Calibrating Self-supervised Monocular Depth Estimation

Robert McCraith robert@robots.ox.ac.uk University of Oxford Lukas Neumann lukas@robots.ox.ac.uk University of Oxford

Andrea Vedaldi vedaldi@robots.ox.ac.uk University of Oxford

Abstract

In the recent years, many methods demonstrated the ability of neural networks to learn depth and pose changes in a sequence of images, using only self-supervision as the training signal. Whilst the networks achieve good performance, the often over-looked detail is that due to the inherent ambiguity of monocular vision they predict depth up to a unknown scaling factor. The scaling factor is then typically obtained from the LiDAR ground truth at test time, which severely limits practical applications of these methods.

In this paper, we show that incorporating prior information about the camera configuration and the environment, we can remove the scale ambiguity and predict depth directly, still using the self-supervised formulation and not relying on any additional sensors.

1 Introduction

Depth estimation is an important computer vision problem with applications in robotics, autonomous driving, augmented reality and scene understanding Poggi et al. [2018], Zhou et al. [2017], Fu et al. [2018], Yang et al. [2019], Madhu Babu et al. [2018]. Of particular theoretical and practical interest is estimating depth from a single RGB image, also known as *monocular depth estimation*. When multiple views are available, depth can be inferred from geometric principles by triangulating image correspondences; however, when only a single view is available, triangulation is not possible and the problem is ill posed. Despite this difficulty, reliable and accurate monocular depth estimation is critical for safety in many applications, with autonomous vehicles being a prime example.

Notably, humans are perfectly able to drive a car with just one eye, suggesting that they can infer depth well enough even from a single view. However, doing so requires prior information on the visual appearance and real-world sizes of typical scene elements, which can then be used to estimate the distance of known objects from the camera. In machine vision, this prior can be learned from 2D images labelled with ground-truth 3D information, extracted from a different sensing modality such as a LiDAR. When using additional sensors is impractical, self-supervised learning can be used instead Godard et al. [2017, 2019], Zhou et al. [2017], Mancini et al. [2016]. In *self-supervised learning*, certain relations or consistency of inputs, rather than ground truth labels, are exploited to train the system. In depth estimation, one typically uses the consistency between subsequent video frames or between stereo image pairs Godard et al. [2017], Zhou et al. [2017].

Whilst recent self-supervised monocular depth estimation methods achieve impressive performance, approaching fully-supervised systems, they all share an important practical limitation which limits

Machine Learning for Autonomous Driving Workshop at the 34th Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada.

their usefulness in real-world applications. Since reconstructing 3D geometry from images has an inherent scale ambiguity Faugeras et al. [2001], and since self-supervised methods only use visual inputs for training Godard et al. [2017, 2019], they do not predict the depth map d_I directly, but rather a scaled version $\Phi(I)$ of it relate to the true depth by an unknown scaling factor α_I , in the sense that $d_I = \alpha_I \Phi(I)$. For evaluation, the scaling factor α_I is not predicted but calculated at *test* time from the ground truth, usually as the the ratio between the median of the predicted depth values and the median of the ground-truth depth values. Furthermore, a different scaling factor is computed for each test image individually Godard et al. [2017]. However, in practical applications ground truth 3D data is not available to calibrate the system, especially in production. Thus, the problem is how to *calibrate* self-supervised depth estimation in order to obtain a physically-accurate prediction, without requiring the use of additional sensors. In this paper, we show that, in a driving scenario, this problem can indeed be solved reliably, robustly and efficiently assuming only knowledge of the camera intrinsics and a very limited amount of additional prior information on the geometry of the system. The output of our technique is a properly calibrated depth map, expressed in meters. The method is applicable to any self-supervised training paradigm and does not require any additional 3D ground truth at training or testing time. This is different from previous monocular depth estimation methods which discounting the scale ambiguity at test time, or use additional senors to remove it Guizilini et al. [2020]. Thus, we make two key contributions in this paper. First, we bring to the attention of the computer vision community the problem of calibrating self-supervised monocular depth estimation systems without resorting to additional sensors such as LiDARs. This is of clear importance if we wish these systems to be of direct practical value. Furthermore, we analyze to what extent the state-of-the-art existing methods depend on the availability of such data. Second, we propose a simple and yet very efficient calibration technique that *does not* make use of any additional sensors, especially of a complex and expensive nature such as a LiDAR. Instead, our method is 'vision closed' at training as well as a test time, in the sense that, just like self-supervised monocular depth estimation methods, it only requires images as input. The only additional information required for calibration is the approximate knowledge of a single constant which is trivially obtained from the construction of the system. The rest of the paper is structured as follows. In section 2 we give an overview of the state of the art, in section 3 our method is presented. Experimental validation is given in section 4 and the paper is concluded in section 5.

2 Previous Work

Depth Estimation. Because of the scale ambiguity inherent to predicting the depth from a single image, monocular depth estimation is an ill-posed problem and other (prior) knowledge has to be incorporated to remove the ambiguity. Scharstein et al. [2002] and more recently Flynn et al. [2016] use classical geometry to extract point-to-point matches between images and use triangulation to estimate depth.

The emergence of deep learning re-formulated the problem as a dense scene segmentation problem, where each image pixel is directly assigned a real value corresponding to the depth. A deep network is then trained to predict depth using supervision either from LiDAR, a RGB-D camera or a stereo pair. Eigen and Fergus [2015] regress depth in multiple scales, refining the depth maps from low to high spatial resolution. Xie et al. [2016] improve network architecture by adding skip connections, so that the network can also benefit from high resolution information. Laina et al. [2016] use deconvolutional segments to refine depth in a coarse-to-fine manner Garg et al. [2016], Kuznietsov et al. [2017], while Garg et al. [2016] use CRFs to improve fine details. Fu et al. [2018] introduced a novel discretized depth representation, which modifies the typical regression task into a classification problem, and using a novel loss function (ordinal regression) they significantly improve accuracy. More recently, Lee et al. [2019] use local planar guidance layers to improve performance. In practice, there are indeed other modalities/sensors that can be used to get depth — such as LiDAR, radar or a stereo camera pair. They all however have their own limitations: LiDARs are quite sensitive to weather Bijelic et al. [2018], the depth maps are sparse which may not be enough for far-away or small objects and they have low refresh rate. Similarly, radars suffer from reflections and interference and they struggle to detect small or slowly-moving objects rad. Stereo is very sensitive to precise calibration and the two cameras can become misaligned over time, greatly reducing the depth map accuracy. Generalizing, there is a clear need for redundancy — a moving autonomous vehicle or a robot in urban environment cannot simply stop working if for example one of the cameras becomes occluded by dirt or if the weather is not perfect, as that would be potentially very unsafe. By having a



Figure 1: Scaling factor inferred from LiDAR ground truth for every image in the KITTI test subset, as used in Godard et al. [2019]

reliable monocular method, this can be used in sensor fusion or as a fallback method, thus improving the overall safety of the system.

Self-Supervised Methods. The main struggle for the monocular depth methods is the requirement of vasts amounts of training data. The ground truth is typically captured by LiDARs, but this is expensive especially if large variety of driving scenarios and countries has to be covered, and the output is only a sparse point cloud. To alleviate the requirement of having expensive ground truth, recently there has been a surge in interest in unsupervised methods for depth map prediction. Xie et al. [2016] used stereo images in training discrete values for VR and 3D video applications and Garg et al. [2016] extended this approach to continuous values. More recently, Godard et al. [2017] added a left-right depth consistency and Zhou et al. [2017] generalized the approach to monocular sequences at training time, by predicting pose change between two sequential video frames. The sequential nature of the data however introduced some new challenges especially for non-stationary objects, which was addressed in Godard et al. [2019], that incorporates a loss which automatically excludes pixels which have become occluded or which correspond to moving objects. Similarly, Casser et al. [2019a] decompose the image into rigid and non-rigid component, thus reducing the re-projection error. All the above methods Godard et al. [2017, 2019] however share the same weakness which severely limits their applicability — the depth estimate is not calibrated, unlike the supervised methods Fu et al. [2018]. In other words, the depth values of self-supervised methods are not in meters but in some arbitrary unit, which moreover differs frame by frame (see fig. 1), and therefore these methods cannot be directly used to reason about the surrounding 3D world.

3 Method

In this section, we first describe the self-supervised training paradigm used by state-of-the-art monocular depth estimation methods Godard et al. [2017, 2019]. Then, we discuss how current methods deal with the issue of scale ambiguity by accessing ground truth 3D information at test time. Finally, we detail how basic prior information on a vehicle-mounted can be used to calibrate self-supervised depth map predictions to produce depth map with a correct physical scaling without requiring the acquisition of 3D or other information by means of additional sensors.

3.1 Self-supervised Depth Estimation

Given a pair of subsequent video frames I^t and I^{t+1} captured by a moving camera, under mild conditions such as Lambertian reflection, the image I^t is (approximately) a warp (deformed version) of image I^{t+1} Hartley and Zisserman [2004]. Moreover, the warp depends only on the *geometry* and motion of the scene, captured by the depth map D_t and the viewpoint change (R^t, T^t) . In other words, we can write $I^t \approx W(I^{t+1}, D_t, R_t, T_t, K)$, where W is a warp Jaderberg et al. [2015] which depends only on the depth D_t , the viewpoint change (R_t, T_t) , and the camera intrinsics K (which we assume known and constant).

The equation above provides a constraint that can be used to self-supervise a monocular depth estimation network Φ from knowledge of the video frames I^t and I^{t+1} alone. In more detail, we task two networks Φ and Ψ to predict respectively the depth $D^t = \Phi(I^t)$ from the first image and the motion $(R_t, T_t) = \Psi(I^t, I^{t+1})$ from the pair of images so as to correctly warp I^{t+1} into I^t , thus establishing the expected visual consistency (see fig. 2). This is done by minimizing the appearance loss between the original I^t and the synthesized image $\hat{I}^t = \mathcal{W}(I^{t+1}, D_t, R_t, T_t, K)$:

$$\mathcal{L}(I^t, I^{t+1}) = \alpha E_{\mathsf{p}}(I^t, \hat{I}^t) + E_{\mathsf{dis}}(I^t, \hat{I}^t) \tag{1}$$

$$\hat{I}^{t} = \mathcal{W}(I^{t+1}; D_t, R_t, T_t, K), \quad D_t = \Phi(I^t), \quad (R_t, T_t) = \Psi(I^t, I^{t+1})$$
(2)

The photometric loss term E_p in eq. (1) is the SSIM loss Wang et al. [2004], Zhao et al. [2016], whilst the E_{dis} term enforces smoothness Godard et al. [2017]. The whole network is trained end-to-end using standard back-propagation.

An analysis of the warp operator W Jaderberg et al. [2015] shows that the operator is invariant to multiplying the depth and the translation parameters by a constant α :

$$\mathcal{W}(I; D, R, T, K) = \mathcal{W}(I; \alpha D, R, \alpha T, K) \quad \alpha \in \mathbb{R}.$$
(3)

This shows that the network can only learn depth and translation up to an undetermined scaling factor α ; in particular, there is no reason for the learned scale to corresponds to the true physical scale of the scene. As a matter of fact, the model is not even forced to learn a scaling factor consistently across different pairs of frames (I^t, I^{t+1}) , which we show in empirically is in fact not the case. In particular, fig. 1 shows that the variation in scaling factor for different frames can be up to a factor of two.

3.2 Ground Truth Data used to Scale Depth at Test Time

l

Since the scale of the predicted depth $\Phi(I^t)$ is arbitrary, its use in downstream tasks that require a physical understanding of the scene (e.g. in robotics) impossible. Equally, all benchmarks for depth estimation Geiger et al. [2012] also require measurements in real units (meters), and therefore the depth map predictions $\Phi(I^t)$ cannot be assessed directly against these benchmarks. Instead, the common approach is to just marginalize out the scale at test time, finding the factor that best matches the predicted and ground-truth depth for each test image independently Godard et al. [2017], Casser et al. [2019b], Luo et al. [2018a], Godard et al. [2019]. Since this ground-truth information is obtained via a sensors such as a LiDAR, this is equivalent to calibrating the method against an additional sensor, which is not a realistic setup.

More formally, given an image I, the network outputs prediction a $\Phi(I)$ which is transformed to the final depth estimate as $d_I = \alpha_I \Phi(I)$ where:

$$\alpha_I = \frac{\text{median} \, d_I^{\text{gt}}}{\text{median} \, \Phi(I)} \tag{4}$$

where d_I^{gt} is the ground-truth depth map, usually created by projecting sparse LiDAR points onto the image plane, projected with the same viewpoint as the input image I.¹

¹Both images are masked such that the scaling factor is only calculated on points where the LiDAR has read data



Figure 2: Self-supervised monocular depth estimation pipeline Godard et al. [2017, 2019]



Parametric fit

Figure 3: Road model estimation. First we take the uncalibrated depth of the input image, combined with scene segmentation to extract only depth values belonging to the road. After further refinement, we project the points into 3D and fit a plane to them

3.3 Road Model Estimation

In order to remove the need for LiDAR ground truth at test time, we exploit prior knowledge of the environment and of the camera setup, especially the camera height. Because cars drive on roads and we know that the camera is at certain height above the road, we can exploit this constraint to calibrate the depth map to real-world values.

In order to do so, we first need to automatically estimate a road model in every test image. In order to account for the fact that many roads are not perfectly flat, more typically they slope up/down or are higher on one side than the other, or that the car tilts during acceleration and deceleration, we estimate the pitch and roll of the road by fitting a plane to the raw depth map. We only use values for pixels that are classified as road by a pre-trained semantic segmentation model Zhou et al. [2018], and whose |X| and Z co-ordinate² is below a certain threshold (see Section 4.3).

We then fit these points using Least Median of Squares regression Rousseeuw [1984] to get the road plane estimate in the 3D world $a_1X + a_2Y + a_3Z + c = 0$. We know that the 3D point on the road right below the camera has the co-ordinate [0, -h, 0], where h is the camera height, and therefore we can infer the following relation for the scaling factor

$$\alpha_I = \frac{c}{h} \tag{5}$$

Compared to eq. (4), the scaling factor eq. (5) now only relies on the visual information, and therefore can be obtained without any LiDAR input.

4 **Experiments**

In this section, we compare our calibration method (Ours) to: (1) the un-calibrated outputs of the network $D = \Phi(I)$ (Raw); (2) computing a per-frame calibration factor using eq. (4) with

²The world co-ordinates X, Z to filter out pixels above the threshold are obtained using camera intrinsics and assuming the road is perfectly flat, i.e. Y = -h



Figure 4: Qualitative depth estimation examples from the KITTI dataset (inverse depth shown). Monodepth2 Godard et al. [2019] output values (Raw Output) are scaled by comparing the output to the ground truth for every test image (GT Scale). Using a single scaling factor from the training set (Single Scale) is significantly worse. Using road model (ours) to estimate the scaling factor achieves significantly better results. All images use the same color coding.

access to the 3D ground-truth (GT Scaling); and (3) the same as (2), but by computing a single scaling factor from either all the training or testing frames (GT Single Scaling). After a qualitative and quantitative comparison with these techniques and state-of-the-art monocular depth estimation networks (both supervised and unsupervised), we ablate our method, showing the importance of the various components, and study sensitivity to its parameters.

Implementation details. In all our experiments, we used the Godard et al. [2019] pre-trained model. In line with prior work, we use the Eigen and Fergus [2015] data split of KITTI dataset Geiger et al. [2012].

4.1 Qualitative comparison

We first look at the optimal scaling factor determined via GT Scaling via access to the ground-truth at test time (section 3.2, Godard et al. [2017, 2019], Zhou et al. [2017]). As is observable in Figure 1,

Method	Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
(a) GT Single Scaling (training set)	0.125	0.942	5.045	0.208	0.84	0.953	0.979
(b) GT Single Scaling (testing set)	0.126	0.952	4.999	0.204	0.848	0.954	0.98
(c) Fixed road plane	0.132	1.073	5.035	0.203	0.86	0.954	0.977
(d) Ours	0.113	0.916	4.974	0.199	0.857	0.945	0.968

Table 1: Comparison of different depth map scaling methods on the KITTI testing subset.

the factor selected by GT Scaling varies wildly even in a single video sequence. This is illustrated in fig. 4 for two different input images. GT Scaling chooses factor 19.63 for the left image vs GT Single Scaling (median on the training set) of 30.462. This means that, for this image, the network predicts a depth map where objects are 50% farther away than for the median case. The image to the right is the opposite, as GT Scaling determines the best factor to be 40.55, so objects are predicted to be 33% than the median case. Given these differences, it is clear that there is no single scaling factor that results in a good fit for all test frames; hence, below we find it unsurprising that using a single scaling factor over the entire test set (GT Single Scaling) produces inaccurate results overall.

By comparison, our scaling technique predicts scaling factors of 19.63 and 36.4 for the two images respectively, which are close to the output of GT Scaling. Hence, our system produces results significantly closer to the per-frame GT Scaling factors than GT Single Scaling while having *no* access to ground-truth (LiDAR) 3D information at training or test time. This useful for autonomous vehicles that wish to adapt to scenes where it frequently drives rather than examples in a training set as is done in McCraith et al. [2020].

4.2 Quantitative comparison

Method	Train	GT@Test	Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Eigen and Fergus [2015]	D	X	0.203	1.548	6.307	0.282	0.702	0.890	0.890
Liu et al. [2015]	D	X	0.201	1.584	6.471	0.273	0.680	0.898	0.967
Klodt and Vedaldi [2018]	D*M	X	0.166	1.490	5.998		0.778	0.919	0.966
Nath Kundu et al. [2018]	D*	X	0.167	1.257	5.578	0.237	0.771	0.922	0.971
Kuznietsov et al. [2017]	DS	X	0.113	0.741	4.621	0.189	0.862	0.960	0.986
Yang et al. [2018a]	D*S	X	0.097	0.734	4.442	0.187	0.888	0.958	0.980
Luo et al. [2018b]	DS	X	0.094	0.626	4.252	0.177	0.891	0.965	0.984
Guo et al. [2018]	DS	X	0.096	0.641	4.095	0.168	0.892	0.967	0.986
Fu et al. [2018]	D	X	0.072	0.307	2.727	0.120	0.932	0.984	0.994
Zhou et al. [2017]†	M	1	0.183	1.595	6.709	0.270	0.734	0.902	0.959
Yang et al. [2018c]	M	1	0.182	1.481	6.501	0.267	0.725	0.906	0.963
Mahjourian et al. [2018]	M	1	0.163	1.240	6.220	0.250	0.762	0.916	0.968
Yin and Shi [2018]†	M	1	0.149	1.060	5.567	0.226	0.796	0.935	0.975
Wang et al. [2018]	M	1	0.151	1.257	5.583	0.228	0.810	0.936	0.974
Zou et al. [2018]	M	1	0.150	1.124	5.507	0.223	0.806	0.933	0.973
Yang et al. [2018b]	M	1	0.162	1.352	6.276	0.252			
Ranjan et al. [2019]	M	1	0.148	1.149	5.464	0.226	0.815	0.935	0.973
Luo et al. [2018a]	M	1	0.141	1.029	5.350	0.216	0.816	0.941	0.976
Casser et al. [2019b]	M		0.141	1.026	5.291	0.215	0.816	0.945	0.979
Godard et al. [2019]	M	1	0.115	0.903	4.863	0.193	0.877	0.959	0.981
Ours	M	X	0.113	0.916	4.974	0.199	0.857	0.945	0.968

Table 2: Depth estimation accuracy on the KITTI test set. D — Depth supervision, D* — Auxiliary depth supervision, M — Self-supervised mono, GT@Test — uses elements of LiDAR ground truth at test time, \dagger — Newer results from GitHub,+ pp — With post-processing. For red metrics, the lower is better; for blue metrics, the higher is better. Best results in each category are in **bold**; second-best underlined

First, in table 1 we contrast our method (d) to GT Single Scaling, fixing the scaling factor using respectively the training and testing subset of the data (a) and (b). We note that our approach is substantially better than both (0.113 vs \geq 0.125 AbsRel). This is because, while GT Single Scaling has access to 3D ground-truth, it uses a fixed scaling factor for all frames, and, as shown above, no single scaling factor can work well. Remarkably, our method is comparable to GT Scaling *as well* (the latter corresponds to the penultimate row of table 2), matching it, in particular, in the Abs Rel metric, despite the fact that GT Scaling chooses the best possible scaling factor for each frame

Length	Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
6	0.338	2.857	7.47	0.353	0.068	0.132	0.155
10	0.12	0.968	5.013	0.202	0.838	0.933	0.957
15	0.116	0.942	5	0.2	0.853	0.944	0.968
20	0.115	0.932	4.986	0.2	0.856	0.945	0.968
25	0.117	0.956	5.002	0.201	0.856	0.944	0.969
30	0.114	0.926	4.98	0.199	0.856	0.945	0.968
40	0.116	0.933	4.985	0.2	0.856	0.946	0.971
60	0.115	0.928	4.979	0.2	0.857	0.947	0.971
80	0.116	0.941	4.99	0.2	0.857	0.945	0.97

Table 3: Road model distance (length) ablation

Width (m)	Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
0.5	0.119	0.944	5.02	0.206	0.838	0.932	0.957
1	0.116	0.926	4.989	0.201	0.845	0.937	0.961
2	0.114	0.918	4.981	0.2	0.856	0.944	0.968
3	0.113	0.916	4.974	0.199	0.857	0.945	0.968
4	0.115	0.936	4.99	0.2	0.856	0.945	0.968
5	0.114	0.923	4.974	0.199	0.858	0.946	0.97
10	0.116	0.933	4.986	0.2	0.855	0.946	0.969
15	0.116	0.933	4.987	0.2	0.855	0.946	0.969

Table 4: Road model width ablation. Points are considered to create the model if |X| < Width and the distance is below 30 meters.

individually against the ground-truth. From the same table, we see that this is obtained against a model, Monodepth V2, which is state-of-the-art, resulting for the first time in excellent *calibrated* self-supervised monocular depth estimation from *vision alone*.

4.3 Ablation and tuning

Road model. Recall that our method is based on estimating the full 3 DoF of the ground plane. First, we test whether this is necessary. In order to do so, we assume instead that the plane is exactly horizontal and at the fixed canonical height below the camera. Then, we use the fixed plane to generate a pseudo-LiDAR map for the road pixels and use GT Scaling against those pseudo-ground-truth values (instead of the actual GT value) in order to determine the scaling factor for each frame. A similar fixed pseudo-LiDAR plane was also used, for example, in Choe et al. [2019] in order to perform 3D object detection. The result of this is shown in table 1 row (c): the fact that this simple fixed-plane model ignores the tendency of real roads to have inclines and declines as well as the cameras ability to have non-negligible pitch and roll during regular car motion which greatly effects it's depth prediction in the far range.

Tuning. Next, we assess the sensitivity of our methods to various parameters and determine their optimal values. In section 3.3 we first take to varying the maximum distance of points on the road we use for the plane fitting. Similar to Man et al. [2019] we find that using points predicted to be under 30 meters from our camera works best for fitting our ground plane as seen in table 4. In a similar fashion we explore the maximum left-right distance from which we consider points from our fit. Typical roads are between 2.75 and 3.75 meters wide so it is within reason that only points within 3m left of right of the car work best with a gradual drop off in performance above this and an insufficient number of points below. Note that the method is not overly sensitive to a specific parameter setting, likely due to the use of robust estimator.

5 Conclusion

In this paper, we highlighted the limitation of self-supervised depth estimation methods and their reliance on LiDAR data at test time. We additionally showed how to overcome this issue by incorporating prior information about camera configuration and the environment, and we achieved comparable performance to the state of the art through vision only, without relying on any additional sensors.

References

- Understanding radar for automotive (ADAS) solutions. https://www.pathpartnertech. com/understanding-radar-for-automotive-adas-solutions/. Accessed: 2019-11-20.
- Mario Bijelic, Tobias Gruber, and Werner Ritter. A benchmark for lidar sensors in fog: Is detection breaking down? In 2018 IEEE Intelligent Vehicles Symposium (IV), pages 760–767. IEEE, 2018.
- Vincent Casser, Soeren Pirk, Reza Mahjourian, and Anelia Angelova. Unsupervised monocular depth and ego-motion learning with structure and semantics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019a.
- Vincent Casser, Soeren Pirk, Reza Mahjourian, and Anelia Angelova. Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos. In AAAI, 2019b.
- Jaesung Choe, Kyungdon Joo, François Rameau, Gyu Min Shim, and In So Kweon. Segment2regress: Monocular 3d vehicle localization in two stages. In *Robotics: Science and Systems*, 2019.
- David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, pages 2650–2658, 2015.
- Olivier D. Faugeras, Quang-Tuan Luong, and Théodore Papadopoulo. *The geometry of multiple images the laws that govern the formation of multiple images of a scene and some of their applications*. MIT Press, 2001.
- John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. Deepstereo: Learning to predict new views from the world's imagery. In *ICCV*, pages 5515–5524, 2016.
- Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *CVPR*, June 2018.
- Ravi Garg, Vijay Kumar BG, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *European Conference on Computer Vision*, pages 740–756. Springer, 2016.
- Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In *CVPR*, 2012.
- Clément Godard, Oisin Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *ICCV*, pages 270–279, 2017.
- Clément Godard, Oisin Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into selfsupervised monocular depth estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3828–3838, 2019.
- Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- Xiaoyang Guo, Hongsheng Li, Shuai Yi, Jimmy Ren, and Xiaogang Wang. Learning monocular depth by distilling cross-domain stereo networks. In *ECCV*, 2018.
- R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.
- Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015.
- Maria Klodt and Andrea Vedaldi. Supervising the new with the old: learning SFM from SFM. In *ECCV*, 2018.
- Yevhen Kuznietsov, Jorg Stuckler, and Bastian Leibe. Semi-supervised deep learning for monocular depth map prediction. In *ICCV*, pages 6647–6655, 2017.

- Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In 2016 Fourth international conference on 3D vision (3DV), pages 239–248. IEEE, 2016.
- Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326*, 2019.
- Fayao Liu, Chunhua Shen, Guosheng Lin, and Ian Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE TPAMI*, 38(10):2024–2039, 2015.
- Chenxu Luo, Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, Ram Nevatia, and Alan Yuille. Every pixel counts++: Joint learning of geometry and motion with 3D holistic understanding. *arXiv*, 2018a.
- Yue Luo, Jimmy Ren, Mude Lin, Jiahao Pang, Wenxiu Sun, Hongsheng Li, and Liang Lin. Single view stereo matching. In *CVPR*, 2018b.
- V. Madhu Babu, K. Das, A. Majumdar, and S. Kumar. Undemon: Unsupervised deep network for depth and ego-motion estimation. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 1082–1088, Oct 2018. doi: 10.1109/IROS.2018.8593864.
- Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and egomotion from monocular video using 3D geometric constraints. In *CVPR*, 2018.
- Yunze Man, Xinshuo Weng, Xi Li, and Kris Kitani. Groundnet: Monocular ground plane normal estimation with geometric consistency. In *Proceedings of the 27th ACM International Conference* on Multimedia, MM '19, page 2170–2178, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450368896. doi: 10.1145/3343031.3351068. URL https://doi.org/ 10.1145/3343031.3351068.
- Michele Mancini, Gabriele Costante, Paolo Valigi, and Thomas A Ciarfuglia. Fast robust monocular depth estimation for obstacle detection with fully convolutional networks. In 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 4296–4303. IEEE, 2016.
- Robert McCraith, Lukas Neumann, Andrew Zisserman, and Andrea Vedaldi. Monocular depth estimation with self-supervised instance adaptation, 2020.
- Jogendra Nath Kundu, Phani Krishna Uppala, Anuj Pahuja, and R. Venkatesh Babu. AdaDepth: Unsupervised content congruent adaptation for depth estimation. In *CVPR*, 2018.
- Matteo Poggi, Filippo Aleotti, Fabio Tosi, and Stefano Mattoccia. Towards real-time unsupervised monocular depth estimation on cpu. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 5848–5854. IEEE, 2018.
- Anurag Ranjan, Varun Jampani, Kihwan Kim, Deqing Sun, Jonas Wulff, and Michael J Black. Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation. In CVPR, 2019.
- Peter J Rousseeuw. Least median of squares regression. *Journal of the American statistical association*, 79(388):871–880, 1984.
- Daniel Scharstein, Richard Szeliski, and Rick Szeliski. A taxonomy and evaluation of dense twoframe stereo correspondence algorithms. *International Journal of Computer Vision*, 47:7–42, May 2002.
- Chaoyang Wang, Jose Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *CVPR*, 2018.
- Zhou Wang, Alan Conrad Bovik, Hamid Rahim Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 2004.
- Junyuan Xie, Ross Girshick, and Ali Farhadi. Deep3d: Fully automatic 2d-to-3d video conversion with deep convolutional neural networks. In *European Conference on Computer Vision*, pages 842–857. Springer, 2016.

- Gengshan Yang, Peiyun Hu, and Deva Ramanan. Inferring distributions over depth from a single image. In 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2019.
- Nan Yang, Rui Wang, Jörg Stückler, and Daniel Cremers. Deep virtual stereo odometry: Leveraging deep depth prediction for monocular direct sparse odometry. In *ECCV*, 2018a.
- Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, and Ram Nevatia. LEGO: Learning edge with geometry all at once by watching videos. In *CVPR*, 2018b.
- Zhenheng Yang, Peng Wang, Wei Xu, Liang Zhao, and Ramakant Nevatia. Unsupervised learning of geometry with edge-aware depth-normal consistency. In *AAAI*, 2018c.
- Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1983–1992, 2018.
- Hang Zhao, Orazio Gallo, Iuri Frosio, and Jan Kautz. Loss functions for image restoration with neural networks. *IEEE Transactions on computational imaging*, 3(1):47–57, 2016.
- Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *International Journal on Computer Vision*, 2018.
- Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1851–1858, 2017.
- Yuliang Zou, Zelun Luo, and Jia-Bin Huang. DF-Net: Unsupervised joint learning of depth and flow using cross-task consistency. In *ECCV*, 2018.