

Pixel-Accurate Depth Evaluation in Realistic Driving Scenarios

Tobias Gruber^{1,3} Mario Bijelic^{1,3} Felix Heide^{2,4} Werner Ritter¹ Klaus Dietmayer³
¹Daimler AG ²Algolux ³Ulm University ⁴Princeton University

Abstract

This work introduces an evaluation benchmark for depth estimation and completion using high-resolution depth measurements with angular resolution of up to 25" (arcsecond), akin to a 50 megapixel camera with per-pixel depth available. Existing datasets, such as the KITTI benchmark [12], provide only sparse reference measurements with an order of magnitude lower angular resolution – these sparse measurements are treated as ground truth by existing depth estimation methods. We propose an evaluation methodology in four characteristic automotive scenarios recorded in varying weather conditions (day, night, fog, rain). As a result, our benchmark allows us to evaluate the robustness of depth sensing methods in adverse weather and different driving conditions. Using the proposed evaluation data, we demonstrate that current stereo approaches provide significantly more stable depth estimates than monocular methods and lidar completion in adverse weather. Data and code are available at <https://github.com/gruberto/DepthBenchmark>.

1. Introduction

3D scene understanding is one of the key challenges for safe autonomous driving, and the critical depth measurement and processing methods are a very active areas of research. Depth information can be captured using a variety of different sensing modalities, either passive or active. Passive methods can be classified into stereo methods [4, 18, 22, 41], which, inspired by the human visual system, extract depth from parallax in intensity images, structure from motion (SfM) [23, 30, 56, 58, 64], and monocular depth prediction [5, 9, 14, 27, 45]. Monocular depth estimation methods attempt to extract depth from cues such as defocus [55], texture gradient and size perspective from a single image only. All of these passive sensing systems suffer in low light and at night, when the measured intensity is too low to robustly match image content. Active sensing methods, such as lidar systems and time-of-flight (ToF) cameras, overcome this challenge by relying on active illumination for depth measurements. Specifically, lidar sys-

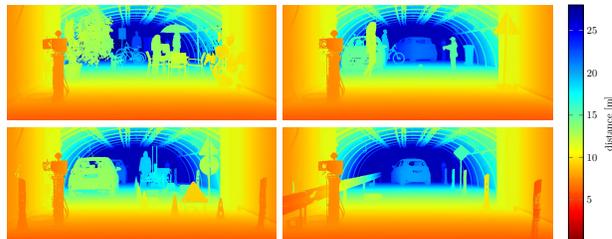


Figure 1: We propose a depth evaluation framework using ground truth depth data acquired at an angular resolution of up to 25" (arcsecond). We evaluate depth sensing under realistic weather situations and different automotive scenarios, including a pedestrian zone, residential area, construction, area and highway.

tems [49] achieve large distances by focusing light into multiple beams which are mechanically scanned. This sequential acquisition fundamentally limits the angular resolution by the scanning mechanics, prohibiting semantic scene understanding at large distances where only a coarse sample distribution is available. As a result, a variety of algorithms for depth completion [6, 21, 31] have been proposed recently. Existing correlation ToF cameras [16, 24, 28] or structured light cameras [1, 36, 37, 48], provide accurate high-resolution depth for close ranges indoors but suffer at long ranges in outdoor scenes due to strong ambient illumination and modulation frequency limitations.

In order to evaluate the performance of such diverse sensing approaches for autonomous driving, empirical datasets under realistic road conditions and meaningful metrics are required. Captured with previous generation hardware, the NYUdepth v2 [53] and KITTI [12] datasets are established for developing and evaluating depth estimation methods. While NYUdepth v2 contains a large variety of close-range indoor scenes captured by a Microsoft Kinect RGB-D camera [63], the KITTI dataset provides real-world street view scenarios recorded with a stereo camera and a lidar sensor. Although the sparse depth measurements of the KITTI dataset have been accumulated over multiple laser scans [57], the resulting “ground truth” depth provides only 300" horizontal and 700" vertical angular sampling. To detect the legs of a pedestrian at 150m distance an angular resolution of less than 25" would be required. In this work,

we present an evaluation framework with pixel-wise annotated ground truth depth, as visualized in Figure 1, at an angular resolution of 25° . This resolution corresponds to a 50 mega-pixel camera (similar to KITTI [12]: $4.65 \mu\text{m}$ pitch, 4.4 mm focal length) which enables semantic understanding tasks at large distances, such as pedestrian detection. We also record and provide stereo and lidar measurements acquired with recent state-of-the-art automotive sensors.

The robustness of the sensing and processing chain is critical for autonomous driving, which mandates reliability in adverse weather situations such as rain, fog and snow. Note that the evaluation in these corner cases is particularly challenging because lidar sensors *fail in severe adverse weather*. For example, in dense fog or snow, the first peak lidar measurements are unusable due to severe back-scatter in these scenarios. As a result, existing driving datasets [12, 19, 60] do not cover these severe conditions. To acquire accurate depth in challenging weather situations, and with reproducible scenes and conditions, we record sensor data with the proposed sensor setup in a weather chamber [7] that provides reproducible fog and rain. Using pixel-wise ground truth depth from clear conditions, the proposed approach allows us to accurately evaluate the performance under varying fog visibilities and rain intensities.

Specifically, we make the following contributions:

- We introduce an automotive depth evaluation dataset (1600 samples) in adverse weather conditions with high-resolution annotated ground truth of angular resolution 25° – an order of magnitude higher than existing lidar datasets with angular resolution of 300° .
- We evaluate the performance state-of-the-art algorithms for stereo depth, depth from mono, lidar and sparse depth completion, in reproducible, finely adjusted adverse weather situations.
- We demonstrate that stereo vision performs significantly more stable in fog than lidar and monocular depth estimation approaches.

2. Related Work

Depth Estimation and Depth Completion Depth estimation algorithms can be categorized based on their input data. While SfM approaches [23, 30, 56, 58, 64] rely on sequentially captured image data, multi-view depth estimation [4, 18, 22, 41] uses at least two different views of a scene which are simultaneously captured. Monocular depth estimation methods [5, 9, 14, 27, 45] tackle the severely ill-posed depth reconstruction by monocular cues such as texture variation, gradients, object size, defocus, color or haze. Over the last years, convolutional neural networks (CNNs) have been shown to be well-suited for both SfM, multi-view and monocular approaches, and they can

be trained supervised by large RGB-D indoor datasets acquired by consumer ToF cameras. Since the acquisition of ground truth depth in large outdoor environments is challenging, semi-supervised [26] or even self-supervised approaches [2, 11, 14, 40] have been proposed, tackling this challenge by solving proxy tasks such as stereo matching. Another body of work focuses on completing sparse lidar point clouds. Existing depth completion methods rely on contextual information from RGB images to obtain high-resolution depth [6, 21, 31].

Depth Datasets The development and evaluation of depth estimation algorithms require a large amount of representative data, particularly for learned estimation methods. Scharstein and Szeliski [47] provided the Middlebury data set as an early testing environment for quantitative evaluation of stereo algorithms, where the ground truth was obtained by structured light. While the Middlebury dataset contained, in the first version, only two samples, subsequent datasets, such as Make 3D [46] provided around 500 samples with ground truth measured by a custom-build 3D scanner, though with lower resolution of 55×305 . Facilitated by consumer depth cameras such as the Microsoft Kinect [63], a number of depth data sets for indoor scenes have been proposed [8, 51, 52, 53, 54]. In particular, the NYUdepth v2 data set [53] is a widely used dataset with around 1500 samples. However, due to the limitations of consumer depth cameras in severe ambient light and modulation frequency limitations, these data sets only include indoor scenarios with limited ranges. The KITTI Stereo 2015 benchmark [34] has introduced 400 images of street scenes with ground truth depth acquired by a lidar system. In order to mitigate the sparsity of raw ground truth depth (4 % coverage), 7 laser scans are registered and accumulated, and moving objects are replaced with geometrically accurate 3D models leading to 19 % coverage. In order to obtain sufficient data for learning algorithms, Uhrig et. al [57] presented a method to generate denser depth maps (16 % coverage) by automatically accumulating 11 laser scans and removing outliers by stereo comparison. Note that, even with accumulation, the resulting depth maps are still not providing depth at the resolution of the image sensor [34]. Synthetic data sets such as the New Tsukuba Stereo Data set [33], Virtual KITTI [10] and Synthia [43] offer the possibility to create a theoretically unlimited amount of data with dense and accurate ground truth depth. While these data sets are extremely valuable for pretraining algorithms [42], they cannot replace real recordings for performance evaluation of real-world applications due to the synthetic-to-real domain gap [38]. In this work, we aim to close the gap between range-limited indoor and ground-truth-limited outdoor scenarios.

Robust Perception in Adverse Weather Robust environment perception is critical for enabling autonomous driv-

ing (without remote operators) over the world and in all environmental conditions. To this end, it is critical that self-driving cars should not stop working when being faced with unknown sensor distortions and situations that have not been in the training distribution. To characterize sensor distortions, previous methods have been focused on evaluating automotive sensors in challenging situations such as dust, smoke, rain, fog and snow [39]. For these evaluations, testing facilities such as Cerema [7] and Carissima [17] provide reproducible adverse weather situations with defined and adjustable severity. In particular, cameras suffer from reduced contrast because particles in the air cause scattering [3, 35]. Recently, the Robust Vision Challenge [13] promotes the development of robust algorithms benchmarked on a variety of datasets. However, the generalization and robustness to real adverse weather is significantly more difficult than dataset generalization because challenging weather conditions occur rarely and change quickly [59]. As a result, recent approaches evaluate robustness using synthetically extended datasets such as Foggy Cityscapes [44]. In this work, we depart from these synthetic datasets and instead propose a non-synthetic, but reproducible, benchmark for depth estimation in adverse weather conditions such as rain and fog under controlled conditions.

3. Sensor Setup and Calibration

To acquire realistic automotive sensor data, which serves as input for the depth evaluation methods assessed in this work, we equipped a research vehicle with a RGB stereo camera (Aptina AR0230, 1920x1024, 12bit) and a lidar (Velodyne HDL64-S3, 905 nm), see Figure 2. All sensors run in a robot operating system (ROS) environment and are time-synchronized by a pulse per second (PPS) signal provided by a proprietary inertial measurement unit (IMU). For the Velodyne lidar, both last and strongest return are recorded. Next, we describe the acquisition of the ground truth depth dataset that we use to evaluate the depth estimates obtained from the automotive sensor suite.

Ground Truth Acquisition We acquire ground truth depth for static scenarios using a Leica ScanStation P30 laser scanner (360°/290° FOV, 1550 nm, with up to 1M points per second, up to 8" angular accuracy, and 1.2 mm + 10 parts per million (ppm) range accuracy). To mitigate the effects of occlusions and to further increase the resolution, we accumulate multiple pointclouds at different overlapping positions. At least three white sphere lidar targets with a diameter of 145 mm at defined positions have to be detected for registration of the raw pointclouds. Each raw scan lasts about 5 min, which limits the proposed high-resolution acquisition approach to static scenarios. Using the known positions of the targets in every raw pointcloud, transforma-



Figure 2: Test vehicle equipped with a sensor suite including an Aptina AR0230 stereo camera and a Velodyne HDL64-S3 lidar system.

tions between these raw pointclouds are obtained by solving a linear least-squares problem. As a result, we obtain a dense pointcloud ($\approx 50M$ points) with a uniformly distributed *mean distance between neighboring points of only 3 mm*, corresponding to an angular resolution of 25". We use the middle of the rear axis, identified by measuring the positions of the wheel hubs, as the origin of the target point clouds. All sensors of the automotive sensor suite from Figure 2 have been calibrated with respect to this ground truth pointcloud, which we describe in the following.

Camera Calibration The intrinsic calibration of the stereo cameras is performed by detecting checkerboards with predefined field size [62]. We recorded these checkerboards at different distances and viewpoints in order to obtain the camera matrix and distortion coefficients. In order to register the very dense ground truth point clouds to the camera coordinate systems (extrinsic calibration), multiple black-white targets are placed at known 3D positions. By labeling the target positions in the images, an extrinsic calibration is obtained by solving the perspective-n-point problem using Levenberg-Marquardt non-linear least-squares optimization [29, 32].

Lidar Registration While the resolution of scans from the Leica laser scanner facilitates the detection of the white sphere targets, it is challenging to detect these targets at larger distances in the Velodyne laser scan. Therefore, these lidar targets cannot be used for registration of both Leica and Velodyne pointclouds. We use generalized iterative closest point (ICP) [50] for registration by minimizing the difference between two pointclouds, with the initial iterate shifted to the manually measured mounting position.

4. Adverse Weather Dataset

We model the typical automotive outdoor scenarios from [12]. Specifically, we setup the following four realistic



Figure 3: Representative automotive scenarios, see [12], covered in this evaluation benchmark.

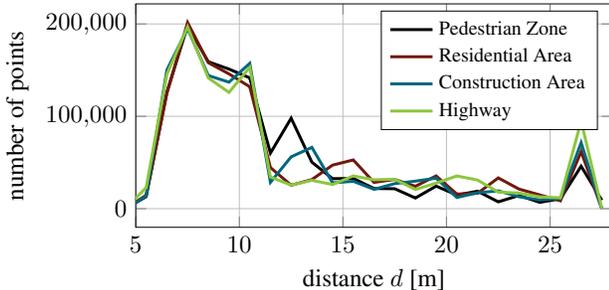


Figure 4: The depth distribution of the proposed scenarios is primarily driven by the camera frustum. The scenario itself changes the distribution only slightly.

scenarios: pedestrian zone, residential area, construction area and highway. Real-world examples of these scenarios are shown in Figure 3, while our setup scenes are shown in Figure 5. These scenarios are recorded under different controlled weather and illumination conditions in a weather chamber [7]. The visible back part of the weather chamber, see Figure 5, is constructed as greenhouse that is either transparent or covered by a black tarp and allows to achieve realistic daytime and night conditions. After acquiring reference measurements in clear conditions, the whole chamber is flooded with fog. The fog density is tracked by the meteorological visibility V defined by $V = -\ln(0.05)/\beta$, where β is the atmospheric attenuation. As the fog slowly dissipates, visibility increases and the recordings are stopped after reaching $V = 100$ m. In order to obtain a larger number of samples, three dissipation runs have been performed. For measurements in rain, two particular intensities at 15 and 55 mm/h/m² represent light and heavy rain. In total, this benchmark consists of 10 samples of each scenario (day/night) in clear, light rain, heavy rain and 17 visibility levels in fog (20-100 m in 5 m steps), resulting in 1600 samples in total.

5. Benchmark

In this section, we introduce all quantitative metrics and visualization methods used as evaluation methods in the remainder of this paper. The evaluation is performed on 2.5D depth images because this domain allows for an immediate

and intuitive comparison with the ground truth obtained by a depth camera or projected lidar points.

Metrics We adopt established metrics for benchmarking depth estimation algorithms. Specifically, we use the metrics from the KITTI benchmark [57], that is scale invariant logarithmic error (SIlog) [9], squared relative distance (SRD), absolute relative distance (ARD) and root-mean-squared error (RMSE). In addition, we evaluate mean-absolute error (MAE) and the threshold metric $\delta_i < 1.25^i$ for $i \in \{1, 2, 3\}$. Recently, Imran et al. [20] proposed root-mean-squared thresholded error (tRMSE) and mean-absolute thresholded error (tMAE) as variants of the established RMSE and MAE metrics which we also add to our evaluation framework. Moreover, as this work offers dense depth ground truth, we also assess depth map accuracy using dense image metrics, such as structural similarity (SSIM) and peak signal-to-noise ratio (PSNR). Since PSNR is based on absolute distances, we introduce a variant using relative depth relative peak signal-to-noise ratio (rPSNR), see supplemental document. We provide an in-depth description and formal definition of each metric in the supplemental document.

Binned Metrics The depth in a depth map is typically not uniformly distributed, as shown in Figure 4. By calculating the mean error of a depth image, errors at shorter distances contribute more to the mean than errors at larger distances. For a fair comparison of algorithms, we also provide binned evaluations where metrics are calculated in bins of approximately 2 m and the mean of the bins gives the final result. This ensures that every distance contributes equally to the evaluation metric.

Top-Down View As additional qualitative visualization, we provide top-down views generated by projecting 2.5D depth images into 3D. However, projecting all points to the ground plane does not provide any meaningful top-down view because in such a top-view, points from the ground, from the ceiling, and from objects in between cannot be distinguished. For improved visualization of the top-down view, we discretize the x-y plane into 10x10 cm grid cells and squeeze the height by counting the number of points in each grid cell [25, 61]. As the number of points in a cell is decreasing with distance (see Figure 4), we normalize the number of points according to their distance.

6. Evaluation

Evaluated Methods A large body of work on depth estimation and depth completion has emerged over the recent years. For brevity, we focus in this benchmark on one algorithm per algorithm category. Specifically, we compare *Monodepth* [14] as a representative method for monocular depth estimation, semi-global matching (SGM) [18] as traditional stereo and *PSMnet* [4] as deep stereo algorithms,

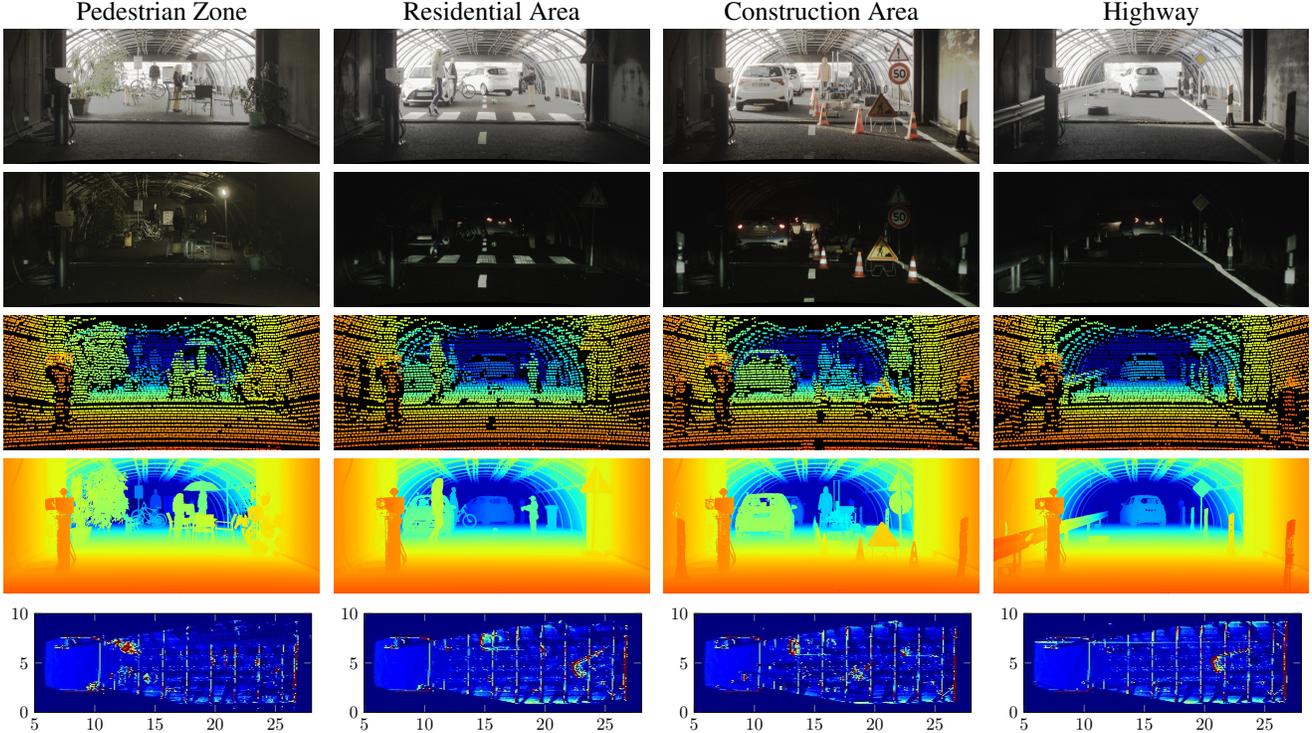


Figure 5: The proposed benchmark covers four different scenarios at day (first line) and night (second line). We provide state-of-the-art stereo camera and lidar sensors (third line) together with pixel-wise annotated ground truth (forth line). The last line shows a top-down view of each scenario. The large sensor on the left belongs to the visibility measurement system. The depth color coding is the same as in Figure 1.

and *Sparse2Dense* [31] as a depth completion method for lidar measurements using RGB image data. We set all algorithms up to estimate full-resolution depth maps, except for *Monodepth* [14] where we observed a substantial drop in performance and therefore resized the images to the native size the model was trained on. Additionally, we fine-tuned the model on the training split of [15] where the same sensor setup has been utilized. For *Sparse2Dense* [31], we trained the network with 8000 points on KITTI in order to apply it to our projected lidar points. Similar to the KITTI Depth Prediction benchmark [57], we interpolate the results of methods that provide less than 100% density with nearest neighbor interpolation. Additionally, we crop the images for 270 pixels from the top, 170 pixels from the left, 20 pixels from the right and 20 pixels from the bottom to avoid boundary artifacts, e.g. missing lidar points in the top part of the image.

Public Benchmark We will make all sensor data and the high-resolution ground truth data publicly available. All code for calculating the error metrics and for generating qualitative results (color-coded depth map, error map and top-down view) will be provided. The dataset is high-resolution, and enables fine-grained evaluation in controlled adverse weather conditions.

Clear Weather Evaluation First, we describe the evaluation results of all benchmarked algorithms in clear weather conditions. Figure 6 shows the 2.5D depth maps, the top-down views and error maps as described in Section 5 for the residential area scenario at daylight and night in clear conditions. Further qualitative results are shown in the supplemental document. While the stereo and depth completion approaches produce wrong depth estimates at the edges of objects, monocular depth estimation shows by far worst performance and does not generalize to our data. The quantitative metrics are averaged over all scenarios and can be found in Table 1. This evaluation shows that the monocular and stereo methods slightly decrease in performance at night, while lidar only and lidar depth completion perform stable. As described in Section 5, these metrics are often calculated over the whole image without considering the depth distribution, and, therefore, we also show a binned variant of the metrics. The performance slightly decreases as measurements from different distances are now equally weighted, which represents long-range driving scenarios. In Figure 7, MAE is plotted with respect to depth. Comparing traditional stereo [18] with deep stereo [4], Table 1 shows that deep stereo methods have difficulties to estimate the correct scale: Figure 6 shows that the distance to the back wall is estimated approximately 4 m closer than

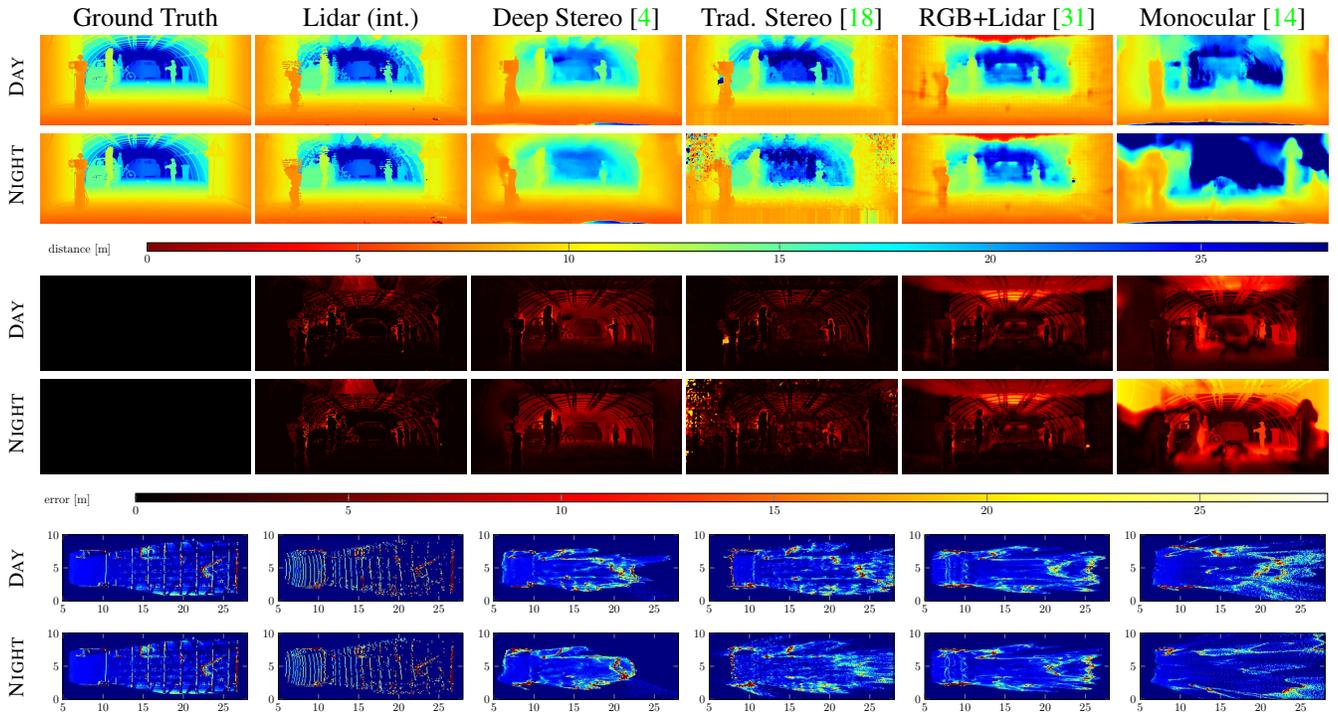


Figure 6: Qualitative evaluation results of the residential area scenario at day and night comparison in clear conditions. For each benchmarked algorithm, we include the color-coded depth map (top), error map (middle) and top-down view (bottom).

	Method	absolute				relative			scale-invariant	full-depth						
		RMSE [m]	tRMSE [m]	MAE [m]	tMAE [m]	logRMSE ↓	SRD ↓	ARD [%]	SIIlog 100 log (m) ↓	δ_1 [%] ↑	δ_2 [%] ↑	δ_3 [%] ↑	SSIM [0 - 1] ↑	PSNR [dB] ↓	rPSNR [dB] ↓	
DAY	not binned	Lidar (int.)	1.89	1.36	0.70	0.59	0.13	0.23	4.79	12.60	93.62	98.13	99.36	0.49	19.67	15.05
		Deep Stereo [4]	2.75	1.96	1.44	1.22	0.18	0.56	9.91	16.07	89.14	97.21	98.80	0.64	17.93	17.53
		Trad. Stereo [18]	1.90	1.40	0.96	0.86	0.14	0.27	8.12	13.32	90.74	98.44	99.50	0.47	20.82	16.43
		RGB+Lidar [31]	3.05	2.04	1.61	1.29	0.26	0.53	10.85	24.01	84.69	94.77	97.05	0.46	16.39	11.46
		Monocular [14]	5.01	2.67	2.93	2.01	0.33	2.78	27.12	29.12	73.73	88.87	95.10	0.47	14.45	18.96
	binned	Lidar (int.)	2.41	1.91	1.51	1.33	0.23	1.07	24.24	14.85	82.03	89.10	91.46	0.54	15.08	4.40
		Deep Stereo [4]	3.82	2.63	3.10	2.38	0.35	4.18	50.56	15.66	69.50	84.12	88.01	0.52	11.92	1.10
		Trad. Stereo [18]	2.52	2.01	1.83	1.60	0.26	1.92	36.31	11.81	76.11	89.21	91.16	0.41	14.90	4.32
		RGB+Lidar [31]	3.82	2.61	2.89	2.24	0.39	1.75	33.07	24.52	66.30	82.63	86.36	0.44	11.73	0.92
		Monocular [14]	7.49	3.42	6.33	3.07	0.54	31.62	143.44	26.30	56.63	75.92	82.59	0.36	8.15	-2.24
NIGHT	not binned	Lidar (int.)	1.90	1.36	0.70	0.58	0.13	0.23	4.83	12.77	93.57	98.03	99.32	0.49	19.09	14.89
		Deep Stereo [4]	2.94	2.12	1.65	1.40	0.19	0.61	11.59	16.73	84.38	97.25	98.95	0.61	17.40	18.69
		Trad. Stereo [18]	3.13	2.03	1.73	1.40	0.23	1.03	16.08	22.98	79.80	93.94	97.22	0.26	17.02	12.76
		RGB+Lidar [31]	3.03	2.03	1.60	1.29	0.25	0.52	10.77	23.38	84.81	94.71	97.14	0.47	16.46	11.57
		Monocular [14]	7.02	3.17	4.48	2.54	0.43	5.09	41.45	37.06	59.26	79.24	88.39	0.41	11.47	16.31
	binned	Lidar (int.)	2.42	1.91	1.50	1.31	0.24	1.07	24.30	15.13	81.95	88.92	91.36	0.54	14.74	4.07
		Deep Stereo [4]	4.11	2.78	3.39	2.50	0.37	5.97	58.18	16.14	64.59	84.89	88.47	0.53	10.93	0.15
		Trad. Stereo [18]	3.46	2.51	2.50	2.05	0.33	2.85	44.01	17.72	69.45	85.89	89.62	0.28	12.89	2.42
		RGB+Lidar [31]	3.80	2.61	2.89	2.24	0.38	1.75	33.09	23.88	66.31	82.34	86.45	0.44	11.77	0.95
		Monocular [14]	9.27	3.92	7.82	3.61	0.63	33.34	154.51	33.14	41.66	65.87	78.08	0.39	5.85	-4.42

Table 1: Quantitative comparison of all benchmarked algorithms based on a variety of 14 metrics averaged over all four scenarios. Although binned results show slightly worse performance, it provides a fairer comparison as depth is not equally distributed.

what it is. Nevertheless, for a scale-invariant metric such as SIIlog, deep stereo [4] performs better than traditional stereo [18], especially at night, matching the qualitative results in Figure 6.

Evaluation in Fog Next, we show how depth estimation algorithms degrade with increasing levels of fog. We have

extracted images at 17 different visibility levels and show the depth estimation performance with increasing visibility in Figure 9. Figure 8a shows the MAE for varying fog densities. While contrast in the camera stream drops with decreasing visibility, many cluttered points appear in the lidar point cloud due to severe back-scatter. As stereo is

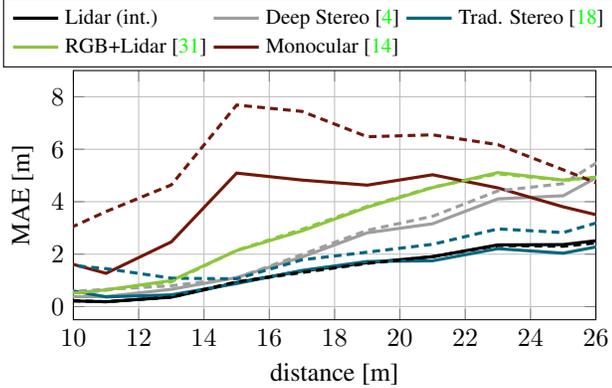
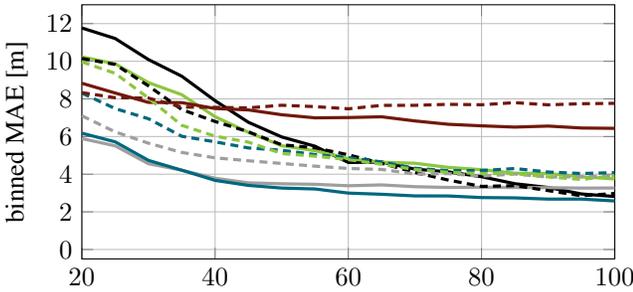
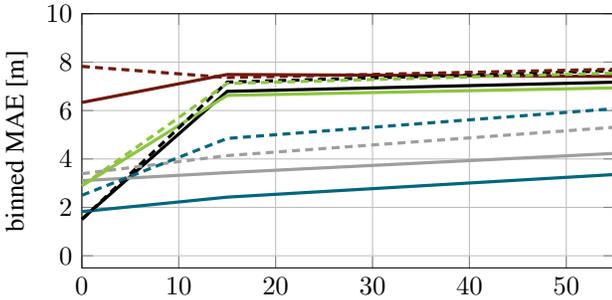


Figure 7: MAE calculated over depth bins of approximately 2 m averaged over all scenarios. Solid lines show daylight performance while dashed lines represent night conditions.



(a) Fog performance with respect to visibility V [m].



(b) Rain performance with respect to rainfall rate [mm/h/m²].

Figure 8: Binned MAE with respect to visibility and rainfall rate. Solid lines show daylight performance while dashed lines represent night conditions. For legend, see Figure 7.

based on finding correspondences, the performance of both stereo approaches decreases only slightly with lower visibility. In contrast, all methods based on lidar data collapse in the presence of fog because cluttered points cannot be distinguished from measurements originating from ballistic photons. The RGB input for lidar depth completion only slightly improves the depth recovery performance. This performance gap can be quantified by comparing clear conditions in daylight with visibility $V = 50$ m: between these two settings, deep stereo drops only by 45 cm in MAE, traditional stereo drops by 84 cm, while the MAE for lidar-

based depth completion reduces by a margin of 2.47 m and interpolated lidar even by 5.17 m. Comparing night and daylight conditions, Figure 8a indicates that camera-based depth estimation performs worse during night due to dark low-signal images, but lidar-based depth estimation shows slightly better performance during night. Additional metric visualizations in foggy conditions can be found in the supplemental document.

Evaluation in Rain We have recorded the scenes at two different rain intensity levels, that is 15 mm/h/m² and 55 mm/h/m². Figure 9 shows the qualitative impressions while Figure 8b visualizes the binned MAE corresponding to different intensities of rain. We note that light rain disturbs lidar already significantly due to path bending and back-scatter at the droplets while both stereo methods are relatively stable.

7. Conclusion

In this work, we provide a high-resolution ground truth depth dataset of angular resolution 25° for representative automotive scenarios in various adverse weather conditions. Using this evaluation data, we present a comprehensive comparison of state-of-the-art algorithms for stereo depth, depth from mono, lidar and sparse depth completion, in reproducible, finely adjusted adverse weather situations. The sensor data for these comparisons has been acquired with a representative automotive test vehicle.

We demonstrate that stereo approaches robustly generalize to the provided depth dataset, whereas monocular depth estimation is very sensitive to a change of scene type and capture conditions. We also evaluate depth estimation in the presence of fog and rain. We find that stereo-based approaches perform significantly more stable in adverse weather than monocular or recent lidar sensing and depth completion methods based on lidar measurements, which fail due to back-scatter in fog and rain.

All sensor data for the proposed benchmark (1600 samples) will be made publicly available, and researchers will be able to evaluate their methods using a benchmark server which keeps the high-resolution ground truth data hidden to the user. By acquiring data in a weather chamber, the proposed benchmark allows for reproducible evaluation in adverse weather – in contrast to existing driving datasets which lack data in adverse weather as current sensing systems fail in these scenarios. As such, this benchmark allows to jointly assess the robustness and resolution of existing depth estimation methods, which is essential for safety in autonomous driving.

References

- [1] S. Achar, J. R. Bartels, W. L. Whittaker, K. N. Kutulakos, and S. G. Narasimhan. Epipolar time-of-flight imaging. *ACM*

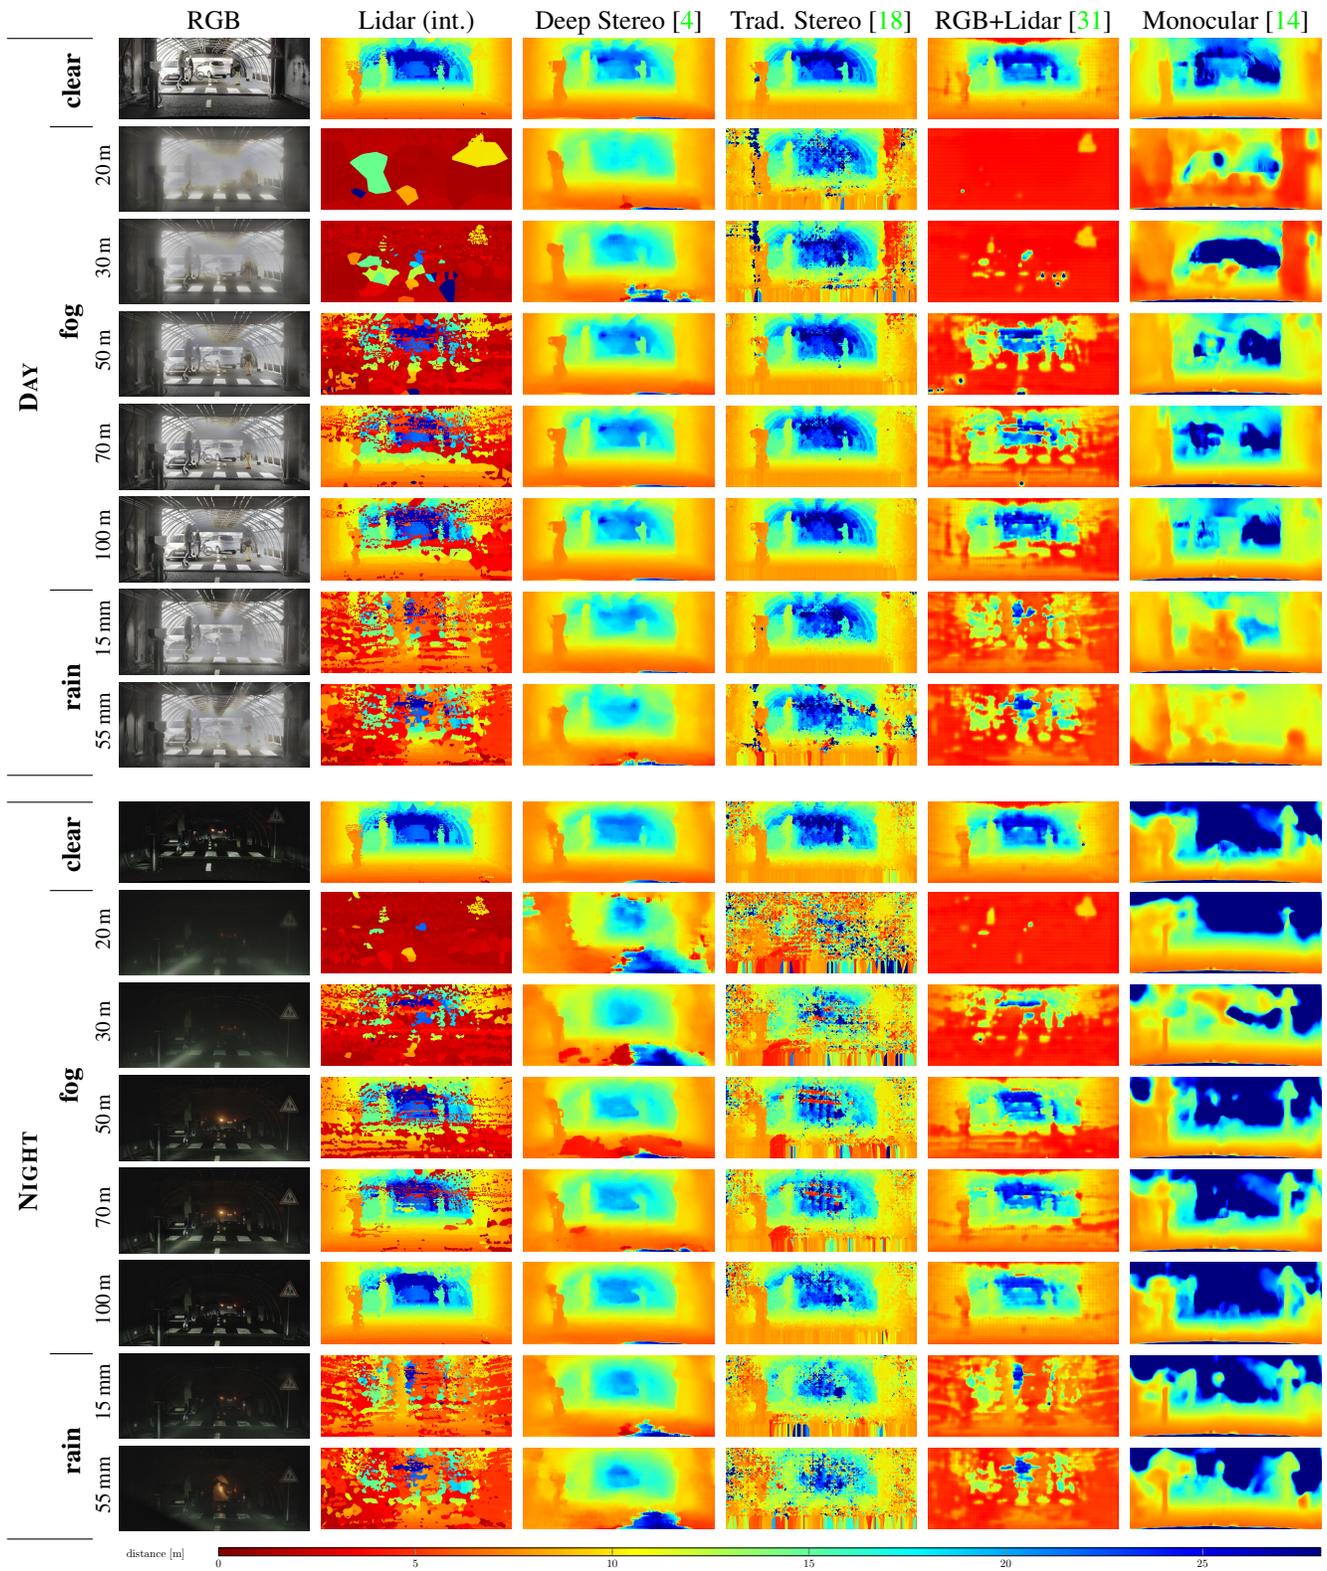


Figure 9: Qualitative results for all benchmarked depth estimation methods in clear conditions, foggy conditions of different visibility levels and rainy conditions. In addition to the resulting 2.5D depth maps, we include a RGB image as reference.

- Transactions on Graphics (ToG)*, 36(4):37, 2017. 1
- [2] F. Aleotti, F. Tosi, M. Poggi, and S. Mattoccia. Generative adversarial networks for unsupervised monocular depth prediction. In *Proceedings of the IEEE European Conference on Computer Vision*, 2018. 2
- [3] M. Bijelic, T. Gruber, and W. Ritter. Benchmarking image sensors under adverse weather conditions for autonomous driving. In *IEEE Intelligent Vehicle Symposium*, 2018. 3
- [4] J.-R. Chang and Y.-S. Chen. Pyramid stereo matching network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5410–5418, 2018. 1, 2, 4, 5, 6, 7, 8
- [5] R. Chen, F. Mahmood, A. Yuille, and N. J. Durr. Rethinking monocular depth estimation with adversarial training. *arXiv preprint arXiv:1808.07528*, 2018. 1, 2
- [6] Z. Chen, V. Badrinarayanan, G. Drozdov, and A. Rabinovich. Estimating depth from RGB and sparse sensing. *Proceedings of the IEEE European Conference on Computer Vision*, Sep 2018. 1, 2
- [7] M. Colomb, J. Dufour, M. Hirech, P. Lacôte, P. Morange, and J.-J. Boreux. Innovative artificial fog production device—a technical facility for research activities. 2004. 2, 3, 4
- [8] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner. Scannet: Richly-annotated 3D reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5828–5839, 2017. 2
- [9] D. Eigen, C. Puhrsch, and R. Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in Neural Information Processing Systems*, pages 2366–2374, 2014. 1, 2, 4
- [10] A. Gaidon, Q. Wang, Y. Cabon, and E. Vig. Virtual worlds as proxy for multi-object tracking analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [11] R. Garg, V. K. BG, G. Carneiro, and I. Reid. Unsupervised CNN for single view depth estimation: Geometry to the rescue. In *European Conference on Computer Vision*, pages 740–756. Springer, 2016. 2
- [12] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The KITTI dataset. *International Journal of Robotics Research*, 32(11):1231–1237, 2013. 1, 2, 3, 4
- [13] A. Geiger, M. Nießner, M. Pollefeys, C. Rother, D. Scharstein, A. A. Hassan, A. Dai, K. Honauer, J. Janai, T. Sattler, N. Schneider, J. Schönberger, T. Schöps, J. Uhrig, J. Wulff, and O. Zendel. Robust vision challenge. www.robustvision.net, 2018. 3
- [14] C. Godard, O. Mac Aodha, and G. J. Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 1, 2, 4, 5, 6, 7, 8
- [15] T. Gruber, F. Julca-Aguilar, M. Bijelic, W. Ritter, K. Dietmayer, and F. Heide. Gated2depth: Real-time dense lidar from gated images. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 5
- [16] M. Hansard, S. Lee, O. Choi, and R. P. Horaud. *Time-of-flight cameras: principles, methods and applications*. Springer Science & Business Media, 2012. 1
- [17] S. Hasirlioglu and A. Riener. Challenges in object detection under rainy weather conditions. In *First International Conference on Intelligent Transport Systems*, pages 53–65. Springer, 2018. 3
- [18] H. Hirschmuller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):328–341, Feb 2008. 1, 2, 4, 5, 6, 7, 8
- [19] X. Huang, X. Cheng, Q. Geng, B. Cao, D. Zhou, P. Wang, Y. Lin, and R. Yang. The apolloscape dataset for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 954–960, 2018. 2
- [20] S. Imran, Y. Long, X. Liu, and D. Morris. Depth coefficients for depth completion. *arXiv preprint arXiv:1903.05421*, 2019. 4
- [21] M. Jaritz, R. De Charette, E. Wirbel, X. Perrotton, and F. Nashashibi. Sparse and dense data with CNNs: Depth completion and semantic segmentation. In *International Conference on 3D Vision (3DV)*, pages 52–60, 2018. 1, 2
- [22] A. Kendall, H. Martirosyan, S. Dasgupta, P. Henry, R. Kennedy, A. Bachrach, and A. Bry. End-to-end learning of geometry and context for deep stereo regression. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017. 1, 2
- [23] J. J. Koenderink and A. J. Van Doorn. Affine structure from motion. *JOSA A*, 8(2):377–385, 1991. 1, 2
- [24] A. Kolb, E. Barth, R. Koch, and R. Larsen. Time-of-flight cameras in computer graphics. In *Computer Graphics Forum*, volume 29, pages 141–159. Wiley Online Library, 2010. 1
- [25] J. Ku, M. Mozifian, J. Lee, A. Harakeh, and S. L. Waslander. Joint 3D proposal generation and object detection from view aggregation. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 1–8. IEEE, 2018. 4
- [26] Y. Kuznetsov, J. Stückler, and B. Leibe. Semi-supervised deep learning for monocular depth map prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2215–2223, 2017. 2
- [27] I. Laina, C. Rupprecht, V. Belagiannis, F. Tombari, and N. Navab. Deeper depth prediction with fully convolutional residual networks. In *International Conference on 3D Vision (3DV)*, pages 239–248, 2016. 1, 2
- [28] R. Lange. 3D time-of-flight distance measurement with custom solid-state image sensors in CMOS/CCD-technology. 2000. 1
- [29] K. Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics*, 2(2):164–168, 1944. 3
- [30] Z. Li, T. Dekel, F. Cole, R. Tucker, N. Snavely, C. Liu, and W. T. Freeman. Learning the depths of moving people by watching frozen people. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Jun 2019. 1, 2
- [31] F. Ma and S. Karaman. Sparse-to-dense: Depth prediction from sparse depth samples and a single image. In *IEEE International Conference on Robotics and Automation*, pages 1–8, 2018. 1, 2, 5, 6, 7, 8

- [32] D. W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics*, 11(2):431–441, 1963. 3
- [33] S. Martull, M. Peris, and K. Fukui. Realistic cg stereo image dataset with ground truth disparity maps. In *ICPR workshop TrakMark2012*, volume 111, pages 117–118, 2012. 2
- [34] M. Menze and A. Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3061–3070, 2015. 2
- [35] J. P. Oakley and B. L. Satherley. Improving image quality in poor visibility conditions using a physical model for contrast degradation. *IEEE transactions on image processing*, 7(2):167–179, 1998. 3
- [36] M. O’Toole, F. Heide, L. Xiao, M. B. Hullin, W. Heidrich, and K. N. Kutulakos. Temporal frequency probing for 5D transient analysis of global light transport. *ACM Transactions on Graphics (ToG)*, 33(4):87, 2014. 1
- [37] M. O’Toole, R. Raskar, and K. N. Kutulakos. Primal-dual coding to probe light transport. *ACM Transactions on Graphics (ToG)*, 31(4):39–1, 2012. 1
- [38] X. Peng, B. Usman, K. Saito, N. Kaushik, J. Hoffman, and K. Saenko. Syn2real: A new benchmark for synthetic-to-real visual domain adaptation. *arXiv preprint arXiv:1806.09755*, 2018. 2
- [39] T. Peynot, J. Underwood, and S. Scheduling. Towards reliable perception for unmanned ground vehicles in challenging conditions. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 1170–1176. IEEE, 2009. 3
- [40] S. Pillai, R. Ambrus, and A. Gaidon. Superdepth: Self-supervised, super-resolved monocular depth estimation. *arXiv preprint arXiv:1810.01849*, 2018. 2
- [41] A. Pilzer, D. Xu, M. Puscas, E. Ricci, and N. Sebe. Unsupervised adversarial depth estimation using cycled generative networks. In *International Conference on 3D Vision (3DV)*, pages 587–595, 2018. 1, 2
- [42] S. R. Richter, V. Vineet, S. Roth, and V. Koltun. Playing for data: Ground truth from computer games. In *Proceedings of the IEEE European Conference on Computer Vision*, pages 102–118. Springer, 2016. 2
- [43] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. M. Lopez. The SYNTHIA dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Jun 2016. 2
- [44] C. Sakaridis, D. Dai, and L. Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, pages 1–20, 2018. 3
- [45] A. Saxena, S. H. Chung, and A. Y. Ng. Learning depth from single monocular images. In *Advances in neural information processing systems*, pages 1161–1168, 2006. 1, 2
- [46] A. Saxena, M. Sun, and A. Y. Ng. Make3D: Learning 3D scene structure from a single still image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):824–840, 2009. 2
- [47] D. Scharstein and R. Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision*, 47(1-3):7–42, 2002. 2
- [48] D. Scharstein and R. Szeliski. High-accuracy stereo depth maps using structured light. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 1. IEEE, 2003. 1
- [49] B. Schwarz. Lidar: Mapping the world in 3D. *Nature Photonics*, 4(7):429, 2010. 1
- [50] A. Segal, D. Haehnel, and S. Thrun. Generalized-ICP. In *Robotics: science and systems*, volume 2, page 435, 2009. 3
- [51] J. Shotton, B. Glocker, C. Zach, S. Izadi, A. Criminisi, and A. Fitzgibbon. Scene coordinate regression forests for camera relocalization in RGB-D images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2930–2937, 2013. 2
- [52] N. Silberman and R. Fergus. Indoor scene segmentation using a structured light sensor. In *2011 IEEE international conference on computer vision workshops (ICCV workshops)*, pages 601–608. IEEE, 2011. 2
- [53] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor segmentation and support inference from RGBD images. In *Proceedings of the IEEE European Conference on Computer Vision*, pages 746–760. Springer, 2012. 1, 2
- [54] S. Song, S. P. Lichtenberg, and J. Xiao. Sun RGB-D : A RGB-D scene understanding benchmark suite. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 567–576, 2015. 2
- [55] M. Subbarao and G. Surya. Depth from defocus: a spatial domain approach. *International Journal of Computer Vision*, 13(3):271–294, 1994. 1
- [56] P. H. Torr and A. Zisserman. Feature based methods for structure and motion estimation. In *International workshop on vision algorithms*, pages 278–294. Springer, 1999. 1, 2
- [57] J. Uhrig, N. Schneider, L. Schneider, U. Franke, T. Brox, and A. Geiger. Sparsity invariant CNNs. In *International Conference on 3D Vision (3DV)*, pages 11–20. IEEE, 2017. 1, 2, 4, 5
- [58] B. Ummenhofer, H. Zhou, J. Uhrig, N. Mayer, E. Ilg, A. Dosovitskiy, and T. Brox. DeMoN: Depth and motion network for learning monocular stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, July 2017. 1, 2
- [59] G. J. van Oldenborgh, P. Yiou, and R. Vautard. On the roles of circulation and aerosols in the decline of mist and dense fog in europe over the last 30 years. *Atmospheric Chemistry and Physics*, 10(10):4597, 2010. 3
- [60] H. Xu, Y. Gao, F. Yu, and T. Darrell. End-to-end learning of driving models from large-scale video datasets. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2174–2182, 2017. 2
- [61] B. Yang, W. Luo, and R. Urtasun. Pixor: Real-time 3D object detection from point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7652–7660, 2018. 4
- [62] Z. Zhang. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22, 2000. 3
- [63] Z. Zhang. Microsoft kinect sensor and its effect. *IEEE multimedia*, 19(2):4–10, 2012. 1, 2

- [64] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. [1](#), [2](#)