

## Camera-based vehicle velocity estimation from monocular video

Moritz Kampelmühler<sup>1</sup>

Michael G. Müller<sup>2</sup>

Christoph Feichtenhofer<sup>1</sup>

<sup>1</sup>Institute of Electrical Measurement and Measurement Signal Processing

<sup>2</sup>Institute of Theoretical Computer Science

Graz University of Technology

kampelmuehler@student.tugraz.at, mueller@igi.tugraz.at, feichtenhofer@tugraz.at

**Abstract.** *This paper documents the winning entry at the CVPR2017 vehicle velocity estimation challenge. Velocity estimation is an emerging task in autonomous driving which has not yet been thoroughly explored. The goal is to estimate the relative velocity of a specific vehicle from a sequence of images. In this paper, we present a light-weight approach for directly regressing vehicle velocities from their trajectories using a multilayer perceptron. Another contribution is an explorative study of features for monocular vehicle velocity estimation. We find that light-weight trajectory based features outperform depth and motion cues extracted from deep ConvNets, especially for far-distance predictions where current disparity and optical flow estimators are challenged significantly. Our light-weight approach is real-time capable on a single CPU and outperforms all competing entries in the velocity estimation challenge. On the test set, we report an average error of 1.12 m/s which is comparable to a (ground-truth) system that combines LiDAR and radar techniques to achieve an error of around 0.71 m/s.*

### 1. Introduction

Camera sensors provide an inexpensive yet powerful alternative to range sensors based on LiDAR or radar. While LiDAR systems can provide very accurate measurements, they may also malfunction under adverse environmental conditions such as fog, snow, rain or even exhaust gas fumes [30, 14]. Arguably vision based sensing is more closely related to how humans engage in driving situations and it should thus be possible to solve any task in autonomous driving based on visual input.

This work addresses monocular vehicle velocity estimation, an emerging task in autonomous driving




Figure 1: A sample image from a training sequence. Velocity and position ground truth are provided for the vehicles surrounded by the green bounding boxes.

which has not yet been thoroughly explored. The specific task, which forms the base for this work, was introduced as the Autonomous Driving Velocity Estimation Challenge<sup>1</sup> at CVPR2017. The goal is to estimate the relative velocity of a specific vehicle from a sequence of monocular RGB images to aid autonomous driving algorithms such as for example collision avoidance [1] or adaptive cruise control [20]. Figure 1 shows an example image from the data<sup>1</sup>.

Vehicle velocity estimation as such is not a new subject of interest, since it is extensively studied in the context of traffic surveillance [17, 4], where, however, a stationary camera is employed. Under the restriction of a fixed camera pose the problem becomes significantly less complex, since with a calibrated camera system angular measurements can be obtained and from these measurements velocity estimates can readily be established. In contrast in our case the observer resides on a moving platform and inferring velocity in a similar fashion would re-

<sup>1</sup><http://benchmark.tusimple.ai/#/t/2>

quire additional information such as camera pose, ego-motion and foreground-background segmentation. Very recent research [41] shows that estimating ego-motion as well as disparity maps from monocular camera images by means of *structure from motion* is indeed possible, but still limited. Semantic segmentation of scenes, which is a fundamental problem in computer vision, has also more recently been tackled using deep neural networks [6, 24].

In a more general sense the given task can be seen as a lightweight version of *object scene flow* as for example provided in the KITTI benchmark [11, 28]. Object scene flow aims at estimating dense 3D motion fields, which in their temporal evolution carry highly valuable information about the geometric constellation of a given scene. Recent approaches [36, 35] yield impressive results, but they rely on the availability of stereo image data. Furthermore, they come at the price of very high computational cost, such that the estimation for a temporal frame pair might take 5-10 minutes on a single CPU core. In autonomous driving scenarios computational resources are in general highly limited [13], which makes object scene flow currently not practically feasible.

In this work we adopt recent deep learning architectures [18, 12] for depth and motion estimation to leverage a mapping of the video input into a beneficial feature space for learning from the few training samples provided. Our approach employs a two-stage process for monocular velocity estimation. In a first step we extract vehicle tracks as well as dense depth and optical flow information, followed by locally aggregating these depth and motion cues at the tracked vehicle locations and concatenating over the temporal dimension. After this feature extraction procedure we use the spatiotemporal depth, flow and location features to train a fully connected regression network for velocity estimation of the respective vehicles.

Further on we conduct an extensive ablation study to investigate the impact of the individual features and combinations thereof on the regression performance as well as on the runtime of the estimation. We show that a light weight implementation can achieve excellent results, and that leveraging deep motion and depth cues does not necessarily improve performance for this task on the given data.

## 2. Related Work

**Tracking.** Object tracking is one of the fundamental problems in computer vision and has been extensively studied [40] and applied in many different tasks.

*Median Flow* [21] is a method building on top of the *Lucas-Kanade* [25] method, which is an early optical flow algorithm operating on local intensity changes. This method is extended by a Forward-Backward error, which denotes the deviation between the trajectories obtained by tracking from  $I_{t-1} \rightarrow I_t$  and  $I_t \rightarrow I_{t-1}$ . Using Forward-Backward error, robust predictions can be identified.

The *Multiple Instance Learning* tracker [2] first transforms the image into an appropriate feature space, and uses a classifier as well as a motion model to determine the presence of an object in a frame, which is referred to as tracking by detection.

More recent methods like [15] employ convolutional neural networks to learn motion and appearance of objects. The feature maps of higher convolutional layers provide robust and accurate appearance representations, but lack spatial resolution. Lower layers on the other hand provide higher spatial resolution and less refined appearance representations. This hierarchical structure is used in [26] by inferring responses of correlation filters on each corresponding layer pair.

**Monocular Depth Estimation.** Estimating the depth information of a scene seen from a given angle using only a single camera is not a well-posed problem. Some of the methods tackling this problem [31, 39] apply supervised learning regimes requiring ground truth depth data. Since acquiring such ground truth data with sufficient accuracy requires immense effort, several methods are being developed that require little to no supervision.

In [23] a semi-supervised approach is described that provides a fusion between using sparse ground truth data from a LiDAR sensor and stereo view synthesis, *i.e.* estimating one image in a stereo pair from the other, to infer dense depth maps. Others [12, 38] in turn rely solely on stereo as a supervision signal, which comes with the benefit of easily available or obtainable data.

Some recent work [41, 10] shows the capability of inferring single image depth from monocular video only. This is achieved by leveraging the (small) temporal motion of the camera and its thus changing pose to learn from multiple views of the scene. Via

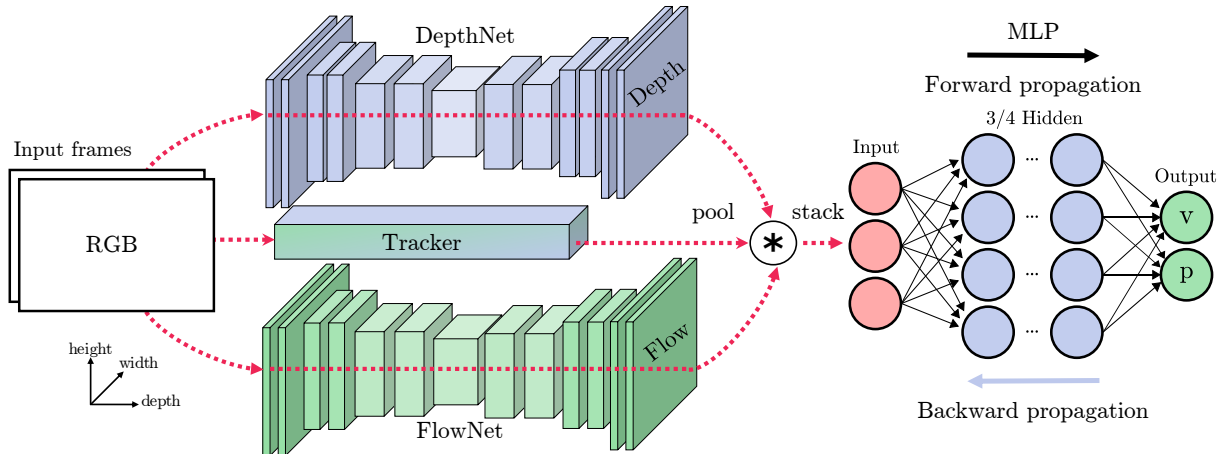


Figure 2: Overview of our overall architecture (see Section 3 for details). Our proposed light-weight architecture only uses the trajectory features estimated by the tracker.

novel view synthesis the future camera frames can be used as a self-supervision signal. The mapping thus learned implicitly carries information about the 3D scene geometry.

**Optical Flow Estimation.** Optical flow is widely used in computer vision, *e.g.* for video object detection [42], to quantify pixel-wise motion in between frames of a moving scene. Traditional methods [16, 37] employ variational motion estimation approaches, that regard the estimation of pixel level correspondences between frames as an optimization problem.

With the growing interest in deep learning also optical flow estimation is now often successfully treated as a supervised learning problem [7]. This is achieved by employing convolutional neural networks for feature extraction and aggregation, followed by an 'upconvolutional' network, which concatenates the feature maps from the corresponding convolutional layers and jointly applies transposed convolution to increase spatial resolution. Further improvements on this approach have since been made [18] that provide improved performance and robustness as well as scalability.

### 3. Vehicle velocity estimation

We present an explorative study over features for vehicle velocity estimation from monocular camera videos and discuss effectiveness as well as significance of the methods employed. The task is to estimate the relative velocity as well as position of given vehicles seen in short dashcam video snippets. Our overall approach, shown in Figure 2, consists of a two stage process: First, features subsidiary to the

task are extracted to subserve the second stage, which is a light-weight Multilayer Perceptron (MLP) architecture working on these features to regress velocity and positions of vehicle instances.

#### 3.1. Feature Extraction

For the estimation of vehicle velocity, the raw RGB video data is first transformed into a beneficial feature space. We use three feature types of complementary nature: Vehicle tracks (*i.e.* trajectories of the 2D object outline over time), depth (*i.e.* disparity estimates from monocular imagery) and motion (*i.e.* optical flow estimates between consecutive frames). The remainder of this section describes the specific algorithmic instances used for extracting these cues.

**Tracks.** For a given vehicle defined by a bounding box in a single frame, tracking over the temporal extent of the input serves for all further processing steps. A variety of well functioning tracking algorithms are readily available in literature. Since we aim for a lightweight tracker that should precisely localize the object outline, we employ fast trackers that operate directly at the pixel level, the *Median Flow* [21] and *MIL* [2] trackers, both implemented in the `OpenCV` library [3]. The *Median Flow* tracker comes with the benefit of being able to adapt bounding boxes over the trajectories, and provides a tight bounding box that can be a very useful feature when estimating the relative velocity of the objects. However, this tracker is unstable for occlusions since it employs forward-backward tracking. Whenever *Median Flow* detects a tracking failure the missing bounding boxes are substituted with *MIL* tracks.

**Depth.** For dense depth map prediction we employ a recently described deep architecture [12], that




Figure 3: Optical flow (left, Middlebury flow encoding), bounding box and depth map (right, darker values represent larger distances) of a close range sample vehicle overlaid over the RGB image.

learns monocular depth map prediction via novel view synthesis in a stereo environment. This is achieved by synthesizing *e.g.* the left camera image from the right, where the left camera image is used as supervision signal. The warping operation that is thus learned implicitly carries information on the disparities in the source image. We use a model pre-trained on *KITTI* [11] and *Cityscapes* [5] stereo images for predicting the dense disparity maps (note that this model is trained without disparity ground-truth). Owing to architectural constraints and limited computational capacity the RGB input images are resized to 512x256p.

**Motion.** Finally, we extract motion information by extracting dense optical flow maps using a state-of-the-art neural network architecture, *FlowNet2* [18]. *FlowNet2* treats optical flow estimation as a supervised learning problem, where a convolutional neural network is trained on a volume of two stacked input frames with ground truth optical flow as a supervision signal. In our case we use a *FlowNet2* architecture pre-trained on the synthetic *Flying Chairs* [8] dataset to calculate 39 dense  $u, v$  flow maps from 512x256p input images. A sample of the extracted feature maps is shown in Figure 3. In both feature maps the vehicle can be segmented from the background, but the capabilities are limited to close ranges.

**Transformation into feature space.** The depth and motion cues are computed globally and further processed to serve as lightweight input features for a regression model. This is achieved by locally aggregating the dense predictions within the bounding boxes established by the tracking stage. The local aggregation is achieved by calculating the mean over the estimates within each tracked bounding box, after shrinking the box by 10% in width and height, which reduces the variance since flow and depth cues tend to be inaccurate at the object boundaries. Sub-

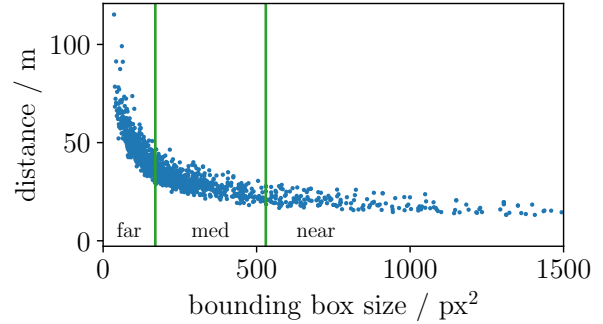


Figure 4: Split of distance categories for neural network. For the labeled training set, the true distance is known. The empirically chosen boundaries are shown in green.

sequently, the aggregated feature vectors are temporally smoothed using a Gaussian kernel of width 5, which is chosen in correspondence to the frame skip of 5 in the learning stage. For optical flow, this procedure is carried out individually for each the horizontal and vertical component  $u, v$ . The temporal smoothing provides robustness to short-term deviations of the camera orientation (*e.g.* caused by bumps when driving on a highway).

### 3.2. Neural Network Model

These pre-computed features allow us to use relatively shallow fully-connected neural networks. This is especially advantageous for the task at hand, in which the number of given training examples is relatively small and learning can take advantage of highly abstract features. We employ a comparatively simple, rather small and thus efficient 4-layer MLP architecture to regress from feature space to vehicle velocities.

**Data Split.** The relationship of the pre-computed features to the learning output is highly nonlinear. Consider for instance the size of bounding boxes around other vehicles: they rapidly decrease when the vehicle is close to the camera, while remaining more or less constant if the vehicle is far away regardless of the velocity. To facilitate learning in these distinct regimes, we split the data set into three disjoint parts, based on the distance from the observer (near/mid/far), and train separate models for each. Figure 4 shows the bounding box size as a function of the distance from the observer for the training set, as well as our chosen splits. As indicator for the relative distance of the target vehicle we use the bounding box area of the last frame that the vehicle is shown.

Features	(a)	(b)	(c)	(d)	(e)	(f)	(g)
Tracking	✓				✓	✓	✓
Motion		✓		✓	✓		✓
Depth			✓	✓		✓	✓
$E_V$	1.86	4.68	4.59	3.74	1.90	<b>1.83</b>	1.86
$E_{V,near}$	0.93	1.42	0.87	<b>0.66</b>	0.91	0.70	0.77
$E_{V,medium}$	<b>0.76</b>	4.84	4.88	4.29	0.87	0.86	0.90
$E_{V,far}$	3.43	6.87	7.72	6.20	3.24	3.23	<b>3.06</b>

Table 1: Best results achieved in ablation study for each individual range.

Although there is some variance near the borders between different sets, this split simplifies the learning task compared to a single model that has to work across all distances.

**Architecture.** On each of these three sets, we performed cross-validation to obtain a good network architecture. The resulting network topologies are (number of hidden layers  $\times$  number of units per layer)  $3 \times 40$  (near),  $4 \times 60$  (medium), and  $4 \times 70$  (far). Extensive validation experiments revealed that best results could be obtained with concatenated rectified linear units (CReLU) [32] as hidden unit activations; thus, all layer output sizes are twice the input size. The output layer activation is linear. We use an L2 loss to train the network by using the ground-truth velocity values, obtained from a combination of LiDAR and radar sensors.

**Training.** Each individual network is trained on the MSE between network output and targets for 2000 epochs using minibatches of 50 samples. For regularization, weight decay of  $10^{-5}$  and Dropout [34] of 0.2 were used. We train using the ADAM optimizer [22].

To make use of all available training data, we use a partitioning scheme similar to k-fold cross-validation and train multiple networks per distance set. Each of them is split up into five partitions. Four of these are used as training set, the fifth is used for validation. After training for 2000 epochs on each possible combination, the model with the lowest validation error is saved. This results in  $3 \times 5$  models for the entire dataset. Note that the number of examples per neural network is quite small, so overfitting may occur easily as the validation score can no longer be regarded as an accurate estimate of the error on the test set in most cases.

The training data not only contains targets for the target vehicle velocity, but also for the relative vehicle position. In other machine learning domains, it is well-known that auxiliary targets can increase

the learning performance (see *e.g.* for reinforcement learning [29, 19, 9]). We also make use of all targets during training and simultaneously regress for vehicle velocities and distance. Early stopping is performed using only the performance on the actual targets. This leads to slight increases in performance.

**Model Averaging.** Above, we described how  $3 \times 5$  partitions are trained. When evaluating the test set, we split the data according to the computed bounding box areas using the same procedure as for the training set. Then, the average over all five models for the respective distance is computed. This gives the final estimation for the relative vehicle velocity.

## 4. Experiments

### 4.1. Dataset

The provided dataset includes 1074 driving sequences in freeway traffic, recorded by a single HD (1280x720p) camera, each 40 frames long captured at 20fps, as well as camera calibration matrices. For each sequence specific vehicles are annotated with a bounding box as well as position and velocity in both  $x$  and  $y$  coordinates in the last frame only, resulting in a total of 1442 annotated vehicles. For evaluation the individual vehicles are classified into three clusters according to their ground truth relative distance  $d$  in the last frame.  $d < 20$  m is considered near range (12% of samples),  $20 \text{ m} \leq d < 45$  m medium range (65% of samples) and  $d \geq 45$  m far range (23% of samples). For each of those ranges different difficulties arise in the estimation. While in near range examples the perspective on vehicles can shift drastically in between instances and over time for individual instances, for far range samples the estimation is limited by the pixel resolution of the data.

From the provided data a method should be developed, which is able to infer velocity as well as position of vehicles specified by a given bounding box in the last frame. For evaluation a test set consisting of 269 clips, or 375 vehicle tracks is provided, which

is structurally identical to the training data, with the only difference being the absence of ground truth position and velocity data.

## 4.2. Ablation study

To investigate the impact of the features used in our initial approach on velocity estimation accuracy, we have conducted an ablation study to test combinations of all features and activation functions.

**Validation split.** In order to be able to evaluate the performance of our approach we had to first split the available training data into a training and validation set. Earlier we used 80/20 random splits, but for more robust evaluation we decided to take into account the distribution of near, medium and far range samples (12/65/23%) and also the similarity of the test sequences. Since the training samples are sequences from a fixed set of drives, we decided to split training and validation splits such that the validation samples stem from unseen drives.

To achieve this we compute the 4096 dimensional *VGGNet fc2* feature vector [33], reduce it to 2D space using *t-sne* [27] and then cluster the resulting features into 7 clusters using k-means. We have then chosen a cluster consisting of 10% of the total training samples with a (14/63/21%) near, medium, far range distribution as validation set. The remaining data is used for training after defining a fixed 5-fold split for cross validation.

**Training.** For each of the 5 train/test splits, 10 models with fixed  $3 \times 40$  topologies are trained according to the same paradigm as described above. Here, training was stopped early if the velocity mean squared error (1) did not improve in the course of 500 iterations. The ablation was performed both training models on the whole data as well as training individual models on near, medium and far data splits.

**Results.** The results of the ablation study described above are shown in Table 1. It can be seen that depth and motion cues can indeed help improve results of vehicle velocity estimation, but their impact is limited by range. Both depth and optical flow estimates show significantly degrading performances for distances larger than 20m (above near range). Within near range on the other hand, where rapid changes in perspective can occur, they are superior in performance to using only tracking.

For medium range, where vehicles are predominantly viewed from their rear only, tracking cues only yield superior performance. For far range the pixel

level appearance of vehicles changes so slightly at different velocities, that a robust estimation of their velocity can not be achieved by the method used.

	Motion	Depth	Tracking	MLP
Hardware	GPU	GPU	CPU	GPU
Timing	344 ms	69 ms	7 ms	3 ms

Table 2: Runtime of the individual stages of our method for a single frame on Intel®Core™ i7-5930K, 32GB, NVIDIA Titan X. For depth and motion estimation inference times only are indicated.

Using tracking cues only is also the most efficient, as shown in Table 2. Inference of our model is generally very fast and thus the feature extraction is the biggest performance bottleneck of our approach. While using motion, depth and tracking cues provides a theoretical estimation rate of 2 frames per second, with tracking cues 100 frames can be estimated in one second.

## 4.3. Challenge results

On the leaderboard of the CVPR2017 Autonomous Driving velocity estimation challenge, the individual entries are ranked by overall average velocity mean squared error. For each given set of samples  $C$ , in this case the ranges near, medium and far, the error is evaluated according to (1) and then averaged (2).

$$E_{V,C} = \frac{1}{|C|} \sum_{c \in C} \|V_c - \hat{V}_c\|^2 \quad (1)$$

$$E_V = \frac{E_{V,near} + E_{V,med} + E_{V,far}}{3} \quad (2)$$

In correspondence to the given evaluation metric, the results achieved on the test set are depicted in Table 3. The entry *Ours full* is the winning approach described above, using individual models trained on all three ranges separately, while incorporating tracking, flow and depth features. *Ours tracking* are the best performing models using only tracking cues obtained in the ablation study, which are again separate models for each range, but trained on the whole data. Notably, the ground-truth accuracy is at around 0.71 m/s which is obtained from a combination of LiDAR and radar sensors (this number was communicated by the challenge organizers). Therefore, our average estimation 1.12 m/s error (corresponding to 1.25 m<sup>2</sup>/s<sup>2</sup> MSE) is relatively close to a LiDAR and radar solu-

Figure 5: Qualitative results on the test data, near medium and far range (f.l.t.r.). In each example one sequence from the test set (*i.e.* 2 seconds of video at 20fps, see Section 4.1) is shown, annotated with the estimated velocity in m/s for  $(x, y)$  coordinates. **Note:** View with Acrobat or PDF-Xchange reader for animation.

tion, but uses only videos recorded from a monocular dash cam.

	$E_V$	$E_{V,near}$	$E_{V,med}$	$E_{V,far}$
<b>Ours tracking*</b>	<b>1.25</b>	<b>0.12</b>	<b>0.54</b>	3.11
<b>Ours full</b>	1.30	0.18	0.66	<b>3.07</b>
Rank2 team	1.50	0.25	0.75	3.50
Rank3 team	2.90	0.55	2.21	5.94
Rank4 team	3.54	1.46	2.74	6.42

Table 3: Challenge leaderbord top 5. **Ours full** is the winning approach. **Ours tracking** denotes the best performing tracking only model. \*submitted post deadline

## 5. Conclusion

This paper documents the winning entry at the CVPR2017 vehicle velocity estimation challenge. We have proposed a light-weight approach for directly regressing vehicle velocities from their tracks in monocular video sequences. By comparing complementary features for vehicle velocity estimation, we find that light-weight trajectory based features outperform depth and motion cues extracted from deep ConvNets. Our approach is real-time capable on a single CPU and outperforms all competing entries in the velocity estimation challenge. Future work shall address an end-to-end system for joint tracking and estimation.

**Acknowledgments.** We are grateful for discussions with Axel Pinz. The GPUs used for this research were donated by NVIDIA.

## References

- [1] R. Aufrère, J. Gowdy, C. Mertz, C. Thorpe, C.-C. Wang, and T. Yata. Perception for collision avoidance and autonomous driving. *Mechatronics*, 13(10):1149–1161, 2003. 1
- [2] B. Babenko, M.-H. Yang, and S. Belongie. Visual tracking with online multiple instance learning. In *Proc. CVPR*, pages 983–990. IEEE, 2009. 2, 3
- [3] G. Bradski. The OpenCV Library. *Dr. Dobb’s Journal of Software Tools*, 2000. 3
- [4] B. Coifman, D. Beymer, P. McLauchlan, and J. Malik. A real-time computer vision system for vehicle tracking and traffic surveillance. *Transportation Research Part C: Emerging Technologies*, 6(4):271–288, 1998. 1
- [5] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. CVPR*, pages 3213–3223, 2016. 4
- [6] J. Dai, K. He, and J. Sun. Instance-aware semantic segmentation via multi-task network cascades. In *Proc. CVPR*, pages 3150–3158, 2016. 2
- [7] 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 *Proc. CVPR*, pages 2758–2766, 2015. 3
- [8] A. Dosovitskiy, P. Fischer, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. v.d. Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. In *Proc. ICCV*, 2015. 4
- [9] A. Dosovitskiy and V. Koltun. Learning to act by predicting the future. In *Proc. ICLR*, 2016. 5
- [10] R. Garg, G. Carneiro, and I. Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *Proc. ECCV*, pages 740–756. Springer, 2016. 2
- [11] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Proc. CVPR*, 2012. 2, 4
- [12] C. Godard, O. Mac Aodha, and G. J. Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proc. CVPR*, 2017. 2, 3
- [13] A. Grzywaczewski. Training ai for self-driving vehicles: the challenge of scale, 2017. 2
- [14] S. Hasirlioglu, A. Riener, W. Ruber, and P. Wintersberger. Effects of exhaust gases on laser scanner data

- quality at low ambient temperatures. In *Intelligent Vehicles Symposium (IV), 2017 IEEE*, pages 1708–1713. IEEE, 2017. 1
- [15] D. Held, S. Thrun, and S. Savarese. Learning to track at 100 fps with deep regression networks. In *Proc. ECCV*, pages 749–765. Springer, 2016. 2
- [16] B. K. Horn and B. G. Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. 3
- [17] J.-W. Hsieh, S.-H. Yu, Y.-S. Chen, and W.-F. Hu. Automatic traffic surveillance system for vehicle tracking and classification. *IEEE Transactions on Intelligent Transportation Systems*, 7(2):175–187, 2006. 1
- [18] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *Proc. CVPR*, Jul 2017. 2, 3, 4
- [19] M. Jaderberg, V. Mnih, W. M. Czarnecki, T. Schaul, J. Z. Leibo, D. Silver, and K. Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016. 5
- [20] R. K. Jurgen. Adaptive cruise control. Technical report, SAE Technical Paper, 2006. 1
- [21] Z. Kalal, K. Mikolajczyk, and J. Matas. Forward-backward error: Automatic detection of tracking failures. In *Proc. ICPR*, pages 2756–2759. IEEE, 2010. 2, 3
- [22] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *Proc. ICLR*, 2014. 5
- [23] Y. Kuznetsov, J. Stückler, and B. Leibe. Semi-supervised deep learning for monocular depth map prediction. In *Proc. CVPR*, 2017. 2
- [24] Y. Li, H. Qi, J. Dai, X. Ji, and Y. Wei. Fully convolutional instance-aware semantic segmentation. In *Proc. CVPR*, 2016. 2
- [25] B. D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision (darpa). In *Proceedings of the 1981 DARPA Image Understanding Workshop*, pages 121–130, 1981. 2
- [26] C. Ma, J.-B. Huang, X. Yang, and M.-H. Yang. Hierarchical convolutional features for visual tracking. In *Proc. CVPR*, pages 3074–3082, 2015. 2
- [27] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *JMLR*, 9(Nov):2579–2605, 2008. 6
- [28] M. Menze and A. Geiger. Object scene flow for autonomous vehicles. In *Proc. CVPR*, June 2015. 2
- [29] P. Mirowski, R. Pascanu, F. Viola, H. Soyer, A. Ballard, A. Banino, M. Denil, R. Goroshin, L. Sifre, K. Kavukcuoglu, et al. Learning to navigate in complex environments. *arXiv preprint arXiv:1611.03673*, 2016. 5
- [30] R. Rasshofer, M. Spies, and H. Spies. Influences of weather phenomena on automotive laser radar systems. *Advances in Radio Science*, 9(B. 2):49–60, 2011. 1
- [31] A. Saxena, S. H. Chung, and A. Y. Ng. Learning depth from single monocular images. In *NIPS*, pages 1161–1168, 2006. 2
- [32] W. Shang, K. Sohn, D. Almeida, and H. Lee. Understanding and improving convolutional neural networks via concatenated rectified linear units. *CoRR*, abs/1603.05201, 2016. 5
- [33] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6
- [34] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *JMLR*, 15(1):1929–1958, 2014. 5
- [35] C. Vogel, S. Roth, and K. Schindler. View-consistent 3d scene flow estimation over multiple frames. In *Proc. ECCV*, pages 263–278. Springer, 2014. 2
- [36] C. Vogel, K. Schindler, and S. Roth. Piecewise rigid scene flow. In *Proc. CVPR*, pages 1377–1384, 2013. 2
- [37] A. Wedel, D. Cremers, T. Pock, and H. Bischof. Structure-and motion-adaptive regularization for high accuracy optic flow. In *Proc. CVPR*, pages 1663–1668. IEEE, 2009. 3
- [38] J. Xie, R. Girshick, and A. Farhadi. Deep3d: Fully automatic 2d-to-3d video conversion with deep convolutional neural networks. In *Proc. ECCV*, pages 842–857. Springer, 2016. 2
- [39] D. Xu, E. Ricci, W. Ouyang, X. Wang, and N. Sebe. Multi-scale continuous crfs as sequential deep networks for monocular depth estimation. In *Proc. CVPR*, 2017. 2
- [40] A. Yilmaz, O. Javed, and M. Shah. Object tracking: A survey. *Acm computing surveys (CSUR)*, 38(4):13, 2006. 2
- [41] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe. Unsupervised learning of depth and ego-motion from video. In *Proc. CVPR*, 2017. 2
- [42] X. Zhu, Y. Wang, J. Dai, L. Yuan, and Y. Wei. Flow-guided feature aggregation for video object detection. In *Proc. ICCV*, 2017. 3