

Geometry meets semantics for semi-supervised monocular depth estimation

Pierluigi Zama Ramirez, Matteo Poggi, Fabio Tosi,
Stefano Mattoccia, and Luigi Di Stefano

University of Bologna,
Viale del Risorgimento 2, Bologna, Italy
{pierluigi.zama,m.poggi,fabio.tosi5}@unibo.it

Abstract. Depth estimation from a single image represents a very exciting challenge in computer vision. While other image-based depth sensing techniques leverage on the geometry between different viewpoints (*e.g.*, stereo or structure from motion), the lack of these cues within a single image renders ill-posed the monocular depth estimation task. For inference, state-of-the-art encoder-decoder architectures for monocular depth estimation rely on effective feature representations learned at training time. For unsupervised training of these models, geometry has been effectively exploited by suitable images warping losses computed from views acquired by a stereo rig or a moving camera. In this paper, we make a further step forward showing that learning semantic information from images enables to improve effectively monocular depth estimation as well. In particular, by leveraging on semantically labeled images together with unsupervised signals gained by geometry through an image warping loss, we propose a deep learning approach aimed at joint semantic segmentation and depth estimation. Our overall learning framework is semi-supervised, as we deploy groundtruth data only in the semantic domain. At training time, our network learns a common feature representation for both tasks and a novel cross-task loss function is proposed. The experimental findings show how, jointly tackling depth prediction and semantic segmentation, allows to improve depth estimation accuracy. In particular, on the KITTI dataset our network outperforms state-of-the-art methods for monocular depth estimation.

1 Introduction

Depth sensing has always played an important role in computer vision because of the increased reliability brought in by availability of 3D data in several key tasks. In this context, dense depth estimation from images compares favorably to active sensors, such as Time-of-Flight cameras or Lidars, due to the latter either featuring short acquisition ranges or being cumbersome and much more expensive. Although traditional image-based approaches rely on multiple acquisitions from different viewpoints, like in binocular or multi-view stereo, depth

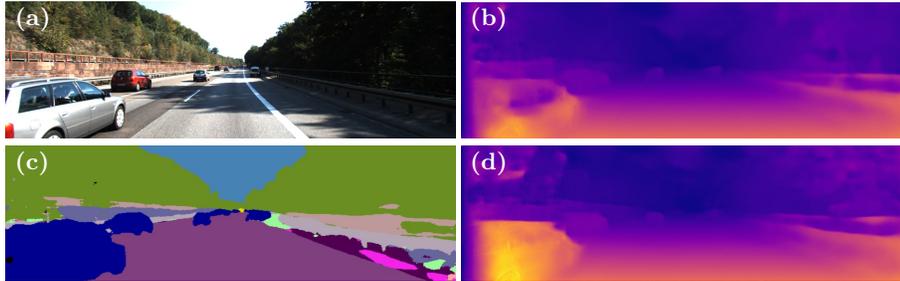


Fig. 1. Joint depth from mono and semantic segmentation. (a) Input image, (b) depth map by state-of-the-art method [1], (c) semantic and (d) depth maps obtained by our network.

estimation from a single image is receiving ever-increasing attention due to its unparalleled potential for seamless, cheap and widespread deployment. Recently proposed supervised learning frameworks based on Convolutional Neural Networks (CNNs) have achieved excellent results in this task, though they require massive amounts of training images labeled with per pixel groundtruth depth measurements. Obtaining these labels is particularly challenging and costly as it relies on expensive active sensors, such as high-end Lidars, which typically provide sparse and noisy measurements requiring further automatic or manual processing [2,3]. To address these issues, multiple acquisitions by a stereo rig [1] or a single moving camera [4] may be used to obtain supervision signals by warping different views according to the estimated depth and measuring the associated image re-projection error.

As for the depth-from-mono task, geometry cues are required at training time only. For inference, the depth estimation network is mainly driven by the learned global image context. Evidence of this can be gathered by running a monocular depth estimator, trained in either supervised or unsupervised manner, on imagery dealing with slightly different environments and observing how it may succeed in yielding reasonable results. These considerations suggest that the feature representation learned to predict depth from a single image is quite tightly linked to the *semantic* content of the scene, thus it leads us to conjecture that guiding the network through explicit knowledge about scene semantics may improve effectiveness in the depth-from-mono task. Moreover, the very recent work by Zamir, Amir R., et al. [5], supports the argument that learning features from multiple tasks is beneficial to performance as there exist relevant dependencies between visual tasks. Although [5] is based on fully-supervised learning, our paper shows that the correlation between semantic segmentation and depth estimation can be exploited also within a semi-supervised learning framework.

Thus, in this paper, we propose to train a CNN architecture to perform both semantic segmentation and depth estimation from a single image. By optimizing our model jointly on the two tasks, we enable it to learn a more effective feature

representation which yields improved depth estimation accuracy. We rely on an unsupervised image re-projection loss [1] to pursue depth prediction whilst we let the network learn semantic information from the observed scene by supervision signals from pixel-level groundtruth semantic maps. Thus, with respect to recent work [1], our proposal requires semantically annotated imagery, thereby departing from a totally unsupervised towards a semi-supervised learning paradigm (*i.e.* unsupervised for depth and supervised for semantics). Yet, though manual annotation of per-pixel semantic labels is tedious, it is much less prohibitive than collecting groundtruth depths. Besides, while the former task may be performed off-line after acquisition, as recently proposed for some images of the KITTI dataset [6], one may very unlikely obtain depth labels out of already collected frames.

To the best of our knowledge, this paper is the first to propose integration of unsupervised monocular depth estimation with supervised semantic segmentation. By applying this novel paradigm, we improve a state-of-the-art encoder-decoder depth estimation architecture [1] according to two main contributions:

- we propose to introduce an additional decoder stream based on the same features as those deployed for depth estimation and trained for semantic segmentation; thereby, the overall architecture is trained to optimize both tasks jointly.
- we propose a novel cross-task loss term, the *depth-semantic-discontinuity* loss \mathcal{L}_{dsd} , aimed at enforcing spatial proximity between depth discontinuities and semantic contours.

Experimental results on the KITTI dataset prove that tackling the two tasks jointly does improve monocular depth estimation. For example, Fig. 1 suggests how recognizing objects like cars (c) can significantly ameliorate depth estimation (d) with respect to a depth-from-mono approach lacking any awareness about scene semantics (b). It is also worth highlighting that, unlike all previous unsupervised frameworks in this field, our proposal delivers not only the depth map (Fig.1 (d)) but also the semantic segmentation of the input image (Fig.1, (c)) by an end-to-end training process.

2 Related work

We review here the literature dealing with unsupervised monocular depth estimation and semantic segmentation, both relevant to our work.

Unsupervised Monocular Depth. Single view depth estimation [7,8,9,10] gained much more popularity in the last years thanks to the increasing availability of benchmarks [11,12]. Moreover, casting depth estimation as an image reconstruction task represents a very attractive way to overcome the need for expansive, groundtruth labels by using a large amount of unsupervised imagery. The work by Garg *et al.*[13] represents the first, pivotal step in this direction, proposing a network for monocular depth estimation by deploying, at training time, view reconstruction loss together with actual stereo pairs as supervision.

Then, Godard *et al.*[1] introduced bilinear warping [14] alongside with more robust reconstruction losses, thereby achieving state-of-the-art performance for monocular depth estimation. This approach was extended to embedded systems [15], using a virtual trinocular setup at training time [16] or a GAN framework [17], Kuznietsov *et al.*[18] trained a network in a semi-supervised manner, by merging the unsupervised image reconstruction error together with the contribution from sparse depth groundtruth labels. While the above mentioned techniques require rectified stereo pairs at training time, Zhou *et al.*[4] proposed to train a network to infer depth from video sequences. This network computes a reconstruction loss between subsequent frames and, at the same time, predicts the relative poses between adjacent frames. Therefore, this method enables a fully-monocular setup whereby stereo pairs are no longer required for training. However, this strategy comes to a price in performance [4], delivering less accurate depth estimations compared to [1]. More recent works aimed at improving the single camera supervision approach because of its easiness of use, introducing 3D point-cloud alignment [19], differentiable visual odometry [20], joint optical flow estimation [21], or combining both stereo and video sequences supervision [20]. Nevertheless, none of them actually outperforms the synergy of stereo supervision and network model deployed by Godard *et al.* [1]. For this reason, in this paper we follow the guidelines of [1], currently the undisputed state-of-the-art for unsupervised monocular depth estimation.

Semantic Segmentation. While most early proposals relied on hand-crafted features and classifiers like Random Forests [22] or Support Vector Machines [23], nowadays pixel-level semantic segmentation approaches mainly exploit fully convolutional neural networks [24]. Compared to previous methods, the key advantage of the present-day strategy concerns the ability to *automatically* learn a better feature representation. Architectures for semantic segmentation focus on exploitation of contextual information and can be divided into five main groups. In the first, we find multi-scale prediction models [25,26,27,28], whereby the same architecture takes inputs at different scales so to extract features at different contextual levels. The second group consists of encoder-decoder architectures. The encoder is in charge of extracting low-resolution features from high-resolution inputs while the decoder should be able to recover fine object details from the feature representation so as to yield a high-resolution output map [24,29,30,31]. The third group accounts for models which encode long range context information exploiting Conditional Random Fields either as a post processing module [27] or as an integral part of the network [32]. The fourth group includes models relying on spatial pyramid pooling to extract context information at different levels [33,27,27]. Finally, the fifth group deals with models deploying atrous-convolutions rather than the standard convolution operator to extract higher resolution features while keeping a large receptive field to capture long-range information [34,35]. Our learning framework deploys an encoder-decoder architecture to fit with the monocular depth model in [1]. Thus, our semantic segmentation network may be thought of as belonging to the second group.

There exist also several works that, akin our paper, pursue joint estimation of depth and semantics from a single image. Ladicky et al. [36] combine depth regression with semantic classification deploying a bag-of-visual-words model and a boosting algorithm. Mousavian et al. [37] deploy a multi-scale CNN to estimate depth and used within a CRF to obtain semantic segmentation. Wang et al. [38] use local and global CNNs to extract pixels and regions potential which are fed to a CRF. More recent works, such as [39], demonstrate that jointly performing multiple task with adequate weighting of each task can be exploited to achieve better results. However, all these methods require groundtruth labels for both depth and semantics and are trained through multiple stages, whereas we propose to boost self-supervised depth estimation with easier to obtain semantic supervision only.

3 Proposed method

In this section, we present our proposal for joint semantic segmentation and depth estimation from a single image. We first explain the main intuitions behind our work, then we describe the network architecture and the loss functions deployed in our deep learning framework.

Estimating the distance of objects from a camera through a single acquisition is an ill-posed problem. While other techniques can effectively measure depth based on features extracted from different view points (*e.g.*, binocular stereo allows for triangulating depth from point matches between two synchronized frames), monocular systems cannot rely on geometry constraints to infer distance unambiguously. Despite this lack of information, modern deep learning monocular frameworks achieved astounding results by learning effective feature representations from the observed environment. Common to latest work in this field [40,8,1,4] is the design of deep encoder-decoder architectures, with a first contractive portion progressively decimating image dimensions to reduce the computational load and increase the receptive field, followed by an expanding portion which restores the original input resolution. In particular, the encoding layers learn a high level feature representation crucial to infer depth. Although it is hard to tell what kind of information the network is actually learning at training time, we argue semantic to play an important role. Recent works like [1,4] somehow support this intuition. Indeed, although the authors trained and evaluated their depth estimators on the KITTI dataset [3], a preliminary training on CityScapes [41] turned out beneficial to achieve the best accuracy with both frameworks, despite the very different camera setup between the two datasets. Common to the datasets is, in fact, the kind of sensed environment and, thus, the overall semantics of the scenes under perception. This observation represents the main rationale underpinning our proposal. By explicitly training the network to learn the semantic context of the sensed environment we shall expect to enrich the feature representation resulting from the encoding module and thus obtain a more accurate depth estimation. This may be realized by a deep model in which a single encoder is shared between two decoders in charge of provid-

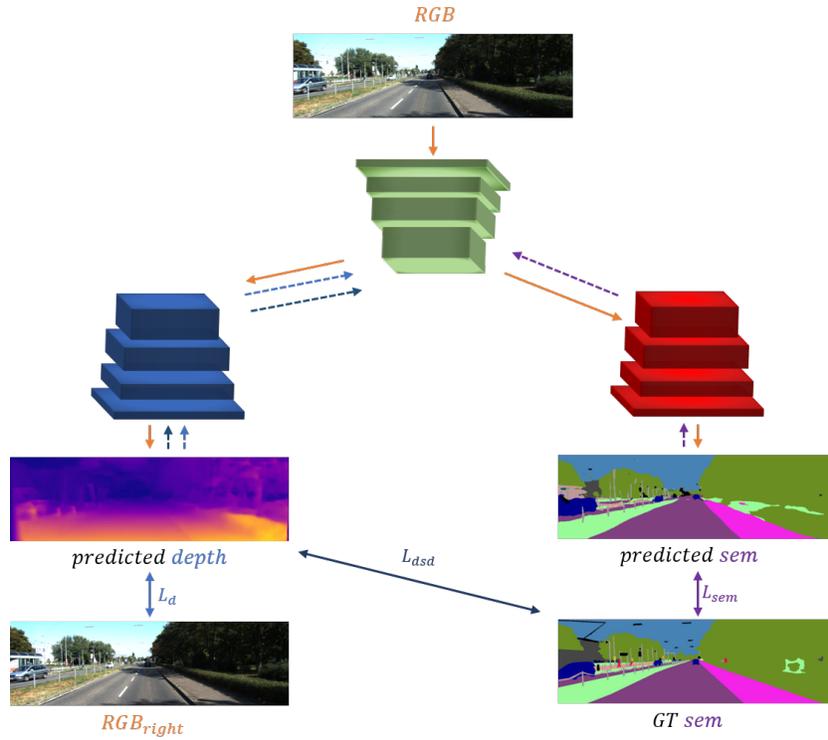


Fig. 2. Schematic representations of the proposed network architecture and semi-supervised learning framework. A single encoder (green) is shared between a depth (blue) and a semantic (red) decoder. The depth decoder is optimized to minimize \mathcal{L}_d and \mathcal{L}_{dsd} , the semantic decoder to minimize \mathcal{L}_s .

ing, respectively, a depth map and a semantic segmentation map. Accordingly, minimization of the errors with respect to pixel-level semantic labels provides gradients that flow back into the encoder at training time, thereby learning a shared feature representation aware of both depth prediction as well as scene semantics. According to our claim, this should turn out conducive to better depth prediction.

Inspired by successful attempts to predict depth from a single image, we design a suitable encoder-decoder architecture for joint depth estimation and semantic segmentation. The encoder is in charge of learning a rich feature representation by increasing the receptive field of the network while reducing the input dimension and computational overhead. Popular encoders for this task are VGG [42] and ResNet50 [43]. The decoder restores the original input resolution by means of up-sampling operators followed by 3×3 convolutions linked by means of skip connections with the encoder at the corresponding resolution. As illustrated in Fig. 2, to infer both depth and semantics we keep relying on a

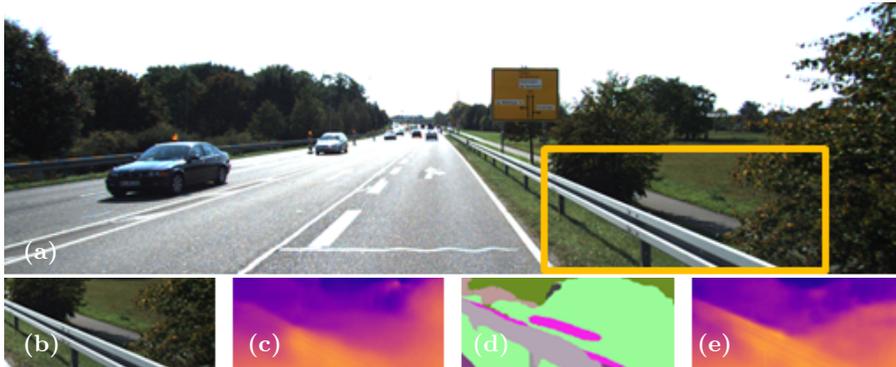


Fig. 3. Example of improved depth estimation enabled by semantic knowledge. (a) input image, (b) region extracted from the scene, (c) depth map predicted by [1], depth (d) and semantic (e) maps predicted by our framework. We can clearly notice how the the structure of the guard rail is better preserved by our method (e) compared to [1] in (c).

single encoder (green) and replicate the decoder to realize a second estimator. The two decoders (blue, red) do not share weights and are trained to minimize different losses, which deal with the depth prediction (blue) and semantic segmentation (red) tasks. While the two decoders are updated by different gradients flows, the shared encoder (green) is updated according to both flows, thereby learning a representation optimized jointly for the two tasks. We validate our approach by extending the architecture proposed by Godard *et al.*[1] for monocular depth estimation: the encoder produces two inverse depth (i.e., disparity) maps by processing the left image of a stereo pair. Then, the right image is used to obtain supervision signals by warping the left image according to the estimated disparities, as explained in the following section. Moreover, we will publicly release both the source code and the trained models once the reviewing process is complete to grant reproducibility.

Figure 3 shows how the shared representation used to jointly tackle both tasks enables to reconstruct better shapes when estimating depth (e) thanks to the semantic context (d) learned by the network compared to standalone learning of depth (c) as in [1].

3.1 Loss functions

To train the proposed architecture, we rely on the following multi-task loss function

$$\mathcal{L}_{tot} = \alpha_d \mathcal{L}_d + \alpha_s \mathcal{L}_s + \alpha_{dsd} \mathcal{L}_{dsd} \quad (1)$$

which consists in the weighted sum of three terms, namely the *depth* (\mathcal{L}_d), *semantic* (\mathcal{L}_s) and *depth-semantic-discontinuity* (\mathcal{L}_{dsd}) terms. As shown in in

Fig. 2, each term back-propagates gradients through a different decoder: in particular, \mathcal{L}_d and \mathcal{L}_{dsd} through the depth (blue) decoder whilst \mathcal{L}_s through the semantic (red) decoder. All gradients then converge so to flow back into the shared (green) encoder.

Depth term The depth term, \mathcal{L}_d , in our multi-task loss is computed according to the unsupervised training paradigm proposed by Godard *et al.*[1]:

$$\mathcal{L}_d = \beta_{ap}(\mathcal{L}_{ap}^l + \mathcal{L}_{ap}^r) + \beta_{ds}(\mathcal{L}_{ds}^l + \mathcal{L}_{ds}^r) + \beta_{lr}(\mathcal{L}_{lr}^l + \mathcal{L}_{lr}^r) \quad (2)$$

where the loss consists in the weighted sum of three terms, namely the *appearance*, *disparity smoothness* and *left-right consistency* terms. The first term measures the image re-projection error by means of the SSIM [44] and L1 difference between the original and warped images, I and \tilde{I} :

$$\mathcal{L}_{ap}^l = \frac{1}{N} \sum_{i,j} \gamma \frac{1 - SSIM(I_{i,j}^l, \tilde{I}_{i,j}^l)}{2} + (1 - \gamma) \|I_{i,j}^l - \tilde{I}_{i,j}^l\| \quad (3)$$

The smoothness term penalizes large disparity differences between neighboring pixels along the x and y directions unless these occur in presence of strong intensity gradients in the reference image I

$$\mathcal{L}_{ds}^l = \frac{1}{N} \sum_{i,j} |\delta_x d_{i,j}^l| e^{-\|\delta_x I_{i,j}^l\|} + |\delta_y d_{i,j}^l| e^{-\|\delta_y I_{i,j}^l\|} \quad (4)$$

Finally, the left-right consistency enforces coherence between the predicted disparity maps, d^l and d^r , for left and right images:

$$\mathcal{L}_{lr}^l = \frac{1}{N} \sum_{i,j} |d_{i,j}^l - d_{i,j+d_{i,j}^l}^r| \quad (5)$$

As proposed in [1], in our learning framework \mathcal{L}_d is computed at four different scales.

Semantic term The semantic term \mathcal{L}_s within our total loss is given by the standard cross-entropy between the predicted and groundtruth pixel-wise semantic labels:

$$\mathcal{L}_s = \mathcal{C}(p_t, \bar{p}_t) = H(p_t, \bar{p}_t) + KL(p_t, \bar{p}_t) \quad (6)$$

where H denotes the entropy and KL the KL -divergence. The semantic term, \mathcal{L}_s , is computed at full resolution only.

Depth-semantic-discontinuity term We also introduce a novel cross-task loss term aimed at enforcing an explicit link between the two learning tasks by leveraging on the groundtruth pixel-wise semantic labels to improve depth prediction. We found that the most effective manner to realize this consists in deploying the observation that depth discontinuities are likely to co-occur with semantic boundaries. Accordingly, we have designed the following *depth-semantic-discontinuity*, \mathcal{L}_{dsd} , term:

$$\mathcal{L}_{dsd} = \frac{1}{N} \sum_{i,j} \text{sgn}(|\delta_x \text{sem}_{i,j}^l|) e^{-\|\frac{\delta_x d_{i,j}^l}{d_{i,j}^l}\|} + \text{sgn}(|\delta_y \text{sem}_{i,j}^l|) e^{-\|\frac{\delta_y d_{i,j}^l}{d_{i,j}^l}\|} \quad (7)$$

where *sem* denotes the ground truth semantic map and *d* the predicted disparity map. Differently from the smoothness term \mathcal{L}_{ds}^l in the disparity domain, the novel \mathcal{L}_{dsd} term detects discontinuities between semantic labels encoded by the sign of the absolute value of the gradients in the semantic map. The idea behind this loss is that there should be a gradient peak between adjacent pixels belonging to different classes. Nevertheless, we do not care about its magnitude since the numeric labels do not have any mathematical meaning.

4 Experimental results

In this section, we compare the performance of our semi-supervised joint depth estimation and semantic segmentation paradigm with respect to the proposal by Godard *et al.*[1], which represents nowadays the undisputed state-of-the-art for unsupervised monocular depth estimation. As discussed in Sec. 3, our method as well as the baseline used in our experiments, *i.e.* [1], require rectified stereo pairs at training time. Suitable datasets for this purpose are thus CityScapes [41] and KITTI [45], which provide a large number of training samples, *i.e.* about 23k and 29k rectified stereo pairs respectively. However, our method requires also pixel-wise groundtruth semantic labels at training time, which limits the actual amount of training samples available for our experiments. In particular, CityScapes includes about 3k finely annotated images, while the KITTI 2015 benchmark recently made available pixel-wise semantic groundtruths for about 200 images [6]. Therefore, to carry out a fair evaluation of the actual contribution provided by semantic information in the depth-from-mono task to the baseline fully unsupervised approach, we trained both methods based on the reduced datasets featuring stereo pairs alongside with semantically annotated left frames.

4.1 Implementation details

Our proposal has been implemented in Tensorflow, starting from the source code made available by the authors of [1]. We adhere to the original training protocol by Godard *et al.*, scheduling 50 epochs on the CityScapes dataset and 50 further on the KITTI 2015 images. For quantitative evaluation, we split the

KITTI 2015 dataset into train and test sets, providing more details in the next section. We train on 256×512 images using a batch dimension of 2, we set the previously introduced hyper-parameters as follows: $\alpha_d = 1$, $\alpha_s = 0.1$, $\alpha_{dsd} = 0.1$, $\beta_{ap} = 1$, $\beta_{lr} = 1$, $\beta_{ds} = \frac{1}{r}$ (being r the down-sampling factor at that resolution) and $\gamma = 0.85$. Models are trained using Adam optimizer [46], with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$. The initial learning rate is set to 10^{-4} , halved after 30 and 40 epochs. We perform data augmentation on input RGB images, in particular random gamma, brightness and color shifts sampled within the ranges $[0.8, 1.2]$ for gamma, $[0.5, 2.0]$ for brightness, and $[0.8, 1.2]$ for each color channel separately. Moreover we flip images horizontally with a probability of 50%. If the flip occurs, the right image in the stereo pair becomes the new reference image and we do not provide supervision signals from semantics (as right semantic maps are not available in the datasets). We implemented our network with both VGG and ResNet50 encoders, as in [1]. The semantic decoder adds about 20.5M parameters, resulting in nearly 50 and 79 million parameters for the two models (31 and 59, respectively, for [1]).

4.2 Monocular depth estimation: evaluation on KITTI 2015

We quantitatively assess the effectiveness of our proposal on the KITTI 2015 training dataset for stereo [3]. It provides 200 synchronized pairs of images together with groundtruth disparity and semantic maps [6]. As already mentioned, to carry out a fair comparison between our approach and [1], we can use only these samples and thus the numerical results reported in our paper cannot be compared directly with those in [1]. Then, we randomly split the 200 pairs from KITTI into 160 training samples and 40 samples used only for evaluation¹. We measure the accuracy of the predicted depth maps after training for 50 epochs on CityScapes and then fine-tuning for 50 more epochs on the samples selected from KITTI.

Table 1 reports quantitative results using VGG or ResNet50 as backbone encoder. Each model, one per row in the table, is trained with four different strategies:

- \mathcal{L}_d uses only the depth term as loss (*i.e.*, equivalently to the baseline approach by Godard *et al.*[1]).
- $\mathcal{L}_d + \mathcal{L}_s$ adds the semantic term to the depth term.
- $\mathcal{L}_d + \mathcal{L}_s + \mathcal{L}_{dsd}$ minimizes our proposed total loss function (Equation 1).
- $\mathcal{L}_d + \mathcal{L}_{dsd}$ minimizes only the losses dealing with the depth decoder.

The table provides results yielded by the four considered networks according to standard performance evaluation metrics [1] computed between estimated depth d and groundtruth D .

¹ The testing samples, belonging to the KITTI 2015 dataset, are: 000001, 000003, 000004, 000019, 000032, 000033, 000035, 000038, 000039, 000042, 000048, 000064, 000067, 000072, 000087, 000089, 000093, 000095, 000105, 000106, 000111, 000116, 000119, 000123, 000125, 000127, 000128, 000129, 000134, 000138, 000150, 000160, 000161, 000167, 000174, 000175, 000178, 000184, 000185 and 000193.

		Lower is better				Higher is better		
	Encoder pp	Abs Rel	Sq Rel	RMSE	RMSE log	δ_1	δ_2	δ_3
\mathcal{L}_d [1]	VGG	0.160	2.707	7.220	0.239	0.837	0.928	0.966
$\mathcal{L}_d+\mathcal{L}_s$	VGG	0.155	2.511	6.968	0.234	0.841	0.931	0.968
$\mathcal{L}_d+\mathcal{L}_s+\mathcal{L}_{dsd}$	VGG	0.154	2.453	6.949	0.235	0.844	0.931	0.967
$\mathcal{L}_d+\mathcal{L}_{dsd}$	VGG	0.161	2.758	7.128	0.240	0.841	0.928	0.964
\mathcal{L}_d [1]	VGG ✓	0.149	2.203	6.582	0.223	0.844	0.936	0.972
$\mathcal{L}_d+\mathcal{L}_s$	VGG ✓	0.147	2.229	6.583	0.223	0.847	0.938	0.972
$\mathcal{L}_d+\mathcal{L}_s+\mathcal{L}_{dsd}$	VGG ✓	0.145	2.040	6.362	0.221	0.849	0.938	0.971
$\mathcal{L}_d+\mathcal{L}_{dsd}$	VGG ✓	0.150	2.278	6.539	0.225	0.843	0.934	0.970
\mathcal{L}_d [1]	ResNet	0.159	2.411	6.822	0.239	0.830	0.930	0.967
$\mathcal{L}_d+\mathcal{L}_s$	ResNet	0.152	2.385	6.775	0.231	0.843	0.934	0.970
$\mathcal{L}_d+\mathcal{L}_s+\mathcal{L}_{dsd}$	ResNet	0.143	2.161	6.526	0.222	0.850	0.939	0.972
$\mathcal{L}_d+\mathcal{L}_{dsd}$	ResNet	0.155	2.282	6.658	0.232	0.840	0.932	0.968
\mathcal{L}_d [1]	ResNet ✓	0.148	2.104	6.439	0.224	0.839	0.936	0.972
$\mathcal{L}_d+\mathcal{L}_s$	ResNet ✓	0.144	2.050	6.351	0.220	0.849	0.938	0.972
$\mathcal{L}_d+\mathcal{L}_s+\mathcal{L}_{dsd}$	ResNet ✓	0.136	1.872	6.127	0.210	0.854	0.945	0.976
$\mathcal{L}_d+\mathcal{L}_{dsd}$	ResNet ✓	0.144	1.973	6.199	0.217	0.849	0.940	0.975

Table 1. Ablation experiments on KITTI 2015 evaluation split, using different configurations of losses, encoders and post-processing (pp). Best setup highlighted in bold for each configuration.

This ablation highlights, for both type of encoder, how introducing the second decoder trained to infer semantic segmentation maps significantly improves depth prediction according to all performance metrics. Moreover, adding the depth-semantic-discontinuity term, \mathcal{L}_{dsd} , leads in most cases to further improvements. On the other hand, minimizing \mathcal{L}_d and \mathcal{L}_{dsd} alone leads to inferior performance compared to the baseline method. We obtain the best configuration according to all metrics using ResNet50 when both \mathcal{L}_s and \mathcal{L}_{dsd} are minimized alongside with the depth term \mathcal{L}_d .

Moreover, we also evaluated the output obtained by all models after performing the post-processing step proposed by [1], that consists in forwarding both the input image I and its horizontally flipped counterpart \hat{I} . This produces two depth maps d_I and $d_{\hat{I}}$, the latter is flipped back obtaining $\hat{d}_{\hat{I}}$ and averaged with the former. We can notice that the previous trend is confirmed. In particular, the full loss $\mathcal{L}_d + \mathcal{L}_s + \mathcal{L}_{dsd}$ leads to the best result on most scores. Furthermore, including the post-processing step allows the VGG model trained with our full loss to outperform the baseline ResNet50 architecture supervised by traditional depth losses only. This fact can be noticed in Table 1 comparing row 7 with row 13, observing that the former leads to better results except for δ_3 metric.

To further prove the effectiveness of our proposed method we compare it with other self-supervised approach as [21],[4],[19]. Thus, we have ran experiments with the source code available from [21],[4],[19] using the same testing data as for [1] and our method. Table 2 shows the outcome of this evaluation. We point out that we used the weights made available by the authors of [21],[4],[19],

	Lower is better				Higher is better		
	Abs Rel	Sq Rel	RMSE	RMSE log	δ_1	δ_2	δ_3
Zhou et al. [4]	0.286	7.009	8.377	0.320	0.691	0.854	0.929
Mahjourian et al. [19]	0.235	2.857	7.202	0.302	0.710	0.866	0.935
Yin et al. [21]	0.236	3.345	7.132	0.279	0.714	0.903	0.950
Godard et al. [1]	0.159	2.411	6.822	0.239	0.830	0.930	0.967
Ours	0.143	2.161	6.526	0.222	0.850	0.939	0.972

Table 2. Comparison with other self supervised method on KITTI 2015 evaluation split. Both [1] and our method use ResNet50 encoder.

trained on a much larger amount of data (i.e., the entire Cityscapes and KITTI sequences, some of them overlapping with the testing split as well) w.r.t. the much lower supervision provided to our network. Despite this fact, monocular supervised works [21],[4],[19] perform poorly compared to both [1] and our approach, confirming our semi-supervised framework to outperform them as well. We also point out that our test split relies on high-quality ground-truth labels for evaluation, available from KITTI 2015 stereo dataset, while the Eigen split used to validate [21],[4],[19] provides much worse quality depth measurements, as also argued by the authors of [1].

As our final test we also compare our method with the recent multi-task learning approach by Kendall et al. [39]. Differently from our approach, they jointly learn depth, semantic and instance segmentation in fully supervised manner. They run experiments Tiny Cityscapes, a split obtained by resizing the validation set of Cityscapes to 128 256 resolution. To compare our results to theirs we have taken our ResNet50 model trained on Cityscapes and validated it following the same protocol. Their depth-only model (trained supervised) achieves 0.640 inverse mean depth error, dropping to 0.522 when trained to tackle semantic and instance segmentation as well. Our ResNet50 network (trained unsupervised) starts with 1.705 error for depth-only, dropping to 1.488. Thus, the two approaches achieve 22% and 15 % improvement respectively. We point out that, besides relying on supervised learning for depth, [39] exploits both semantic and instance segmentation, requiring additional manually annotated labels, while we only enforce our cross-domain depth-semantic-discontinuity loss.

Figure 4 depicts a qualitative comparison between the depth maps predicted by [1] and our semi-supervised framework. In the figure, from top to bottom, we consider images 000019 and 000095 belonging to our evaluation split. We can observe how explicitly learning the semantics of the scene helps to correct wrong depth estimations, especially on challenging objects. For example, we can notice how depth maps predicted by our frameworks provide better car shapes thanks to the contribution given by the semantic. This fact is particularly evident in correspondence of reflective or transparent surfaces like car windows as reported on image 000095. Moreover, the quality of thin structures like poles is improved as well, as clearly perceivable by looking at frame 000019.

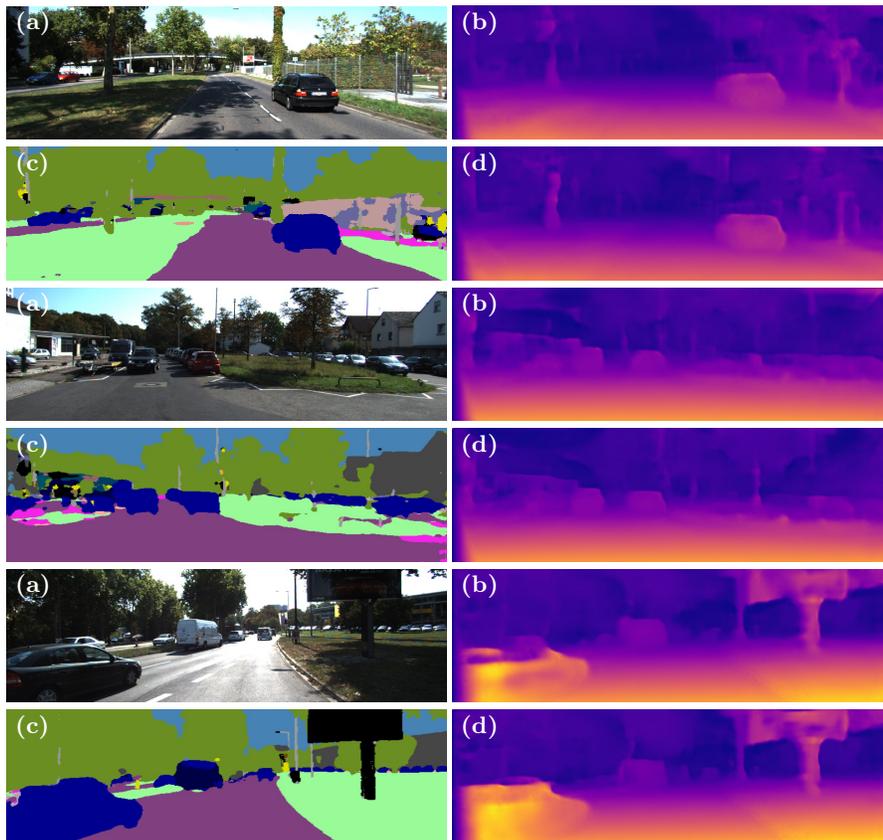


Fig. 4. Qualitative comparison between [1] and our proposal on KITTI 2015 evaluation split [3]. (a) Input image, (b) depth map by [1], (c) and (d) semantic and depth maps by our approach. Both models use Resnet50 as encoder. From top to bottom, results concerning images 000019, 000087 and 000095 belonging to our evaluation split.

4.3 Semantic segmentation: evaluation on KITTI 2015

Although our proposal is aimed at ameliorating depth prediction by learning richer features exploiting semantics, our network also delivers a semantic segmentation of the input image. To gather hints about the accuracy of this additional outcome of our network, we evaluated the semantic maps generated on the same KITTI evaluation split defined before. Differently from the monocular depth estimation task, results concerning semantic segmentation are quite far from the state-of-the-art. In particular, we obtain 88.51% and 88.19% per-pixel accuracy, respectively, with models based on VGG and ResNet50. We ascribe this to our architecture - inspired by [1] - being optimized for unsupervised depth prediction, whereas different design choices are often found in networks pursuing

semantic segmentation (i.e., atrous convolutions, SPP layers ...). We also found that training the basic encoder-decoder for semantic segmentation only yields to 86.72% and 88.18% per-pixel accuracy with VGG and ResNet50, respectively. Thus, while semantics helps depth prediction inasmuch as to outperform the state-of-the-art within the proposed framework, the converse requires further studies as the observed improvements are indeed quite minor. Therefore, we plan to investigate on how to design a network architecture and associated semi-supervised learning framework whereby the synergy between monocular depth prediction and semantic segmentation may be exploited in order to significantly improve accuracy in both tasks.

5 Conclusion and future work

We have proposed a deep learning architecture to improve unsupervised monocular depth estimation by leveraging on semantic information. We have shown how training our architecture end-end to infer semantics and depth jointly enables us to outperform the state-of-the-art approach for unsupervised monocular depth estimation [1]. Our single-encoder/dual-decoder architecture is trained in a semi-supervised manner, *i.e.* using ground truth labels only for the semantic segmentation task. Despite obtaining groundtruth labels for semantic is tedious and requires accurate and time-consuming manual annotation, it is still more feasible than depth labeling. In fact, this latter task requires expensive active sensors to be used at acquisition time and becomes almost unfeasible offline, on already captured frames. Thus, our method represents an attractive alternative to improve self-supervised training without adding more image samples. Future work will i) explore single camera sequences as supervision [4,21,20] and ii) dig into the semantic segmentation side of our framework, to reach top accuracy on this second task as well and to propose a state-of-the-art framework for joint depth and semantic estimation in a semi-supervised manner.

Acknowledgments We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X GPU used for this research

References

1. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: CVPR. (2017)
2. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: CVPR, IEEE (2012) 3354–3361
3. Menze, M., Geiger, A.: Object scene flow for autonomous vehicles. In: Conference on Computer Vision and Pattern Recognition (CVPR). (2015)
4. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: CVPR. Volume 2. (2017) 7
5. Zamir, A.R., Sax, A., Shen, W., Guibas, L., Malik, J., Savarese, S.: Taskonomy: Disentangling task transfer learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2018) 3712–3722

6. Alhaija, H.A., Mustikovela, S.K., Mescheder, L., Geiger, A., Rother, C.: Augmented reality meets deep learning for car instance segmentation in urban scenes. In: BMVC. (2017)
7. Liu, F., Shen, C., Lin, G., Reid, I.: Learning depth from single monocular images using deep convolutional neural fields. TPAMI **38** (2016) 2024–2039
8. 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. (2014) 2366–2374
9. Wang, X., Fouhey, D., Gupta, A.: Designing deep networks for surface normal estimation. In: CVPR. (2015) 539–547
10. Cao, Y., Wu, Z., Shen, C.: Estimating depth from monocular images as classification using deep fully convolutional residual networks. IEEE Transactions on Circuits and Systems for Video Technology (2017)
11. Saxena, A., Sun, M., Ng, A.Y.: Make3d: Learning 3d scene structure from a single still image. TPAMI **31** (2009) 824–840
12. Uhrig, J., Schneider, N., Schneider, L., Franke, U., Brox, T., Geiger, A.: Sparsity invariant cnns. In: 3DV. (2017)
13. Garg, R., BG, V.K., Carneiro, G., Reid, I.: Unsupervised cnn for single view depth estimation: Geometry to the rescue. In: European Conference on Computer Vision, Springer (2016) 740–756
14. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: Advances in neural information processing systems. (2015) 2017–2025
15. Poggi, M., Aleotti, F., Tosi, F., Mattoccia, S.: Towards real-time unsupervised monocular depth estimation on cpu. In: IEEE/JRS Conference on Intelligent Robots and Systems (IROS). (2018)
16. Poggi, M., Tosi, F., Mattoccia, S.: Learning monocular depth estimation with unsupervised trinocular assumptions. In: 6th International Conference on 3D Vision (3DV). (2018)
17. Aleotti, F., Tosi, F., Poggi, M., Mattoccia, S.: Generative adversarial networks for unsupervised monocular depth prediction. In: 15th European Conference on Computer Vision (ECCV) Workshops. (2018)
18. Kuznetsov, Y., Stuckler, J., Leibe, B.: Semi-supervised deep learning for monocular depth map prediction. In: CVPR. (2017)
19. Mahjourian, R., Wicke, M., Angelova, A.: Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018)
20. Wang, C., Buenaposada, J.M., Zhu, R., Lucey, S.: Learning depth from monocular videos using direct methods. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018)
21. Yin, Z., Shi, J.: Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018)
22. Shotton, J., Johnson, M., Cipolla, R.: Semantic texton forests for image categorization and segmentation. In: CVPR, IEEE (2008) 1–8
23. Fulkerson, B., Vedaldi, A., Soatto, S.: Class segmentation and object localization with superpixel neighborhoods. In: ICCV, IEEE (2009) 670–677
24. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR. (2015) 3431–3440
25. Eigen, D., Fergus, R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In: ICCV. (2015) 2650–2658

26. Chen, L.C., Yang, Y., Wang, J., Xu, W., Yuille, A.L.: Attention to scale: Scale-aware semantic image segmentation. In: CVPR. (2016) 3640–3649
27. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. TPAMI **40** (2018) 834–848
28. Liang-Chieh, C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.: Semantic image segmentation with deep convolutional nets and fully connected crfs. In: ICLR. (2015)
29. Badrinarayanan, V., Kendall, A., Cipolla, R.: Segnet: A deep convolutional encoder-decoder architecture for image segmentation. TPAMI (2017)
30. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI, Springer (2015) 234–241
31. Lin, G., Milan, A., Shen, C., Reid, I.: Refinenet: Multi-path refinement networks with identity mappings for high-resolution semantic segmentation. CVPR (2017)
32. Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, V., Su, Z., Du, D., Huang, C., Torr, P.H.: Conditional random fields as recurrent neural networks. In: ICCV. (2015) 1529–1537
33. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: CVPR. (2017) 2881–2890
34. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. CoRR, abs/1703.06211 **1** (2017) 3
35. Wang, P., Chen, P., Yuan, Y., Liu, D., Huang, Z., Hou, X., Cottrell, G.: Understanding convolution for semantic segmentation. WACV (2018)
36. Ladicky, L., Shi, J., Pollefeys, M.: Pulling things out of perspective. In: CVPR. (2014) 89–96
37. Mousavian, A., Pirsaviash, H., Košecá, J.: Joint semantic segmentation and depth estimation with deep convolutional networks. In: 3DV, IEEE (2016) 611–619
38. Wang, P., Shen, X., Lin, Z., Cohen, S., Price, B., Yuille, A.L.: Towards unified depth and semantic prediction from a single image. In: CVPR. (2015) 2800–2809
39. Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. (2018)
40. Laina, I., Rupperecht, C., Belagiannis, V., Tombari, F., Navab, N.: Deeper depth prediction with fully convolutional residual networks. In: 3DV. (2016) 239–248
41. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: CVPR. (2016) 3213–3223
42. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. (2015)
43. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. (2016) 770–778
44. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing **13** (2004) 600–612
45. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. International Journal of Robotics Research (IJRR) (2013)
46. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)