









A New Dataset for Monocular Depth Estimation Under Viewpoint Shifts

Aurel Pjetri^{1,2}, Stefano Caprasecca¹, Leonardo Taccari¹, Matteo Simoncini¹, Henrique Pineiro Monteagudo^{1,3}, Walter Wallace¹, Douglas Coimbra de Andrade^{4*}, Francesco Sambo¹, and Andrew David Bagdanov¹

¹ Verizon Connect Research, Florence, Italy {aurel.pjetri, stefano.caprasecca, leonardo.taccari, matteo.simoncini, francesco.sambo, wally.walter}@verizonconnect.com

² Department of Information Engineering, University of Florence, Florence, Italy andrew.bagdanov@unifi.it

³ University of Bologna, Bologna, Italy henrique.pineiro2@unibo.it

⁴ SENAI Institute of Innovation, Rio de Janeiro, Brazil douglas.coimbra@sc.senai.br

Abstract. Monocular depth estimation is a critical task for autonomous driving and many other computer vision applications. While significant progress has been made in this field, the effects of viewpoint shifts on depth estimation models remain largely underexplored. This paper introduces a novel dataset and evaluation methodology to quantify the impact of different camera positions and orientations on monocular depth estimation performance. We propose a ground truth strategy based on homography estimation and object detection, eliminating the need for expensive lidar sensors. We collect a diverse dataset of road scenes from multiple viewpoints and use it to assess the robustness of a modern depth estimation model to geometric shifts. After assessing the validity of our strategy on a public dataset, we provide valuable insights into the limitations of current models and highlight the importance of considering viewpoint variations in real-world applications.

Keywords: Dataset · Domain Shift · Monocular Depth Estimation

1 Introduction

Recent years have seen significant advancements in the machine learning community, with transformers in particular achieving state-of-the-art results on multiple tasks [6, 22]. Autonomous driving (AD) has been receiving more and more attention from the public with different forms of it being regulated and approved for public implementations, from fully autonomous robotaxis [3] to consumer vehicles with different levels of autonomy [1, 2]. Perception, and computer vision in particular, is the most important building block for these technologies. A fundamental task for the perception of the surroundings is depth estimation. It is

* Work done while at Verizon Connect Research, Florence, Italy.

well known that such tasks are susceptible to out-of-distribution samples, from adverse weather and light conditions to image corruptions [16, 19, 24].

Another distribution shift that is common in this context is viewpoint change, either caused by vehicles with different sizes or by variations in the camera installation. A few works have recently started to shed light on the effect of such shifts on related tasks like birds-eye-view semantic segmentation and object detection [15, 23]. However, quantification of such an effect on monocular depth estimators remains heavily underexplored, with only [21] giving some insights on the behaviours of these models. This could be mainly due to the lack of annotated data and the cost and technical difficulties associated with it. Traditionally, this kind of task is performed using measurements from expensive lidar sensors [5, 8, 26]. These sensors are not ubiquitous and require multiple layers of post-processing, such as outlier detection and removal, data alignment, registration with other sources, and other computationally expensive steps. A possible solution lies in the domain of synthetic data, which introduces a well-known sim-to-real domain shift [12].

Our proposal is to remove the dependency on expensive sensors to evaluate depth estimation techniques by exploiting basic geometric elements such as *homographies*. We estimate a ground truth (GT) homography from an initial calibration session that maps points from the image plane into the road plane. We use the homography to compute distances of road objects that serve as GT for our evaluation. As a matter of fact, our proposal is to use *object distance estimation* as a proxy task for traditional depth evaluation. We are aware that this approach is accurate only if the geometry of the road plane is in line with the one measured during calibration, otherwise introducing an intrinsic error. However, we empirically show that such an error is considerably limited on planar geographies such as the one from the KITTI dataset [8]. Furthermore, experiments on KITTI show that our proxy task provides evaluations that are very close to the ones observed using lidar GT.

Our strategy enabled us to collect a large dataset of video sequences, recorded from different known and calibrated camera settings. We use such data to highlight the impact of the different viewpoint shifts on a modern depth estimator and give insights on their effects on the scale perceived by the model. The dataset can be accessed at kaggle.com/datasets/aurelpjetri/viewpointdepth.

To summarize, the main contributions of this work are:

- we propose a new ground truth strategy based on classical computer vision techniques, which can be used to replace expensive sensors for object distance estimation;
- we discuss in-depth the limitations of the proposed approach and empirically show its effectiveness on a public annotated dataset;
- we release a new dataset of road scenes collected from multiple viewpoints;
- we quantify the effects of different viewpoints on the performance of a modern monocular depth estimation model.

The remainder of this paper is structured as follows: Sec. 2 presents the most relevant works in this context, Sec. 3 outlines the proposed ground truth and

evaluation strategy, Sec. 4 presents the details of the collected dataset, Sec. 5 describes the different experiments conducted and Sec. 6 draws final conclusions.

2 Related Work

In this section we review work from the recent literature most related to our contributions.

2.1 Road Scene Depth Datasets

Many datasets have been proposed to facilitate research on 3D perception for autonomous driving or, in general, for road scene analysis. The KITTI dataset was a pioneering work in this context, in which a test vehicle collected driving data from multiple cameras and sensors, including a laser scanner for distance ground truth [8]. The authors collected different scenes driving in an urban environment in Germany. Inspired by KITTI, other datasets have been released by expanding the variety of geographical regions and scenes collected. Examples are Nuscenes [5] and Waymo Open [20], which are larger than KITTI and cover a wider variety of cities and scenes. Other datasets like WoodScape [26] focus on fisheye cameras, which are not only front-facing, but also cover lateral and rear, and are mounted on many modern cars.

All public datasets, in general, are mainly focused on capturing a variety of scenes, weather, lighting conditions, sensors or image corruptions [16]. None of them, however, provides a broad variety of viewpoints, as the vehicle type and camera position are usually fixed across the whole dataset.

2.2 Self-supervised Monocular Depth Estimation

Depth estimation is a fundamental geometric task for 3D scene understanding. Self-supervised monocular depth estimation is a particular variant of this task that enables training networks without the use of explicit depth labels. Although it is an ill-posed problem, since multiple 3D scenes can generate the same 2D image, some deep learning techniques have shown good results. One of the first proposed approaches used the pretext task of reconstructing frames in a video sequence to force the network to learn a good depth estimation of the scene [29]. The depth network is trained alongside a pose network for the new frame to be rendered. Given a source frame at time t and a target frame at $t + 1$, each pixel p_{t+1} can be back-projected to the source view using depth D_t and pose $P_{t+1 \rightarrow t}$ predicted by the networks:

$$p_t \sim K P_{t+1 \rightarrow t} K^{-1} p_{t+1}. \quad (1)$$

In this way bilinear sampling can be used to synthesise \hat{p}_{t+1} [13]. A direct comparison of each pixel i within each frame t for training loss can then be used [29]:

$$\mathcal{L} = \sum_t \sum_i |p_t^i - \hat{p}_t^i|. \quad (2)$$

Inspired by this work, [10] introduced multiple improvements, such as the use of a photometric error in the loss function combining L1 and SSIM [25]. The results were impressive, not only because they made state-of-the-art, but also given how close they were to stereo-based methods, proving the effectiveness of this paradigm.

Over the years many works have built on top of this paradigm, implementing more effective and modern architectures. One of the first implementations of vision transformers [6] to tackle this problem was proposed in [28]. In particular, inspired by the MPViT architecture [17], the authors propose a Conv-stem block for the network encoder, in which they combine convolutional blocks with transformers. They observe that this enables modeling both local and global information in the frame, achieving state-of-the-art at the time of its release.

2.3 Viewpoint Shift

Depth estimation under geometric shifts remains underexplored, despite being a well-known problem. It has been tackled in the context of supervised depth estimation by introducing a geometric embedding of the inputs and a variational latent representation [11]. Additionally, viewpoint shifts caused by different camera installations for autonomous driving have been studied [15]. In this case, the task was bird’s eye view semantic segmentation and the dataset used was synthetic. Moreover, a viewpoint generalization technique for object detection, based on learning homography transformations of the input images to account for different geometric shifts, has been proposed [23].

Finally, for the specific task of monocular depth estimation, [21] conducted multiple experiments with the aim of understanding which depth cues are actually exploited by the models. The authors observed that vertical position of the object is the main cue used by the models: the lower side of the vehicle and the point where it touches the ground seem to be the main cues exploited by models to estimate its distance. Their experiments also suggest limitations of such models to account for camera orientation changes simulated by image central crops and rotations.

3 Ground Truth and Evaluation Strategy

We propose to use a homography in combination with an object detector as a ground truth source for measuring monocular depth estimation error on road scenes. Depth estimation models are traditionally tested by directly comparing per-pixel predictions with ground truth from expensive sensors, such as lidar, and computing metrics such as absolute relative error:

$$\text{abs-rel} = 100 \cdot \frac{1}{N} \sum_{i=0}^N \frac{|x_{gt}^i - x_{pred}^i|}{x_{gt}^i}. \quad (3)$$

We hypothesize that comparable results can be achieved by evaluating the models on object distance estimation, measuring the error only on road objects




Fig. 1: (a) Manually labeled points of a parking lot pattern (source points for the homography). (b) Target plane with metric distances (target points for the homography).

instead of on the whole frame. Moreover, we observe that a homography that maps ground points in the image to a plane with a metric reference system can be used as distance ground truth source for points on the ground (Fig. 1). Such projection, used in combination with an object detector, can associate a metric distance to objects. As a matter of fact, previous studies have already suggested that monocular depth models use the lower part of the object to determine their distance [21], which further supports our strategy. This ultimately enables experiments and research in conditions in which expensive sensors are not available. We divide our method in three steps: calibration, object detection, and evaluation.

3.1 Calibration

The calibration phase is the only manual step of the method. We select reference ground points and measure their distance to the camera to construct the target plane with a metric reference system centered in the camera (see Fig. 1), so that for any point $X = [x, y]^T$ its distance to the camera can be computed by simply $\sqrt{x^2 + y^2 + h^2}$, where h is the height of the camera. This allows us to estimate a homography mapping points in the image to points on the target plane, and therefore to associate a metric distance to any point in the image, provided that it lies on the ground. Given a point on the image $x = [u, v]^T$, we project it to X in the target plane:

$$[X, Y, W]^T = H \cdot [u, v, 1]^T \Rightarrow X = [X/W, Y/W]^T = [x, y]^T. \quad (4)$$

3.2 Object Detection

We use an object detector to detect the relevant road objects in the image: car, truck, bus, motorcycle, bicycle and person. Given a bounding box with coordinates (x_1, y_1) and (x_2, y_2) , we assume that all objects lie on the ground, and therefore we define their distance as that of the center of the lower side of




Fig. 2: (a) Bounding box of a vehicle with the point \hat{X} used for GT distance. (b) Depth prediction with bounding box. The original box was resized to $\alpha = 75\%$ (dashed line) to extract the percentile β of the prediction.

the box $\hat{X} = ((x_2 - x_1)/2, y_2)$ (Fig. 2(a)). By projecting this point using Eq. (4), we are able to associate a metric distance to each object and use it as ground truth for the model evaluation in the next step.

3.3 Evaluation

To measure the performance of a depth estimation model, we only consider road objects. Therefore, we again use the bounding box information from the object detector and use the depth estimation model to infer the dense distance of the points within the boxes. We then select a percentile β from the estimated distances and use that as the inferred object distance. Moreover, to account for inaccuracies we resize bounding boxes to a certain fraction α of the original dimension before collecting inferred depths, as shown in Fig. 2(b). α and β are treated as hyperparameters that we select in the experimental phase (see Sec. 5.1).

4 Dataset

The aim of this dataset is to measure the effects of different viewpoint shifts on monocular depth estimation models. We collected the data using a 2023 RAM Promaster 2500 cargo van located in Alpharetta (Georgia), United States. Between March 2023 and February 2024, we installed two identical dashcams in the windshield of the vehicle. The first camera, to which we refer as *base*, was fixed in an ideal position for a dashcam, *i.e.*, close to the rear-view mirror at a height of 177.8 cm with an orientation of 0° roll, 0° yaw and -4° pitch (later moved to 6° pitch as shown in Tab. 1). Example frames recorded from these positions are reported in Fig. 4.

The second camera, which we identify as *shifted*, was instead used to simulate different viewpoint shifts and was moved on a weekly basis. Table 1 and Fig. 3 show the different combinations of positions and orientations of the cameras during the collection period. For the camera position in the windshield (X, Y ,




Fig. 3: (a) Reference system on the windshield centered on the *base* camera, where the Z axis orientation is in the direction of movement of the vehicle. The red dots indicate the different positions of the *shifted* camera. (b) Axes and rotation reference systems following the right-hand rule.

Z) we used the reference system illustrated in Fig. 3(a) centered in the position of the *base* camera.

It is worth noting that the yaw angle, present in the camera when installed on the left side (lines 6–9 in Tab. 1), was implicitly generated by the curvature of the windshield itself. Roll angles were computed by manually measuring the camera inclination on the windshield. On the other hand, we measured pitch and yaw angles for each viewpoint by manually labeling the horizontal vanishing point from the straight lines on the road (v_u, v_v) , and computing the angle given the focal length (f_x, f_y) of the camera and its principal point (c_u, c_v) :

$$\text{pitch} = \frac{180}{\pi} \arctan \left(\frac{c_v - v_v}{f_y} \right) \quad (5)$$

$$\text{yaw} = \frac{180}{\pi} \arctan \left(\frac{c_u - v_u}{f_x} \right) \quad (6)$$

A total of 10 different viewpoints were collected. For each viewpoint we release 3 videos of 10 minutes each from different scenes. Given the long time range in which the videos are distributed, we manually selected samples to represent a similar distribution of suburban areas in non-adverse weather and at day-time as shown in Figs. 4 and 5. Dashcams recorded at an average of 30 frames per second in 720p MP4 format and with a 150° diagonal field of view. We down-sampled the data at 10 fps and extracted the frames undistorted in JPEG format, with a quality factor of 92 and 4:2:0 chroma sub-sampling. This resulted in a total of 360K frames.


Note that although the two cameras are simultaneously recording the same scene, no synchronization mechanism was implemented. This can result in the two recordings being shifted by a variable number of frames. Finally, to foster further research, we also release GPS points and accelerometer data at 1 and 100 Hz, respectively.

Table 1: Combinations of different positions in the windshield and orientation angles (in degrees). X, Y, and Z represent the position of the camera in the windshield (in centimeters). Objects represents the number of objects detected in the associated frames. For each row we also present the pitch the *base* camera had during the recording of that viewpoint and he associated objects.

ID	<i>Shifted</i> Dashcam							<i>Base</i> Dashcam	
	X	Y	Z	Pitch	Yaw	Roll	Objects	Pitch	Objects
0	0	30	0	-4	0.0	0	5291	-4	3026
1	5	0	0	-10	0.0	0	4325	6	4146
2	5	0	0	16	0.0	0	5004	6	5458
3	5	0	0	-4	2.5	0	2906	6	2824
4	5	0	0	-4	5.0	0	2281	6	2219
5	5	0	0	16	5.0	0	3068	6	4207
6	-65	-5	-10	5	0.0	-5	4425	6	3916
7	-65	-5	-10	16	0.0	-5	3537	6	3759
8	-65	30	-10	5	0.0	-5	3465	6	3278
9	-65	30	-10	16	0.0	-5	3712	6	4413



(a)



(b)

Fig. 4: (a) Example of *base* camera with -4° of pitch. (b) Example of *base* camera with 6° of pitch.

4.1 Object Detection

We used a Yolov5 object detector [14], trained on a combination of BDD100K [27] and a private dataset collected with same camera model of this work. We do not expect the object detector to be affected by the viewpoint shifts as both datasets used for training have non-fixed viewpoints. Moreover, from a qualitative analysis, we observe that the combinations in Tab. 1 do not induce a significant change in appearance of the objects. This is also confirmed by smaller-scale experiments that showed no evident change in mean average precision between the different viewpoints.

To exclude problematic cases, we filter out all the boxes with a confidence level lower than 0.5 and the lower side above the vanishing point of the camera. Furthermore, we remove occluded objects by detecting overlapping boxes and only keeping the lowest ones in the image. Finally, we filter out small objects

with specific area thresholds for each class: 3000 pixels² for all vehicles, 1000 pixels² for person and bicycle, and 1500 pixels² for motorcycle. Both raw and filtered detections are released with the dataset.

4.2 Homography

As mentioned in Sec. 3, in order to be able to associate a ground truth distance to road objects, we estimated homographies mapping image ground points to a target plane with metric reference system (Fig. 1) that we release alongside the dataset.

Base Camera Homography We exploited the regular pattern of an empty parking lot in order to have several reference points and minimize manual measurements. In Fig. 1 we show the ground points manually labeled on the image and the corresponding points on a plane with metric coordinates. The homography was estimated using all the provided points with a simple least-squares scheme from the OpenCV [4] Python library.

Shifted Camera Homography We selected a set of corresponding points between the cameras after manually synchronizing a portion of the two videos. Using the *base* camera homography, we projected the points from the base camera to the target plane and used the corresponding points from the *shifted* camera to estimate the needed homography. This operation was repeated for each position.

5 Experimental Results

In this section we report the results of our experiments. In particular, we validate our method as ground truth source by comparing it with lidar readings on the KITTI dataset in Sec. 5.1. Then, in Sec. 5.2, we use our ground truth method to measure the effects of the different viewpoints on a self-supervised monocular depth estimator.

In particular, for all our experiments we use the MonoViT depth estimation model [28], which is a modern architecture which yields state-of-the-art performance on the RoboDepth challenge on robustness against adverse weather, light, and sensor corruption [16]. Despite being a challenge on simulated weather conditions and image degradation, we thought it would be a promising candidate for our study.

5.1 Homography Evaluation on KITTI

The goal of this experiment is to assess the margin of error of our geometry-based strategy by comparing it with lidar ground truth. We take the MonoViT model trained on the KITTI dataset [8] for self-supervised monocular depth estimation

and measure its absolute relative error (Eq. (3)) using two ground truth sources: laser scans and our strategy.

As a first step we estimate the homography by exploiting the provided camera projection matrix. We select a set of reference points in the target plane with metric reference system, project them to the image plane and estimate the homography with the same method used in our data. We also exploit the provided object labels [9] and filter out occluded and partial objects.

We perform inference with MonoViT on the test set from the Eigen split [7]. Monocular depth estimators give predictions with an unknown scale factor and need to be manually scaled to metric depth using ground truth. Traditionally, a median scaling is performed at an image level, *i.e.* for each image j predictions are multiplied by $s_j = \text{median}(X_{gt}^j)/\text{median}(X_{pred}^j)$. However, this can be unfair towards other stereo-based or supervised methods that do not use ground truth information at test time, as highlighted by [10]. Therefore, in these experiments we compute a single scale multiplier by computing medians across the whole test set. We replicate the same in our strategy by computing medians of homography GTs and prediction percentiles extracted from the boxes.

To find the best parameters for bounding box resize and percentile of the prediction inside the box (α and β respectively), we performed a grid search over the two parameters and found that the best performance is given by $\alpha = 0.75$ and $\beta = 75$. With these parameters we measured an abs-rel of 16.8 with our homography and 13.6 with laser, giving a difference of 3.22 points.

Moreover, for the task of object distance estimation, we compute Spearman rank correlation scores between laser readings and the homography. To associate a laser distance to an object, we take the sensor reading of the middle point of the lower side of the box. What we observe is that the two sources of ground truth have a high correlation with a Spearman score of 0.97.

Road Geometry These results also suggest that the intrinsic error introduced by the use of a homography, with its strong assumptions of flat ground surface, is limited. To further verify this, we also computed statistics about road inclination in KITTI’s test set exploiting the provided GPS signals. As a matter of fact, using altitude and speed, we are able to compute the road’s slope in the traveled segment. In particular we sample GPS points at 1Hz and keep only those where the horizontal displacement resulted in more than 1m. We observe that the average of the absolute values of the angles is 1.2° , median is 0.4° and 99th percentile is 21.1° . The same statistics on our dataset are: 1.8° , 0° , 19.1° respectively. Moreover, the homography’s assumption is broken only in the case where there is a change in the road geometry, therefore if we compute the number of times the GPS altitude changes of more than 1m (the granularity of our GPS signal), we observe that in KITTI it happens in 1.46% of the points, which is similar to the 1.99% in our dataset. These measurements indicate that the road geometry of the two datasets is similar and encourage us to use this strategy on our dataset expecting an intrinsic error analogous to the one on KITTI.

5.2 Viewpoint Shift Effects

Our aim is to find what are the most downgrading shifts among the different positions that we collected. We build a training set from additional 3 videos of the *base* camera with -4° pitch at 10 fps and exploit the rest of the dataset as test set at 1 fps to measure the effects of different viewpoints on MonoViT. We run our experiments using MonoViT pre-trained on KITTI with a resolution of 1024x320. To fit our images in the network, we centrally crop our frames from 720 to 704, ending up with frames of 1280x704 pixels.

We first ran a zero-shot experiment on both cameras and all the positions. Then we fine-tuned the model on the training set using an Nvidia L4 GPU with 24GB, with a batch size of 2 for 21 epochs. For the optimization we used AdamW [18] with an initial learning rate of 10^{-9} for the depth encoder and 10^{-8} for all the other parts of the networks, and exponential decay.

As mentioned in Sec. 5.1, predictions need to be scaled to metric distances for evaluation. However, in order to highlight the effects of the shifts on the scale perceived by the model, we choose to use the same scaling factor for all positions. This factor is computed exploiting all the objects O in the training set T , both in the zero-shot and fine-tuning experiments:

$$s = \frac{\text{median}(\text{gt}(O^k))}{\text{median}(\text{pred}(O^k))}, \forall O^k \in T. \quad (7)$$

Table 2 gives the absolute relative errors of the model in the different positions, both in zero-shot and fine-tuned inference. To isolate the effect of the viewpoint, for each position, we test on footage from the same scene recorded by the two cameras simultaneously, and compute the difference in performance between *shifted* (S) and *base* (B) camera. The rows with the highest difference between the S and B cameras correspond to the installation positions with the strongest performance degradation.

From the zero-shot results, we observe that the two most affected positions are 7 and 5 (in decreasing order). The combination of pitch and yaw has the most degrading effect, leading to more than doubling the error. Pitch and roll have also the effect of doubling the error of the model. Fine-tuning reduces the error in general, however the degrading effects of the 2 worst positions remain almost unchanged. In general we observe that pure rotations or translations don't affect the model as much as a combination of them, with some more impactful than others.

Scale Effect Given the single scale factor used for all positions, we believe that the main effect of viewpoint shifts is the distortion of the scale perceived by the model, which could not have been detected if per-image or per-position scaling was used. In order to better investigate this phenomenon, we compute scaling factors as in Eq. (7) on the different positions separately. Tab. 3 reports the computed scales on the zero-shot experiment, along with the error degradation from Tab. 2. A correlation between performance degradation and perceived scale

Table 2: Absolute relative error of MonoViT in different camera positions for zero-shot and fine-tuned networks. We report errors on the *shifted* (S) and *base* (B) cameras corresponding to the same scene. [†] *Base* camera with pitch -4.0° , all other rows are with *base* camera at pitch 6° . The top 2 largest degradations are in **bold**, the second 2 are underlined.

ID	<i>Shifted</i> Dashcam Position						Zero-shot			Fine-tuned		
	X	Y	Z	Pitch	Roll	Yaw	B	S	S-B	B	S	S-B
0	0	30	0	-4	0	0	18.3 [†]	19.4	1.1	16.9 [†]	17.3	0.4
1	5	0	0	-10	0	0	19.4	29.0	<u>9.6</u>	18.0	22.9	<u>5.0</u>
2	5	0	0	16	0	0	17.2	26.1	<u>8.8</u>	17.0	21.1	<u>4.1</u>
3	5	0	0	-4	2.5	0	18.5	21.4	2.9	15.1	16.4	1.2
4	5	0	0	-4	5	0	20.9	21.3	0.4	17.2	17.7	0.5
5	5	0	0	16	5	0	20.5	40.4	19.9	19.7	35.8	16.2
6	-65	-5	-10	5	0	-5	19.6	24.0	4.4	16.5	18.1	1.6
7	-65	-5	-10	16	0	-5	19.0	47.5	28.5	17.7	44.1	26.4
8	-65	30	-10	5	0	-5	18.5	23.6	5.1	17.3	21.2	4.0
9	-65	30	-10	16	0	-5	19.2	25.9	6.6	15.1	17.6	2.5

can be seen. In particular, we observe that higher pitch angles correspond to smaller scales perceived by the model (and therefore a bigger scaling factor necessary to align to GT). In the same way, lower pitch angles generate bigger scales perceived. This observation seems to confirm that indeed depth models use vertical position as depth cue for objects as stated by [21]. In general this suggests that adaptation methods tackling this problem should focus on automatic scaling of models at inference time.

6 Conclusion

In this paper, we addressed the underexplored problem of viewpoint shifts in monocular depth estimation for camera systems. We introduced a novel dataset collected using two dashcams in various positions and orientations. Our method leverages homography estimation and object detection as a cost-effective alternative to expensive lidar sensors for ground truth depth generation. We validated our approach on the KITTI dataset, demonstrating a high correlation (Spearman score of 0.97) with lidar measurements and measuring an absolute relative error comparable to the one using lidar. We showed how the intuitive limitations relative to the road geometry of the proposed approach are not impactful in practice on a public dataset like KITTI, and, by extension being the two geometries similar, to our dataset.

Using the MonoViT model, we quantified the effects of viewpoint shifts on depth estimation performance. Our experiments revealed that certain camera positions significantly degrade accuracy. In particular combinations of pitch and

Table 3: Scale perceived by the network in different positions. Scale B and Scale S are the scales perceived in the *base* and *shifted* cameras. Scale S-B is computed as Scale S - Scale B to highlight the effect of the viewpoint shift. The difference in error (Error S-B) is also reported from Tab. 2 for reference. The top 2 largest degradations are in bold, the second 2 are underlined.

ID	<i>Shifted</i> Dashcam Position						Zero-shot			
	X	Y	Z	Pitch	Roll	Yaw	Error S-B	Scale B	Scale S	Scale S-B
0	0	30	0	-4	0.0	0	1.1	3116	3064	-52
1	5	0	0	-10	0.0	0	<u>9.6</u>	3359	2669	<u>-689</u>
2	5	0	0	16	0.0	0	<u>8.8</u>	3167	4313	<u>1145</u>
3	5	0	0	-4	2.5	0	2.9	3310	3086	-224
4	5	0	0	-4	5.0	0	0.4	3580	3327	-252
5	5	0	0	16	5.0	0	19.9	3330	5254	1923
6	-65	-5	-10	5	0.0	-5	4.4	3315	3034	-280
7	-65	-5	-10	16	0.0	-5	28.5	3310	6133	2822
8	-65	30	-10	5	0.0	-5	5.1	3158	2706	-451
9	-65	30	-10	16	0.0	-5	6.6	3397	3970	572

roll and pitch and yaw have the most degrading effects. As a matter of fact, the most impactful shift induces an increased error of above 26 points in both zero-shot and fine tuning, which is more than double from the *base* position. Moreover, a study on inference scale of the model highlighted that camera pitch directly affects perceived scale and that, in general, performance degradation is correlated to scale distortion, suggesting that indeed automatic scaling of depth models at inference time could help mitigate the problem and thus be a promising research direction for generalization.

These results highlight the vulnerability of current depth estimation models to geometric shifts, particularly combinations of pitch and yaw angles. Fine tuning on the *base* camera data improved overall performance but did not fully mitigate the effects of viewpoint changes.

Our work provides insight into the limitations of current self-supervised monocular depth estimation models and emphasizes the importance of considering viewpoint variations in real-world applications. Future research should focus on developing more robust models that can generalize across different viewpoints, potentially by incorporating our dataset and evaluation methodology into training and testing pipelines. Additionally, expanding the dataset to include different vehicle types, camera models and more viewpoints would further enhance its utility.



Fig. 5: Example frames from the *shifted* camera positions following the ID in Tab. 1.

References

1. Automated driving revolution: Mercedes-Benz announces U.S. availability of DRIVE PILOT – the world’s first certified SAE Level 3 system for the U.S. market. <https://media.mbusa.com/releases/automated-driving-revolution-mercedes-benz-announces-us-availability-of-drive-pilot-the-worlds-first-certified-sae-level-3-system-for-the-us-market>
2. Tesla Autopilot, <https://www.tesla.com/support/autopilot>
3. Waymo one, <https://waymo.com/waymo-one/>
4. Bradski, G.: The OpenCV Library. Dr. Dobb’s Journal of Software Tools (2000)
5. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
6. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2021)
7. Eigen, D., Puhrsch, C., Fergus, R.: Depth Map Prediction from a Single Image using a Multi-Scale Deep Network. In: Advances in Neural Information Processing Systems. vol. 27. Curran Associates, Inc. (2014)
8. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The KITTI dataset. The International Journal of Robotics Research **32**(11), 1231–1237 (Sep 2013). <https://doi.org/10.1177/0278364913491297>
9. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2012)
10. Godard, C., Aodha, O.M., Firman, M., Brostow, G.: Digging Into Self-Supervised Monocular Depth Estimation. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 3827–3837. IEEE, Seoul, Korea (South) (Oct 2019). <https://doi.org/10.1109/ICCV.2019.00393>
11. Guizilini, V., Vasiljevic, I., Chen, D., Ambruş, R., Gaidon, A.: Towards zero-shot scale-aware monocular depth estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9233–9243 (October 2023)
12. Hu, C., Hudson, S., Ethier, M., Al-Sharman, M., Rayside, D., Melek, W.: Sim-to-real domain adaptation for lane detection and classification in autonomous driving. In: 2022 IEEE Intelligent Vehicles Symposium (IV). pp. 457–463 (2022). <https://doi.org/10.1109/IV51971.2022.9827450>
13. Jaderberg, M., Simonyan, K., Zisserman, A., kavukcuoglu, k.: Spatial transformer networks. In: Advances in Neural Information Processing Systems. vol. 28 (2015)
14. Jocher, G.: Yolov5 by ultralytics (2020). <https://doi.org/10.5281/zenodo.3908559>, <https://github.com/ultralytics/yolov5>
15. Klinghoffer, T., Philion, J., Chen, W., Litany, O., Gojcic, Z., Joo, J., Raskar, R., Fidler, S., Alvarez, J.M.: Towards viewpoint robustness in bird’s eye view segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 8515–8524 (October 2023)
16. Kong, L., Xie, S., Hu, H., Ng, L.X., Cottureau, B., Ooi, W.T.: Robodepth: Robust out-of-distribution depth estimation under corruptions. In: Advances in Neural Information Processing Systems. vol. 36, pp. 21298–21342 (2023)

17. Lee, Y., Kim, J., Willette, J., Hwang, S.J.: Mpvit: Multi-path vision transformer for dense prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7287–7296 (June 2022)
18. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (2019)
19. Saunders, K., Vogiatzis, G., Manso, L.J.: Self-supervised monocular depth estimation: Let’s talk about the weather. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 8907–8917 (October 2023)
20. Sun, P., Kretschmar, H., Dotiwala, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., Vasudevan, V., Han, W., Ngiam, J., Zhao, H., Timofeev, A., Ettinger, S., Krivokon, M., Gao, A., Joshi, A., Zhang, Y., Shlens, J., Chen, Z., Anguelov, D.: Scalability in perception for autonomous driving: Waymo open dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
21. Van Dijk, T., De Croon, G.: How do neural networks see depth in single images? In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2183–2191 (2019). <https://doi.org/10.1109/ICCV.2019.00227>
22. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
23. Vedit, V., Engilberge, M., Salzmann, M.: Learning transformations to reduce the geometric shift in object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 17441–17450 (June 2023)
24. Wang, K., Zhang, Z., Yan, Z., Li, X., Xu, B., Li, J., Yang, J.: Regularizing night-time weirdness: Efficient self-supervised monocular depth estimation in the dark. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 16055–16064 (October 2021)
25. Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004). <https://doi.org/10.1109/TIP.2003.819861>
26. Yogamani, S., Hughes, C., Horgan, J., Sistu, G., Varley, P., O’Dea, D., Uricar, M., Milz, S., Simon, M., Amende, K., Witt, C., Rashed, H., Chennupati, S., Nayak, S., Mansoor, S., Perrotton, X., Perez, P.: Woodscape: A multi-task, multi-camera fisheye dataset for autonomous driving. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2019). <https://doi.org/10.1109/ICCV.2019.00940>
27. Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
28. Zhao, C., Zhang, Y., Poggi, M., Tosi, F., Guo, X., Zhu, Z., Huang, G., Tang, Y., Mattoccia, S.: MonoViT: Self-Supervised Monocular Depth Estimation with a Vision Transformer. In: 2022 International Conference on 3D Vision (3DV). pp. 668–678. IEEE Computer Society (Sep 2022). <https://doi.org/10.1109/3DV57658.2022.00077>
29. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised Learning of Depth and Ego-Motion from Video. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6612–6619. IEEE, Honolulu, HI (Jul 2017). <https://doi.org/10.1109/CVPR.2017.700>