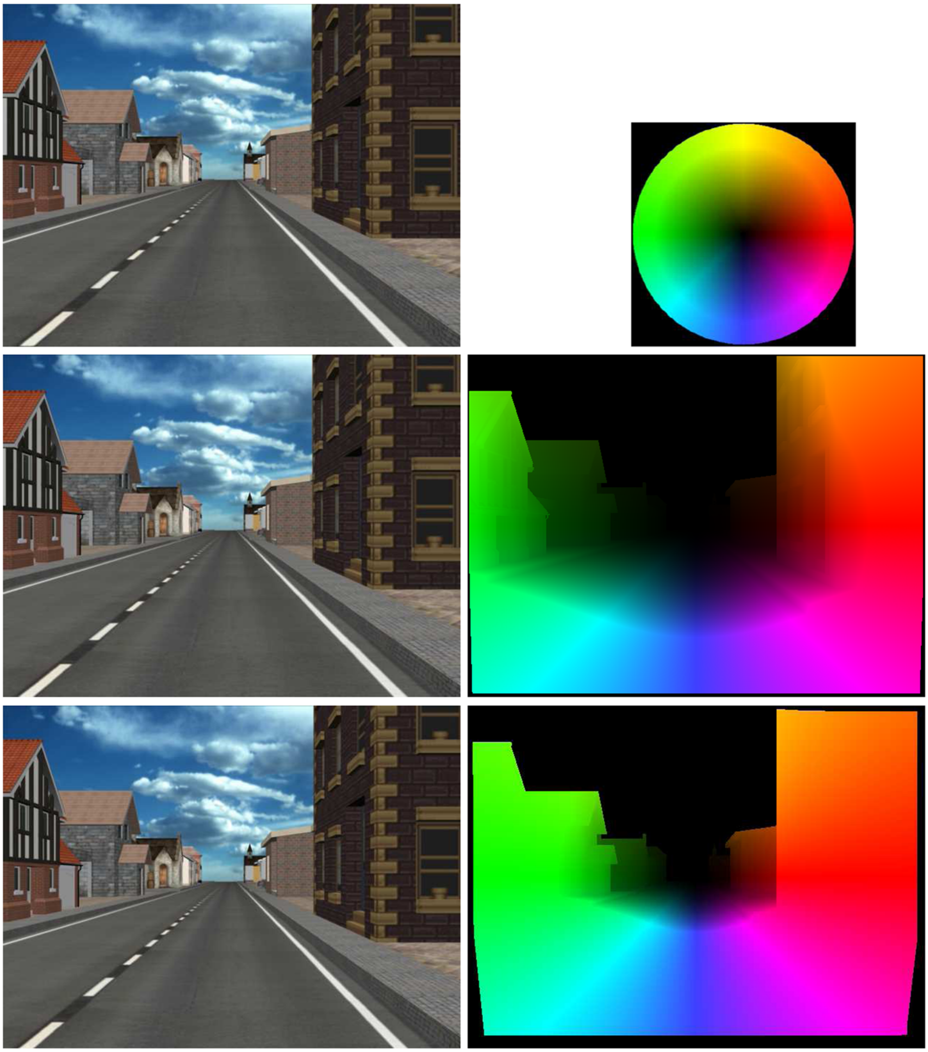

Fig. 2 Images from sequences of different speeds: (*top-left*) first frame common for all sequences; (*top-right*) colormap used to show the flow fields; (*left-column*) second frames from the sequences of different speeds in the increasing order (*2nd* and *3rd rows*); (*right-column*) the ground-truth flow fields between the respective first and second frames

with different road textures are rendered. Figure 3 shows images with three different textures on the road surface.

In 3D design, the designer has full control over all the things such as motion of camera, lighting, textures and motion of different objects in the scene. With this capability, for case iii) we have added two moving vehicles in the scene. One vehicle coming towards the on-board camera vehicle and another coming from a cross road. To add complexity, pitch and yaw variations to the on-board camera are also incorporated. The yaw is 0.25 degrees to the left/right and the pitch is 0.25 degrees to the up/down. Rendered images from this sequence are shown in Fig. 4. In all these above cases, the camera focal length is 35 mm and the framerate is 24 fps.

4 Optical flow ground-truth

This section describes the generation of ground-truth flow fields. It is based on the well known ray-tracing approach. Ray-tracing is basically a technique of tracing the path of a light ray. It is being used in 3D computer graphics discipline to render photorealistic images. The same idea has been used in the current work to estimate displacement vector of each pixel. The complete information of designed 3D model enables us to use such a technique to compute the distplacement vector. A pixel P in an image plane at time t is traced to its position in the 3D scene. Then this 3D point is projected back to the image plane at time

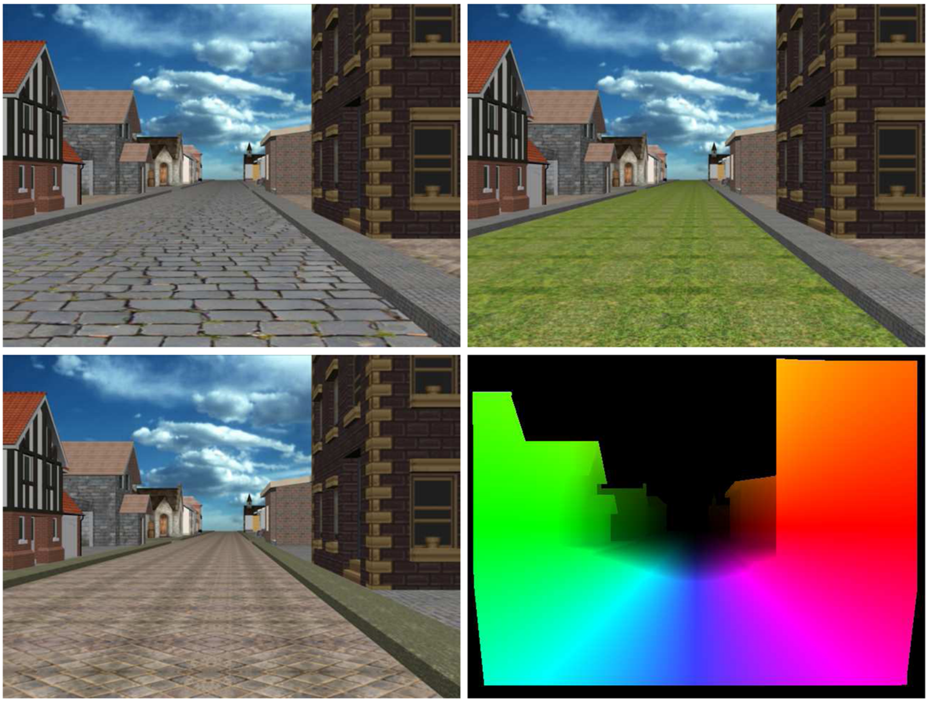

Fig. 3 (*top-left*), (*top-right*) and (*bottom-left*) frames with different texture from different sequences; (*bottom-right*) Ground-truth flow field for all the pairs of images in (*top-left*), (*top-right*) and (*bottom-left*); all of them have the same scene geometry and same speed but with different textures

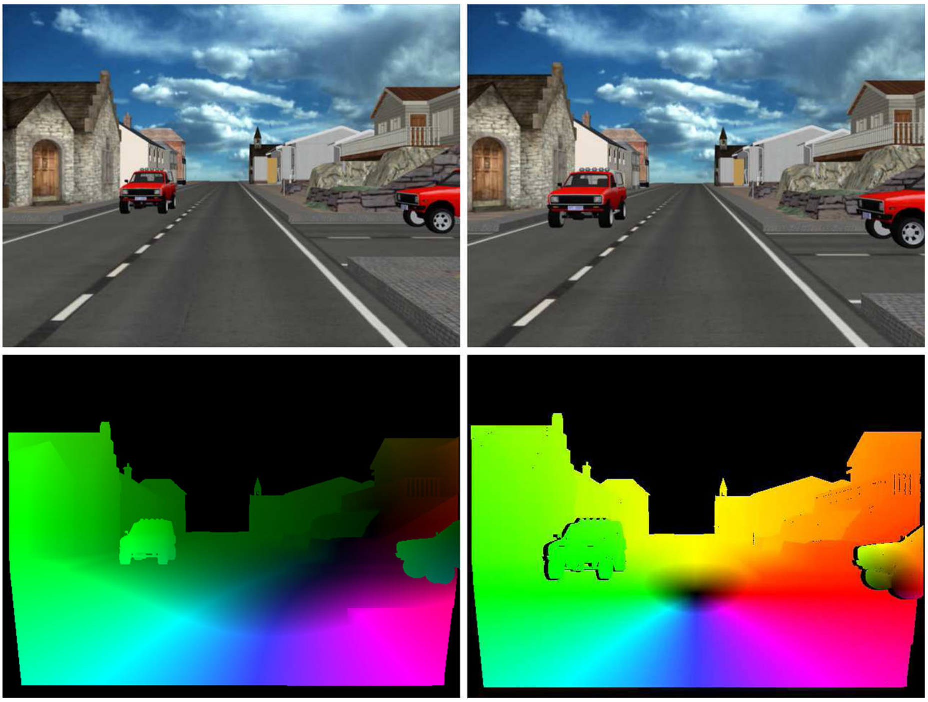

Fig. 4 (*top*) Two different image frames from a sequence with independently moving vehicles and different egomotion. (*bottom*) Ground-truth flow fields between the above frames and to their next ones in the sequence

$t + 1$. Since we know the simulated camera motion in time, difference in pixel position at different times on the image plane gives the displacement vector. Figure 5 depicts the ray-tracing approach. A vector from P_t to P_{t+1} is the flow vector at P_t with respect to image at time t.

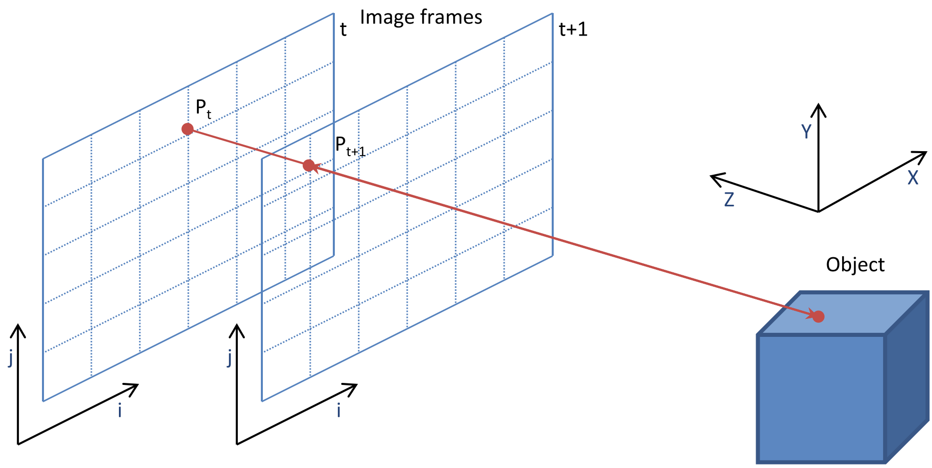

Fig. 5 Ray-tracing to estimate displacement vector

Table 1 Different synthetic sequences

Sequence		Number of frames (For each of the four textures T1, T2, T3 and T4)
Speed S1		40
Speed S2		20
Speed S3		13
Speed S4		10
Complex sequences	(For each of the four speeds)	

The computed ground-truth flow fields for different speeds are shown in Fig. 2. Figure 2 (*top-right*) is the colormap used to depict flow fields. The color indicates the direction and intensity indicates the magnitude. Figure 2 (*middle-right*) is the ground-truth flow field between the frames in (*top-left*) and (*middle-left*). Similarly, Fig. 2 (*bottom-right*) is the ground-truth flow field between top-left and bottom-left. One can notice large blank space at the bottom of Fig. 2 (*bottom-right*) that indicates occluded area that is larger at a higher speed than the one at lower speed in (*middle-right*). The maximum displacement in lower speed sequence is 8.31 pixels and that in higher speed sequence is 33.67 pixels. The ground-truth flow fields for differently textured image pairs are shown in Fig. 3 (*bottom-right*). Since the scene geometry is the same, the ground-truth flow fields for all

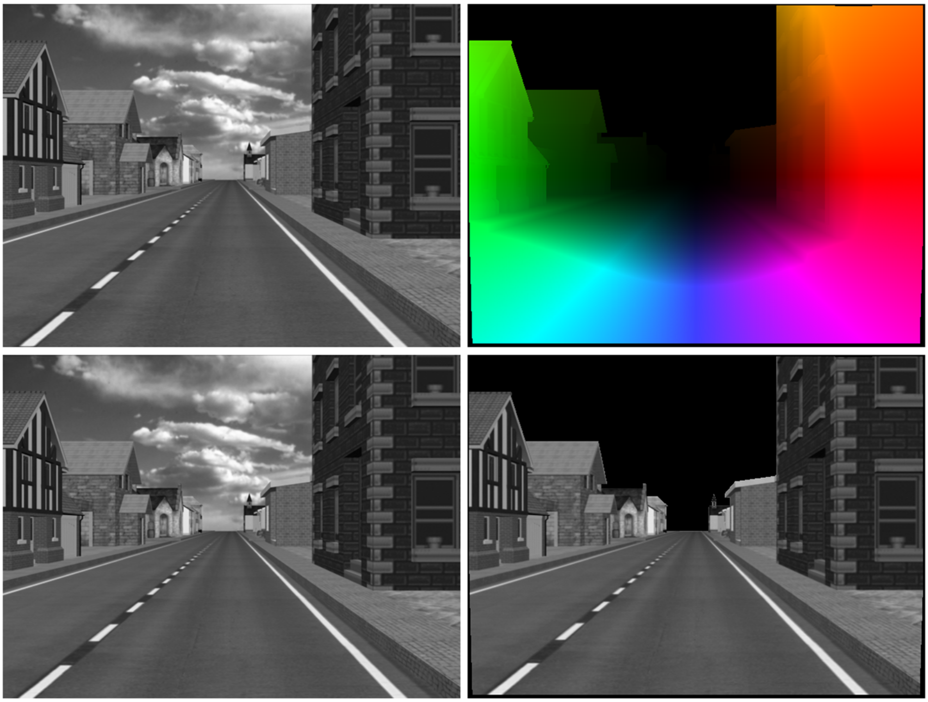

Fig. 6 Backward projection of a pair from S1T1 sequence in our sequnces; (*left*) An image pair. (*top-right*) Ground-truth flow field. (*bottom-right*) Synthesized image obtained by back-projecting frame two (*bottom-left*) using the flow field (*top-right*)

differently textured sequences are the same. Further ground-truth flow fields for two image pairs from complex sequences with moving vehicles are shown in Fig. 4. Notice that the flow field in Fig. 4 (*bottom-left*) has flow vectors at all pixels except the sky and occluded pixels at the boundary, whereas the flow field in (*bottom-right*) does not have values at the edge of moving vehicles depicting the occluded regions. Hence in synthetic sequence generation one has full control of all the possible scenarios and it is very useful.

5 Dataset

Using the framework presented in the current work, we have generated four sequences of different speeds. The sequence with higher speed has displacement four times the displacement of the sequence with lowest speed. Then, we have also created four sequences of different road textures for each speed. Hence we have generated sixteen sequences of different combinations of speeds and textures. In the case of analysis of optical flow accuracy for different speeds, if we generate equal number of frames for each speed, then the scene geometry covered by the distance varies and it might affect estimated flow accuracy. Thus we have generated frames for a constant distance and hence sequences of different speeds have different number of frames. The generated sequences with different textures have different textural properties, particularly in the increasing order of texture contrast. The third set contains complex sequences with two independently moving vehicles; pitch and yaw motion of

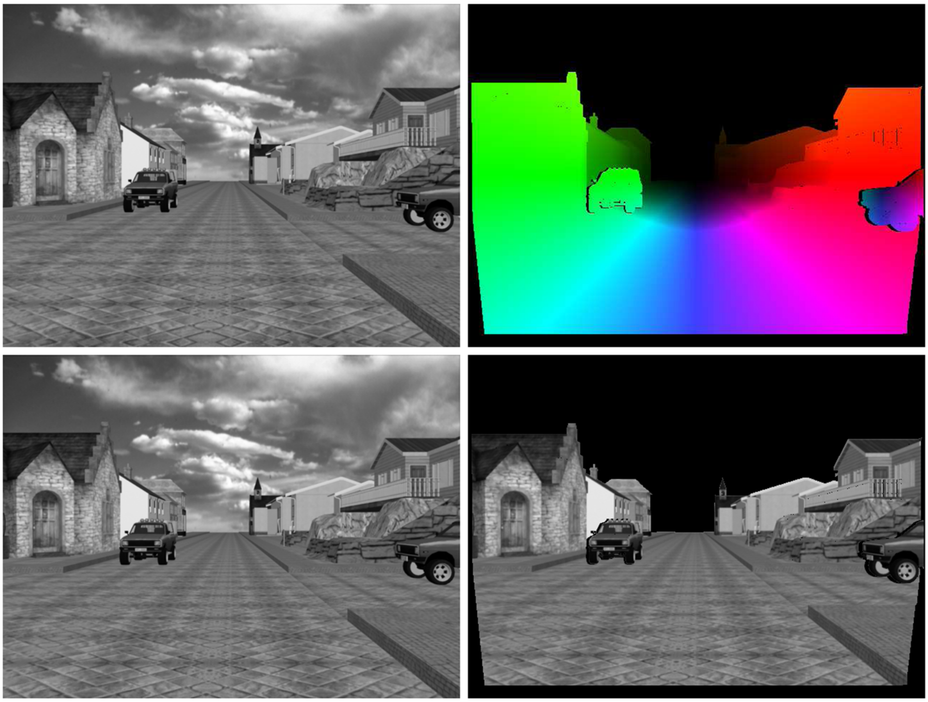

Fig. 7 Backward projection of a pair from complex-S4T4 sequence in our sequences; (*left*) An image pair. (*top-right*) Ground-truth flow field. (*bottom-right*) Synthesized image obtained by back-projecting frame two (*bottom-left*) using the flow field (*top-right*)

Table 2 NRMSE for several pairs from EISATS and our sequences

Image pairs from our sequences	NRMSE	Image pairs from EISATS	NRMSE
S1T1	0.0079	Pair1	0.0174
S4T4	0.0089	Pair2	0.0346
Complex S1T1	0.0100	Pair3	0.0139
Complex S4T4	0.0115	Pair4	0.0193

on-board camera gives an opportunity to study more dynamic scenes. These sequences are also generated for different combinations of road textures and on-board camera speeds without the constraint of constant distance covered. In this case, all the complex sequences have ten frames. The number of frames in all the sequences generated in the proposed dataset are depicted in Table 1. All the sequences generated with the proposed framework, together with the corresponding ground-truth data, are available through our website [4].

6 Analysis and validation

As described in Sections 4 and 5, we have generated several sequences with different characteristics of a driving scenario. In this section, we evaluate how good is the developed

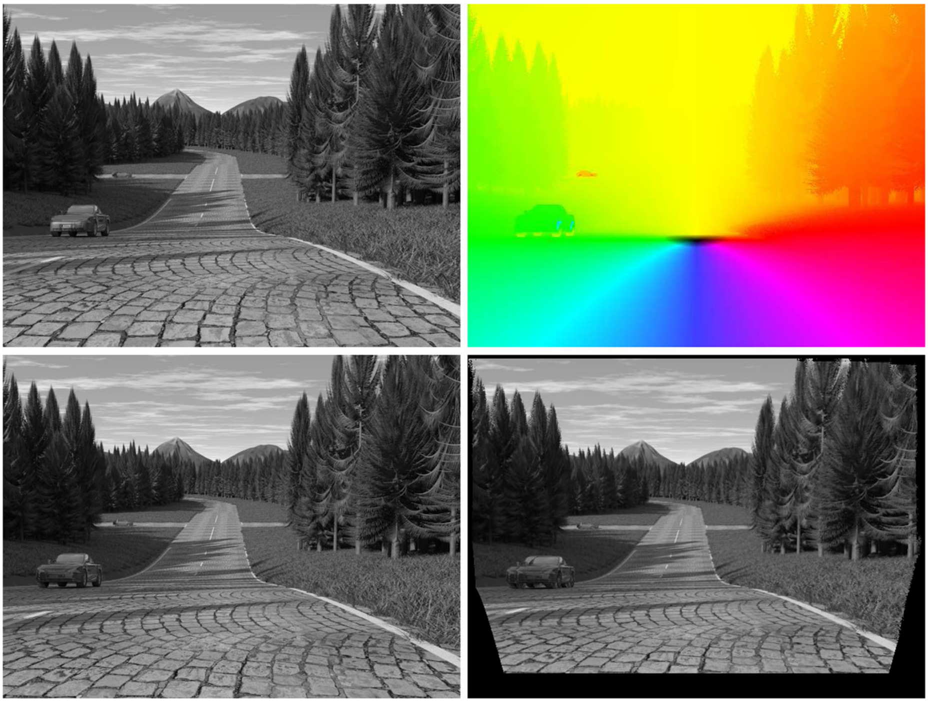

Fig. 8 Backward projection of Pair2 in EISATS sequence; (*left*) An image pair. (*top-right*) Ground-truth flow field provided by EISATS. (*bottom-right*) Synthesized image obtained by back-projecting frame two (*bottom-left*) using the flow field (*top-right*)

dataset. For this, we examine the generated ground-truth optical flow. In order to evaluate the accuracy, we consider several pairs of images from the generated sequences. For a given pair, we back-project the second image using the ground-truth optical flow to match the first frame. Figures 6 and 7 show two of such examples. Further, we calculate normalized-root-mean-square-error (NRMSE) between the first image and the corresponding back-projected one, considering only the region where the back-projected pixels exist. NRMSE for several pairs from our sequences are shown in second column in Table 2. It is expected that NRMSE should be lower near zero.

As presented in Section 2, the EISATS dataset is the most appropriated one to perform comparisons with the proposed dataset. Hence, we took several pairs of images from sequence-2 of set-2 in EISATS. The similar back-projection procedure has been performed and NRMSE are computed. Figures 8 and 9 show two examples of back-projection from the EISATS sequence. Quantitative results are presented in Table 2 in last column. Similar to the results for our dataset the NRMSE is also not zero, as could be expected; actually in the EISATS dataset errors are higher than in our sequences. One reason for higher errors in EISATS could be the complex scene containing trees compared to the building blocks in our case. By comparison, this analysis confirms that the generated ground-truth is of reasonable accuracy.

The important characteristics of our dataset is that occluded areas (black region at the boundary of moving vehicles in Fig. 7 (*top-right*) flow field) are identified while generating ground-truth flow vectors in those regions, whereas EISATS dataset does not give any

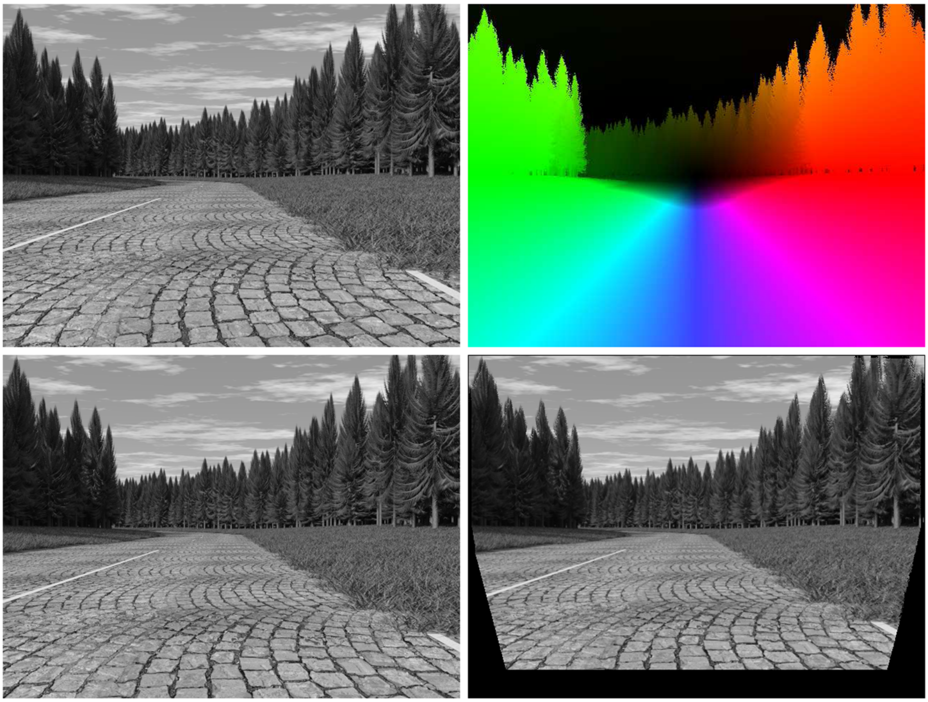

Fig. 9 Backward projection of Pair4 in EISATS sequence; (*left*) An image pair. (*top-right*) Ground-truth flow field provided by EISATS. (*bottom-right*) Synthesized image obtained by back-projecting frame two (*bottom-left*) using the flow field (*top-right*)

Table 3 Average Angular Error (AAE) and Average End Point Error (EPE), computed as in [7] and [14], for two pairs by different optical flow algorithms [14]

An image pair from sequence	TV-L1[19]		Classic+NL-Fast[17]		PolarOF[5]		MDP-Flow2[20]	
	AAE	EPE	AAE	EPE	AAE	EPE	AAE	EPE
S1T1	3.3133	0.2442	1.9884	0.0841	2.4817	0.1014	1.9697	0.0888
S4T4	1.4535	0.2746	0.9879	0.1261	1.1579	0.1402	0.7870	0.1292
Complex S1T1	3.8310	0.1625	2.6277	0.1185	3.6925	0.1544	2.5182	0.1206
ComplexS4T4	2.0844	0.5386	1.5853	0.1352	1.7304	0.1374	1.3722	0.1403

information about occluded areas. The drawback of this can be seen in Fig. 8 (*bottom-right*) back-projected image. The moving vehicle boundary looks double at that position. These kind of issues are taken care in the proposed dataset where occluded regions are accurately identified.

On further curiosity that how would the state of the art optical flow methods perform on our dataset, we have computed optical flow on the same selected pairs of images from our sequences. We have considered four optical flow algorithms. Both average-angular-error (AAE) and average-end-point-error (EPE) are calculated and are shown in Table 3. These results reveal that the method MDP-Flow2 [20], which is top-rank in Middlebury evaluation,

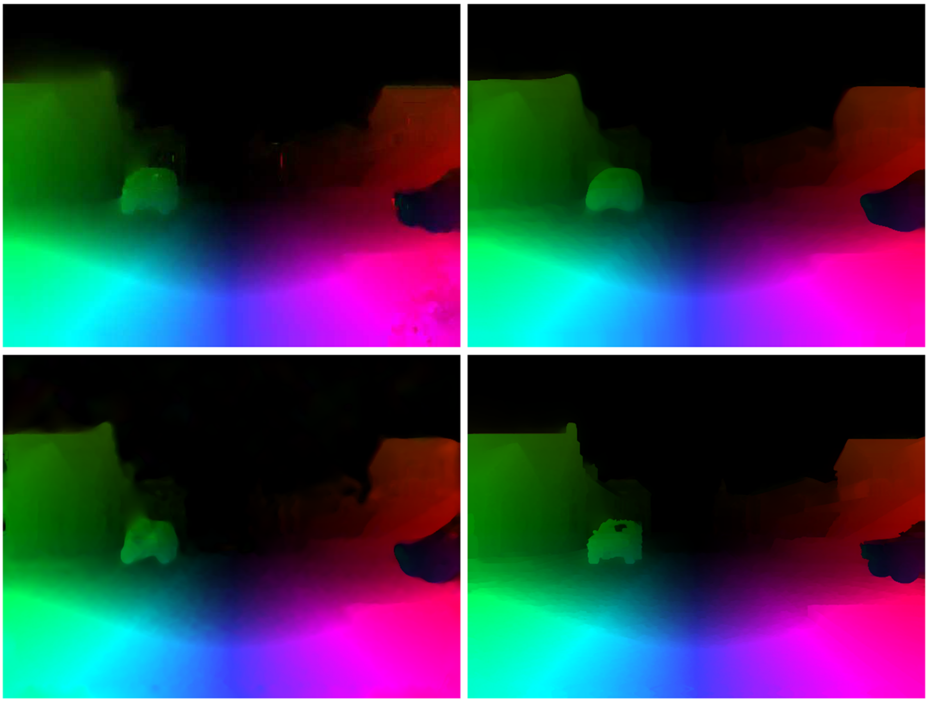

Fig. 10 Estimated optical flow fields by different methods on a an pair from Complex S4T4 sequence; (*top-left*) TV-L1 [19]. (*top-right*) Classic+NL-Fast [17]. (*bottom-left*) Polar optical flow [5]. (*bottom-right*) MDP-Flow2 [20]

has little better performance to Classic+NL-Fast approach [17] considering AAE. Whereas, considering EPE, Classic+NL-Fast has little better performance to MDP-Flow2. One can also observe that there are significant changes in errors across all methods when there are changes in complexity, texture and speed of the vehicle. This opens new questions about the robustness of existing optical flow approaches with respect to the considered factors here. There are plenty of other factors in driving scenarios which needs to be analyzed. Computed flow fields from all of the four methods for the pair from Complex-S4T4 sequence are shown in Fig. 10.

7 Conclusions

A framework to generate synthetic sequences using Maya is presented. The computation of ground-truth flow fields corresponding to the generated sequence is also detailed in the current work. This framework is used to generate sequences of driving scenarios. The scenarios include different speeds, different road textures, independently moving objects and complex motion of the on-board camera. The generated ground-truth data are validated by computing NRMSE and comparing them with the state of the art synthetic datasets of driving scenarios. We anticipate that the proposed framework and dataset will create interest in the driving assistance systems community to explore and improve current optical flow approaches. The obvious future goal tends towards incorporating more realism by motion blur, material characteristics and atmospheric effects.

Acknowledgments This work was supported in part by the Spanish Government under Project TIN2011-25606. The work of N. Onkarappa was supported in part by the Catalan Government through the Agency for Management of University and Research Grants (AGAUR) under an FI Grant.

References

1. http://vision.middlebury.edu/flow/
2. http://sintel.is.tue.mpg.de/results/
3. http://www.cvlibs.net/datasets/kitti/
4. http://www.cvc.uab.es/adas
5. Adato Y, Zickler T, Ben-Shahar O (2011) A polar representation of motion and implications for optical flow. In: IEEE conference on computer vision and pattern recognition. Colorado Springs, CO, pp 1145–1152
6. Baker S, Scharstein D, Lewis JP, Roth S, Black MJ, Szeliski R (2011) A database and evaluation methodology for optical flow. Int J Comput Vis 92(1):1–31
7. Barron JL, Fleet DJ, Beauchemin SS (1994) Performance of optical flow techniques. Int J Comput Vis 12(1):43–77
8. Butler DJ, Wulff J, Stanley GB, Black MJ (2012) A naturalistic open source movie for optical flow evaluation. In: European Conference on Computer Vision (6), vol 7577. Florence, pp 611–625
9. Geiger A, Lenz P, Urtasun R (2012) Are we ready for autonomous driving? the KITTI vision benchmark suite. In: Computer Vision and Pattern Recognition (CVPR). Providence
10. Horn BKP, Schunk BG (1981) Determining optical flow. Artif Intell 17:185–203
11. Liu C, Freeman WT, Adelson EH, Weiss Y (2008) Human-assisted motion annotation. In: IEEE conference on computer vision and pattern recognition. Anchorage
12. Lucas BD, Kanade T (1981) An iterative image registration technique with an application to stereo vision (DARPA). In: DARPA image understanding workshop, pp 121–130
13. Mac Aodha O, Humayun A, Pollefeys M, Brostow GJ (2013) Learning a confidence measure for optical flow. IEEE Trans Pattern Anal Mach Intell 35(5):1107–1120

14. McCane B, Novins K, Crannitch D, Galvin B (2001) On benchmarking optical flow. Comput Vis Image Underst 84(1):126–143
15. Meister S, Kondermann D (2011) Real versus realistically rendered scenes for optical flow evaluation. In: 14th ITG Conference on Electronic Media Technology (CEMT). Dortmund, pp 1–6
16. Otte M, Nagel HH (1994) Optical flow estimation: advances and comparisons. In: European conference on computer vision (1), Lecture notes in computer science, vol 800. Stockholm, pp 51–60
17. Sun D, Roth S, Black MJ (2010) Secrets of optical flow estimation and their principles. In: IEEE conference on computer vision and pattern recognition. San Francisco, pp 2432–2439
18. Vaudrey T, Rabe C, Klette R, Milburn J (2008) Differences between stereo and motion behaviour on synthetic and real-world stereo sequences. In: Image and vision computing new zealand. Christchurch, pp 1–6
19. Wedel A, Pock T, Zach C, Cremers D, Bischof H (2008) An improved algorithm for TV-L1 optical flow. In: Dagstuhl motion workshop. Dagstuhl Castle, pp 23–45
20. Xu L, Jia J, Matsushita Y (2012) Motion detail preserving optical flow estimation. IEEE Trans Pattern Anal Mach Intell 34(9):1744–1757

Naveen Onkarappa received the B.Sc. degree from Kuvempu University, Shimoga, India, in 1999 and the M.Sc. degree in computer science and the M.Sc.Tech. degree (by research) in computer science and technology from the University of Mysore, Mysore, India, in 2001 and 2007, respectively. He obtained Ph.D. degree in computer vision from Computer Vision Center, Autonomous University of Barcelona, Barcelona, Spain, in 2013. From 2001 to 2005, he was a Guest Lecturer with Kuvempu University and the University of Mysore. From 2007 to 2009 he was with HCL Technologies, Bangalore, India. He was a member of the Advanced Driver Assistance Systems Group, Computer Vision Center, Autonomous University of Barcelona. His research interests include optical-flow estimation and its applications to driver assistance systems.

Angel D. Sappa received the electromechanical engineering degree from National University of La Pampa, General Pico, Argentina, in 1995 and the Ph.D. degree in industrial engineering from the Polytechnic University of Catalonia, Barcelona, Spain, in 1999. In 2003, after holding research positions in France, the UK, and Greece, he joined the Computer Vision Center, where he is currently a senior researcher. He is a member of the Advanced Driver Assistance Systems research group. His research interests span a broad spectrum within the 2D and 3D image processing. His current research focuses on stereoimage processing and analysis, 3D modeling, and dense optical flow estimation. He is a senior member of the IEEE.