Monocular semantic occupancy grid mapping with convolutional variational encoder-decoder networks

Citation for published version (APA):

Lu, C., van de Molengraft, M. J. G., & Dubbelman, G. (2019). Monocular semantic occupancy grid mapping with convolutional variational encoder-decoder networks. *IEEE Robotics and Automation Letters*, 4(2), 445-452. [8603774]. <https://doi.org/10.1109/LRA.2019.2891028>

Document license:
TAVERNE

DOI:
[10.1109/LRA.2019.2891028](https://doi.org/10.1109/LRA.2019.2891028)

Document status and date:
Published: 01/04/2019

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Monocular Semantic Occupancy Grid Mapping With Convolutional Variational Encoder–Decoder Networks

Chenyang Lu^{1b}, Marinus Jacobus Gerardus van de Molengraft, and Gijs Dubbelman^{1b}

Abstract—In this letter, we research and evaluate end-to-end learning of monocular semantic-metric occupancy grid mapping from weak binocular ground truth. The network learns to predict four classes, as well as a camera to bird’s eye view mapping. At the core, it utilizes a variational encoder–decoder network that encodes the front-view visual information of the driving scene and subsequently decodes it into a two-dimensional top-view Cartesian coordinate system. The evaluations on Cityscapes show that the end-to-end learning of semantic-metric occupancy grids outperforms the deterministic mapping approach with flat-plane assumption by more than 12% mean intersection-over-union. Furthermore, we show that the variational sampling with a relatively small embedding vector brings robustness against vehicle dynamic perturbations, and generalizability for unseen KITTI data. Our network achieves real-time inference rates of approx. 35 Hz for an input image with a resolution of 256×512 pixels and an output map with 64×64 occupancy grid cells using a Titan V GPU.

Index Terms—Semantic scene understanding, object detection, segmentation and categorization, computer vision for transportation.

I. INTRODUCTION

ENVIRONMENT perception is a key task in mobile robot and intelligent vehicle operation. In the past decade, significant progress has been made, mainly due to increased computational power that has unlocked deep learning-based approaches for real-time usage, such as semantic segmentation [1]–[5] and object detection [6]–[9]. However, it can be argued that, for higher levels of robot and vehicle autonomy, perception and the incorporation of information derived from perception into a consistent world-model, is still a bottleneck. In this letter, we therefore research and evaluate the usage of semantic occupancy

Fig. 1. An illustration of the proposed variational encoder-decoder approach. From a single front-view RGB image, our system can predict a 2-D top-view semantic-metric occupancy grid map.

grid maps, as a means for end-to-end learning of monocular input data to form a world-model.

A world-model typically consists of multiple conceptual layers [10], e.g. layers of dynamic objects, permanent static objects, and movable static objects. Furthermore, one can distinguish layers that contain a priori knowledge from the environment, e.g. a global topological map, and layers that are estimated locally while the vehicle is traversing the environment. An occupancy grid map is particularly well-suited to represent the local free-space around the vehicle that is estimated in real-time from sensory input. This is also how we use it and we extend it with three different semantic sub-classes for free-space, namely road, sidewalk, and terrain, besides the usual non free-space class.

A particular branch of deep learning research focuses on convolutional neural networks (CNNs), which have significantly advanced computer vision in the past decade [11]–[13]. At a specific intermediate layer in CNNs, the feature map contains the semantic abstraction of the pixels as well as the inter-pixel 2-D spatial relations between them. The same inter-cell relations also hold for occupancy grids, thereby CNNs are potentially well-suited for end-to-end learning of occupancy grid maps with semantics from image data, which is proposed in this letter. We discuss the related work on occupancy grid maps and neural network approaches in more detail in Section II.

Our approach, which is detailed in Section III, contains the following contributions:

- To the best of our knowledge, we are the first to perform end-to-end learning on monocular imagery to produce a semantic-metric occupancy grid map and to achieve real-time inference rates.
- We show that this end-to-end monocular approach is intrinsically robust to pitch and roll perturbations,

Manuscript received September 10, 2018; accepted December 19, 2018. Date of publication January 7, 2019; date of current version January 16, 2019. This letter was recommended for publication by Associate Editor E. E. Aksoy and Editor T. Asfour upon evaluation of the reviewers’ comments. (Corresponding author: Chenyang Lu.)

C. Lu and G. Dubbelman are with the Mobile Perception Systems research cluster of the SPS/VCA Group, Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven 5600 MB, The Netherlands (e-mail: c.lu.2@tue.nl; g.dubbelman@tue.nl).

M. J. G. van de Molengraft is with the Control System Technology Group, Department of Mechanical Engineering, Eindhoven University of Technology, Eindhoven 5600 MB, The Netherlands (e-mail: m.j.g.v.d.molengraft@tue.nl).

This letter has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors. A video is included in the package, which visualizes that the system can work properly in real time.

Digital Object Identifier 10.1109/LRA.2019.2891028

and can generalize on unseen data from different Cameras.

- We show that, our approach can be trained from weak ground truth and is inherently robust to the sparseness of input data.

Considering the above, end-to-end learning of occupancy grids is a promising extension of, or even potentially can partially replace, traditional point-cloud processing techniques. Our approach is evaluated on Cityscapes [14] and KITTI [15] and the details on this are provided in Section IV after which our conclusions are put forward in Section V.

II. RELATED WORK

The occupancy grid map [16] is one of the most popular local metric map representations for mobile robots. Besides range sensors such as RaDAR and LiDAR, occupancy grid maps can also be generated from RGB-D cameras [17], stereo vision [18], and from fusion of multiple sensors [19]. However, the classical occupancy grid maps are without semantics, *i.e.* cells only have two possible states: occupied or not occupied.

More efficient and reliable navigation can be realized if semantics of the environment are utilized. Semantic segmentation is a potential approach to provide additional semantic scene understandings. Most semantic segmentation research has been carried out on RGB images with the goal to estimate a semantic class label for each individual pixel. For this particular task, it can be noted that deep learning methods are surpassing other classical methods in terms of both accuracy and efficiency. One state-of-the-art framework is the fully convolutional network (FCN) [1] that utilizes the convolutional feature extractor from other classification networks, such as VGG [12] or ResNet [13]. Another framework named SegNet [2], has a similar structure of auto-encoders. Further research shows that the segmentation quality can be enhanced by applying a conditional random field (CRF) as a post-processing step [4]. To integrate this in an end-to-end manner, CRFasRNN [20] is proposed to form a CRF as a recurrent neural network (RNN) that can be trained directly. Recent research has also performed semantic segmentation in an adversarial manner to produce improved result in terms of segmentation accuracy [21].

Image-based semantic segmentation methods are usually not directly compatible with vehicle mapping and planning systems. The reason is that in the mainstream state-of-the-art, metric mapping of the environment is deterministically performed in parallel with the semantic segmentation, by which the image-based semantic 3-D mapping is achieved. Sengupta *et al.* [22] project the image semantic labels into 3-D using stereo vision. Based on this, dense pairwise CRFs [23], [24] and higher order CRFs [25], [26] are proposed to optimize the 3-D labels. Recently, CNNs are integrated into the mapping pipeline [27]–[29] for a better image segmentation, and furthermore, even depth and pose estimation, which are used in deterministic metric mapping.

Instead of conducting metric mapping and semantic scene understanding separately, our long-term aim is to develop a holistic approach that can estimate metric, semantic, and

topological information simultaneously and in real-time. For this we take inspiration from recent work that has shown that deep learning approaches excel in estimating 3-D depth information from monocular [30]–[32] and binocular data [33], which means that the metric information can be learned from photometric data directly. This motivates us to research mapping the environment into semantic-metric occupancy grid maps directly from monocular input data in an efficient, end-to-end manner with deep neural networks.

III. SEMANTIC OCCUPANCY GRID MAPPING

In this section, we discuss the details of the aforementioned semantic-metric occupancy grid representation and the detailed structure and training of the proposed deep neural network.

A. Map Representation

We extend the classical definition of occupancy grid maps [16] to make the map representation contain semantic and metric information as well as suitable for modern deep neural networks.

Grid size and perceiving distance: Sensors mounted on autonomous vehicles such as cameras, RaDARs, and LiDARs usually have a fixed field of view (FOV), and the perception reliability decreases when the perceiving distance increases. To ensure each cell in the grid map, which is represented in 2-D vehicle coordinates, has a reliable status even at large distance, we set each grid map to contain 64×64 cells, with the size of each cell being 0.5×0.5 meters. As the region within 5 meters in front of the vehicle center is never visible, due to the camera's point of view, we apply a 5-meter offset in the grid map w.r.t. the vehicle center.

Semantic encoding: Each cell in the grid map is encoded with one of the following four semantic classes: *road*, *sidewalk*, *terrain*, and *non free-space* (including undetected grids that are behind the foreground objects and out of the camera's FOV). In this configuration, instead of a binary occupancy grid map (free-space or non free-space), the ground area in the map is extended with semantics, which potentially benefits the navigation of mobile robots and autonomous vehicles.

B. Network Structure and Training

In this letter, instead of implementing a deterministic point cloud based mapping algorithm, we propose an end-to-end learning approach. The proposed system is composed of two components: a low-level feature extractor and a modified version of variational auto-encoder (VAE) [34] network on top of the extracted feature map. As in our usage the input and output are not the same, as with a traditional VAE, we refer to our network as a variational encoder-decoder (VED) network. The input of this network is one front-view monocular RGB image, and the output is the top-view occupancy grid map in which each cell is assigned with a semantic class. The network is implemented in PyTorch [35] and Fig. 2 shows the detailed structure of the network.

Feature extractor: We use a modern CNN model, *e.g.* VGG-16 [12], pre-trained on ImageNet [36], to extract the low level

Fig. 2. The proposed network structure during the training phase. Every colored block represents a feature map and the arrows between them are neural network layers. Yellow indicates the encoder part while blue indicates the decoder. A pre-trained VGG-16 Net (without fully connected layers) is utilized for feature extraction on top of the input image. Legend: CP = VGG-like convolutional layers (2 layers) with kernel size 3 and 2×2 max pooling, FC = fully connected layer, SR = sample the latent vector with Normal distribution from z_{mean} and z_{stddev} and reshape, UC = one up-convolutional layer and VGG-like convolutional layers (2 layers) with kernel size 3, C = one VGG-like convolutional layer with kernel size 3. Every convolution layer uses batch normalization except the output layer.

features from the input monocular image. The receptive field of the VGG-16 network is 224×224 pixels. For reasons of efficiency, we use an input resolution of 256×512 pixels. As the receptive field is smaller than the input, the latent features in the output of the VGG-16 network are encoding the semantic information locally instead of on the entire image. This ensures that the spatial information is naturally preserved in the feature map, which is required for decoding the feature map into a top-down view.

Training with variational sampling: The variational auto-encoder [34] is originally proposed for learning variational Bayesian models in a neural network fashion. The learned coding vector contains the high-level representation of the input data, which is sampled from a standard normal distribution for later reconstruction. Recent research has shown that, when ground truth for voxel-based learning is incomplete, VAE can be used to produce a reconstruction output that surpasses the ground truth in term of completeness [37]. In our VED case, the ground truth is relatively imprecise (as will be explained in the following subsection), and we aim to mitigate this by using the variational sampling’s robustness to imperfect ground truth. In contrast to the VAE model in [37], several important modifications are made for our VED model: 1) taking the feature map from a modern feature extractor as input, and 2) training in *supervised* encoder-decoder manner instead of an auto-encoder manner.

We denote the encoding probabilistic model as $q_\phi(z|x)$, where $x = f_\gamma(i)$ is the high-level feature from the input image i and z is the latent embedding combined with spatial information and semantics. On top of the encoder, the probabilistic decoder $p_\theta(m|z)$ produces the 2-D grid semantic map m from the latent embedding z . The models f , q , p are organized as neural networks and their parameters γ , ϕ , θ can be learned simultaneously with end-to-end training. As we enforce the latent embedding z to obey the standard normal distribution, the latent loss $\mathcal{L}_{latent}$ is defined as Kullback-Leibler divergence between z and $\mathcal{N}(0, I)$. The mapping loss $\mathcal{L}_{mapping}$ is defined as

cross-entropy between the softmax output layer and the one-hot semantic coding of the ground truth. Therefore, the overall loss $\mathcal{L}$ for training is twofold, namely latent loss and mapping loss:

$$\mathcal{L} = \lambda_1 \cdot \mathcal{L}_{latent} + \lambda_2 \cdot \mathcal{L}_{mapping} \quad (1)$$

where λ_1 and λ_2 are the weights for the balancing of two objectives, which is set as 0.1 and 0.9 respectively in the experiments. We train the network using Adam [38] optimizer with learning rate = 0.0001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and mini-batch sizes of 8 for 60 epochs.

Weak ground truth for training: One major challenge of our approach is that there is no direct ground truth available, as the top-down view semantic occupancy grid representation is not provided in any publicly available dataset. However, one can utilize datasets that contain front-view image semantic annotations and 3-D information that can be pixel-wised registered as depth/disparity maps. To automatically generate the (weak) ground truth for training, we reconstruct the 3-D point cloud for each frame in vehicle coordinates from the corresponding depth/disparity map, given the intrinsic and extrinsic (the camera’s pose in the vehicle coordinates) camera calibration. For each frame of the generated point cloud, given the corresponding front-view image semantic ground truth annotation, a semantic label can be assigned to each 3-D point. Next, we project the 3-D points to the 2-D ground plane and subsequently fill the occupancy grid with pre-defined size. For each cell, a semantic label is assigned, based on the label statics of the cell’s points (majority vote).

The 3-D information registered for the pixels can be noisy (e.g. a disparity map estimated using a stereo matching method) or sparse (e.g. a depth map from LiDAR measurements). It can be argued that the automatically generated ground truth contains noise mainly from the imprecise depth/disparity map, e.g. grid cells can be missed on the road, due to the corresponding depth/disparity region is invalid. For this reason, we refer to the automatically generated ground truth as *weak ground truth*. Some automatically generated weak ground truth examples can be seen in Fig. 3(c). Note that only for evaluation we have manually annotated 70 top-view grid maps, which is too few for end-to-end training. The ability to train from weak ground truth is an important feature of our neural network based approach.

IV. EXPERIMENTS

We conduct the following experiments to demonstrate our approach and to compare its accuracy and robustness with three baseline approaches being: 1) a traditional monocular method that relies on a flat-plane assumption, 2) a traditional binocular approach, and 3) a CNN based approach which is commonly used for segmentation tasks:

- *Quantitative evaluation:* In this experiment, we use the Cityscapes [14] dataset to measure performances employing metrics from semantic image segmentation.
- *Input disturbance invariance:* We simulate roll and pitch movements of the camera, to investigate the invariance of our approach to such perturbations.

Fig. 3. Some visualized mapping examples on the test set with different methods. (a) is the input image from the left RGB camera mounted on the vehicle. (b) is the satellite image corresponding to the RGB image based on the GPS signal from Cityscapes for a better understanding of our work. The region in the white rectangle is focused in the mapping task. (c) is the weak ground truth map with ground truth semantic segmentation and semi-global matching disparity. (d) is the manually improved ground truth map based on the weak ground truth. (e)–(h) are the mapping result with predicted semantic segmentation and the same disparity, flat-plane assumption geometric transformation, SegNet [2] baseline method, and our proposed VED method. Grids with black mask are ignored in evaluation as they are out of the camera’s FOV or with ignored semantic labels.

- *Generalizability to unseen scenarios:* The unseen KITTI [15] dataset containing a semantic domain gap and different camera parameters is used to evaluate the generalizability of the proposed approach.
- *Ablation study:* We compare the results of the proposed network trained with and without variational sampling to better motivate the usage of the variational sampling.
- *Mapping quality invariance w.r.t. resolutions:* We generate maps using two deterministic baseline methods and the proposed neural network method in different resolutions and investigate the advances of the neural network based approach.
- *Semantic latent embedding:* In this small experiment, we research what high-level information is encoded in the latent embedding of our VED approach.

A. Dataset and Ground Truth

We use the Cityscapes dataset [14] for ground truth generation and experiments, as it provides stereo images with disparity and fine semantic annotations for each pixel. We use the 2975 images in the *training set* for training, and the 500 images in the *validation set* for evaluation and comparison. In our experiments, all the images are resized from 1024×2048 to 256×512 for efficiency. The disparity maps provided from Cityscapes with semi-global matching (SGM) method [39] for weak ground truth generation. As discussed in Section III-B, the automatically generated ground truth contains noise. To perform a valid quantitative evaluation, we also manually improved and annotated 70 top-view grid maps in the validation set, based on the visual cue in the corresponding front-view image, which are referred as *ground truth* and visualized in Fig. 3(d). Furthermore, we use KITTI [15] semantic dataset, which contains 200 images with publicly available semantic annotation, depth maps and camera parameters, to verify the generalizability of different methods.

The same procedure is applied for weak ground truth generation and another 70 KITTI samples are manually improved as ground truth for evaluation.

B. Baseline Methods

Other than our VED approach, there are multiple methods available for mapping sensory data to the proposed map representation. In this letter, we compare our approach with two canonical point cloud based methods and one CNN based SegNet [2] method:

1) *Monocular Mapping With Flat-Plane Assumption (Flat-Plane Assumption):* Our first baseline method does not use direct 3-D information, but instead uses a flat-plane assumption to map the output of the semantic segmentation, obtained with a VGG-16 based FCN [1] on front-view images, to a top-down view. More precisely, in this method, we assume each pixel in the RGB image which is predicted as one of the ground-like classes (*road*, *sidewalk*, and *terrain*) is located on the ground in 3-D (height being 0 in vehicle coordinates if the vehicle local dynamics is compensated by an IMU *w.r.t.* the static world). With this assumption, the free-space pixels’ position in 2-D coordinates can be deterministically solved and mapped using the same method as ground truth generation. However, even given the perfect IMU data, the flat-plane assumption cannot handle scenarios where there is a slope in front of the vehicle, which leads to a map with large offsets. Furthermore, the IMU data can be noisy at each data point, which might further hurt the mapping results, due to the vulnerability of flat-plane assumption. Considering the above reasons and that the data is collected in a relative steady situation, we only use the camera calibration (intrinsic and extrinsic, *i.e.* the camera’s pose in vehicle coordinates) instead of IMU data in the experiments. The aim is to outperform this straightforward monocular baseline using

TABLE I
QUANTIFIED PERFORMANCE FOR DIFFERENT MAPPING METHODS

method	weak ground truth		ground truth	
	mean accuracy	mean IoU	mean accuracy	mean IoU
with disparity	91.8	80.0	85.7	70.6
flat-plane assumption	69.9	47.2	70.1	47.2
SegNet [2]	73.7	61.1	75.9	61.5
VED (ours)	73.8	59.5	74.9	59.6

our end-to-end learning approach. This method is referred to as *flat-plane assumption* in all figures and tables.

2) *Binocular Mapping (With Disparity)*: To provide an upper bound on what we realistically could achieve with our monocular approach, we also validate against a binocular approach. For this baseline, we use the same procedure as for generating the weak ground truth, but the key difference is that now the semantic information is estimated using a VGG-16 based FCN [1] instead of the labeled Cityscapes ground truth annotations. This baseline uses binocular image pairs to obtain the corresponding disparity maps for 3-D point cloud generation. In the implementation, the 3-D point clouds are obtained from the Cityscapes disparity maps with SGM method [39] and used to fill the occupancy grid. However, note that the disparity maps can also be obtained from other methods, such as stereo network-based approaches [33] and monocular network-based approaches [30]–[32]. This method is referred to as *with disparity* in all figures and tables.

3) *Mapping Using Canonical CNNs (SegNet)*: We also use a convolutional encoder-decoder based SegNet [2] to perform the same task. As the size of the input image and output are 256×512 and 64×64 , instead of using the original SegNet, we drop the last two decoder modules and applied an additional max-pooling layer with kernel size and stride 1×2 after the last up-pooling layer. We train the SegNet using the same optimizer settings with mini-batch sizes of 2 for 35 epochs. SegNet uses the same VGG-16 [12] as backbone, and the number of trainable parameters are similar between SegNet (≈ 28.9 M) and the proposed VED network (≈ 27.5 M), which constructs a fair comparison. Note that the SegNet has a significantly larger information bottleneck ($8 \times 16 \times 512$) than the proposed network (1×512). Hence, it can pass more information to the decoder, which is beneficial for map generation while introducing risks in terms of robustness and generalizability. These will be further discussed in the experimental results. This method is referred to as *SegNet* in all figures and tables.

C. Results

1) *Quantitative Evaluation*: As our target maps are organized in an image-like fashion, we evaluate the results in terms of mean accuracy and mean intersection-over-union (IoU), averaged over the test samples. The performances of the three mapping methods are provided in Table I. Note that in this letter, the grid cells out of the camera’s FOV are used in training

but ignored in evaluation and visualization with black mask as they are consistent and trivial for each frame. We report the metrics evaluated on both weak ground truth and manually improved ground truth. The performance of the binocular mapping method (*with disparity*) on weak ground truth is higher than that on manually improved ground truth by a large margin, while the other three methods remain at the same level. This is because the binocular mapping baseline uses exactly the same Cityscapes disparity maps as are also used for weak ground truth generation, which leads to the positive bias when evaluating on the weak ground truth. The aforementioned bias is removed in the metrics evaluated on the manually improved ground truth. In either ground truth setting, it can be seen that the binocular mapping method outperforms the other three monocular methods, as expected. Concerning the monocular methods, the results clearly show that two neural network based methods surpass the flat-plane assumption method by nearly 5% mean accuracy and 12% mean IoU. SegNet [2] provides slightly better performance than our VED network with a margin less than 2%. This is because that it has a significant larger ($8 \times 16 \times 512$ instead of 1×512) information bottleneck without prior distribution regularization, which brings small performance improvements but also introducing disadvantages in terms of robustness and generalizability. These disadvantages will be discussed in the following experiments. Given an input with resolution 256×512 , our method requires about 28 milliseconds and is thereby able to achieve frame-rates of approx. 35 Hz on a Nvidia Titan V GPU.

2) *Input Disturbance Invariance*: While driving, the camera will exhibit roll and pitch perturbations w.r.t. to a stand-still situation. If not accounted for, these perturbations significantly degrade the performance when using a monocular approach based on a flat-plane assumption. Clearly, IMUs can provide orientation information, but the measurement accuracy and time synchronization can be problematic. Ideally, one would want to make the mapping from image coordinates to top-view coordinates intrinsically invariant to such perturbations without using an IMU. We illustrate that our VED network exhibits this invariance. We simulate new input images in the cases of different common orientation disturbances in *pitch* (simulated with vertical pixel offsets) and *roll* (simulated with in-plane rotations around the imaging center) and feed them into different methods. Table II shows the metrics and Fig. 4 visualizes some examples with different orientation disturbances. It can be concluded that VED exhibits intrinsic levels of invariance w.r.t. to pitch and roll perturbations, compared to the other monocular baselines. This is mainly because VED network learns to extract the high level semantic-metric information into a low dimensional space, rather than deterministic mapping or direct large feature map passing. Furthermore, it is interesting to note that these results are obtained *without* data augmentation techniques during training that simulate pitch and roll perturbations, which would probably increase the invariance even further.

3) *Generalizability to Unseen Scenarios*: The neural network based methods, *i.e.*, our VED network and SegNet [2] baseline do not need the camera parameters during testing, which makes the task more challenging when the image is

TABLE II
ROBUSTNESS EVALUATION W.R.T. VEHICLE LOCAL DYNAMICS. THE NUMBERS IN THE BRACKETS INDICATE THE PERFORMANCE DOWNGRADE W.R.T. THE ORIGINAL PERFORMANCE WITHOUT PERTURBATION

	VED (ours)		SegNet [2]		flat-plane assumption	
	mean accuracy	mean IoU	mean accuracy	mean IoU	mean accuracy	mean IoU
no perturbation	74.9	59.6	75.9	61.5	70.1	47.2
$\pm 1.5^\circ$ pitch	72.0 (-2.9)	56.2 (-3.4)	69.5 (-6.4)	54.3 (-7.2)	53.4 (-16.7)	35.3 (-11.9)
$\pm 5^\circ$ roll	70.3 (-4.6)	55.1 (-4.5)	65.7 (-10.2)	52.0 (-9.5)	58.3 (-11.8)	41.5 (-5.7)

Fig. 4. Visualized comparison for different pitch and roll perturbations. We present two examples which are divided by the black horizontal line. For each example, the most left column shows the input RGB image and its corresponding (weak) ground truth. The other columns show the predictions of three monocular approaches in different perturbation settings.

TABLE III
QUANTIFIED PERFORMANCE FOR DIFFERENT MAPPING METHODS EVALUATED ON THE UNSEEN KITTI DATASET

Method	no perturbation		$\pm 1.5^\circ$ pitch	
	mean accuracy	mean IoU	mean accuracy	mean IoU
with disparity	71.1	55.9	71.1	55.9
flat-plane assum.	68.3	51.2	52.1(-16.2)	35.7(-15.5)
SegNet [2]	36.1	19.6	33.1(-3.0)	16.2(-3.4)
VED (ours)	58.2	42.4	56.1(-2.1)	40.1(-2.3)

from a different camera. To investigate this, we evaluate our VED network, which is trained on Cityscapes, on the unseen KITTI dataset, which has different camera parameters and semantic domain gaps. We compare the results with the flat-plane assumption (using updated KITTI camera parameters) and the SegNet baseline, see Table III. Obviously, the flat-plane assumption can perform deterministic mapping with the new camera parameters. For the neural network based approaches, in order to make the KITTI data compatible with the networks

trained on Cityscapes, we align the horizontal field of view (by cropping KITTI images) and align the vanishing point (by vertical image translation). The results, provided in Table III and Fig. 5, show that, although the VED network is trained on camera parameters from Cityscapes [14], the network can still work with some performance degradation, while the SegNet baseline fails. This degradation (nearly 10% compared to the flat-plane assumption) is expected, as the camera pose is significantly different between the two datasets. We also observe that without cropping and alignment of the images, the performance of the VED degrades by nearly 13% for both metrics, which means that these pre-processing steps are necessary when dealing with unseen data. Furthermore, when applying the pitch perturbations on the unseen KITTI dataset, our approach exhibits better performance than the flat-plane assumption, in both degradation and absolute values, which shows the robustness and generalizability of our approach even with changed camera settings and scene domain gaps.

4) *Ablation Study*: To better motivate the usage of variational sampling, we perform an ablation study to investigate the effectiveness of the variational sampling in our mapping task. We train the proposed VED network with only one modification:

Fig. 5. Some visualized mapping examples on the unseen KITTI [15] dataset using different methods. (a) is the manually improved ground truth map based on the weak ground truth. (b) is the deterministic mapping result with predicted semantic segmentation and ground truth depth. (c)–(e) is the mapping result using flat-plane assumption geometric transformation, SegNet [2] baseline method, and our proposed VED method, with and without pitch perturbations. The black FOV mask is the same as in the Cityscapes, as the input images are cropped for aligning the FOVs of two cameras.

TABLE IV
QUANTIFIED PERFORMANCE FOR THE PROPOSED VED APPROACH TRAINED WITH AND WITHOUT VARIATIONAL SAMPLING

Method	no perturbation		$\pm 1.5^\circ$ pitch	
	mean accuracy	mean IoU	mean accuracy	mean IoU
VED (ours)	74.9	59.6	72.0(-2.9)	56.2(-3.4)
VED w/o sampling	74.0	60.6	69.2(-4.8)	55.0(-5.6)

the embedding vector is directly passed from the encoder to the decoder, instead of randomly sampled and regularized based on the outputs of two fully-connected layers. Table IV, shows the performance of the VED network with and without the variational sampling. Two networks exhibit similar performance when no perturbation is applied. However, the usage of variational sampling improve the robustness against perturbations: the performance degradation with sampling is about 2% less than that without sampling.

5) *Mapping Quality Invariance*: In our experiments, the resolution of the map representation is set to be 64×64 pixels, while it can be extended to any other resolution, such as 128×128 pixels or even higher. With the output resolution increasing, the side effects will appear in point cloud based mapping approaches: the artifacts will exhibit because the points registered for the grid at far distance are insufficient for a reliable majority vote. In Fig. 6, we show some prediction examples using deterministic approaches and our VED approach with the map resolution being 128×128 pixels. It can be observed that at large distance, semantic information is lost in some grids with certain patterns in point cloud based methods, which degrades mapping quality, while the network based method will not exhibit this behavior. Our approach is intrinsically invariant to point cloud density as we extract high level semantic-metric information from images directly and achieve higher map resolution with up-convolution operations.

6) *Semantic Latent Embedding*: The latent representation in our proposed network is supposed to encode both high-level semantic and spatial information into an embedding vector with

Fig. 6. Examples of weak ground truth map and predictions from different mapping approaches in high resolution (128×128 pixels) setting. Both point cloud based approaches produce maps with certain artifact patterns, while the VED network produces maps with acceptable quality.

Fig. 7. PCA perturbation analysis. The numbers are indicating the perturbation values applied on the first and second principle axis.

512 dimensions. As our system handles complicated data in real urban environments, some attributes in the vector might be highly correlated, which makes it difficult to perform direct attribute analysis. To analyze the effectiveness of our encoding and decoding system separately, we conduct the principal component analysis (PCA) on 500 test images' embedding vectors. We apply perturbations on the first and second principal axes and visualize the modified map predictions, which are illustrated in Fig. 7. It can be noted that the first principal axis is mainly

encoding the width of the drivable space in front of the vehicle, and the second one is encoding the size of the non free-space area near the center of FOV. This shows that our network indeed learns to encode semantic and spatial understanding from monocular image into a latent embedding vector. As mentioned earlier this spatial understanding provides the network with robustness to pitch and roll perturbations as well allows for up-sampling the resolution of the occupancy grid map.

V. CONCLUSION

In this letter, we proposed a novel real-time neural network based end-to-end mapping system, which requires a single front-view image from a monocular camera and from it estimates a top-view semantic-metric occupancy grid map. It is shown that our VED approach outperforms a monocular system using a flat-plane assumption, and exhibits better robustness and generalizability, due to its variational sampling over a relatively low dimensional embedding space, when compared to the canonical neural network baseline. We have verified that the network can learn semantics as well as metric spatial information, by investigating the latent embedding that it uses. Although further research is certainly required to bring our VED approach to a level at which it can be applied, our work demonstrates that, occupancy grids, although already several decades old, are still a very relevant and powerful representation and that they link very well with state-of-the-art methods from deep learning, which can enhance or even partially replace traditional point cloud processing techniques. In future work, we aim to further leverage on deep learning and predict the road layout beyond the camera's FOV.

REFERENCES

- [1] E. Shelhamer, J. Long, and T. Darrell, "Fully convolutional networks for semantic segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 4, pp. 640–651, Apr. 2017.
- [2] V. Badrinarayanan, A. Kendall, and R. Cipolla, "SegNet: A deep convolutional encoder-decoder architecture for image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 12, pp. 2481–2495, Dec. 2017.
- [3] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 6230–6239.
- [4] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 4, pp. 834–848, Apr. 2018.
- [5] P. Meletis and G. Dubbelman, "Training of convolutional networks on multiple heterogeneous datasets for street scene semantic segmentation," in *Proc. IEEE Intell. Veh. Symp. IV*, 2018, pp. 1045–1050.
- [6] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 580–587.
- [7] R. Girshick, "Fast R-CNN," in *Proc. Int. Conf. Comput. Vis.*, 2015, pp. 1440–1448.
- [8] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, Jun. 2017.
- [9] K. He, G. Gkioxari, P. Dollar, and R. Girshick, "Mask R-CNN," in *Proc. Int. Conf. Comput. Vis.*, 2017, pp. 2980–2988.
- [10] A. Furda and L. Vlacic, "An object-oriented design of a World Model for autonomous city vehicles," in *Proc. IEEE Intell. Veh. Symp. IV*, 2010, pp. 1054–1059.
- [11] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [12] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1, 2014.
- [13] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [14] M. Cordts *et al.*, "The cityscapes dataset for semantic urban scene understanding," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 3213–3223.
- [15] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *Int. J. Robot. Res.*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [16] A. Elfes, "Occupancy grids: A stochastic spatial representation for active robot perception," in *Proc. 6th Conf. Uncertainty Artif. Intell.*, 1990, pp. 136–146.
- [17] M. Himstedt and E. Maehle, "Online semantic mapping of logistic environments using RGB-D cameras," *Int. J. Adv. Robot. Syst.*, vol. 14, no. 4, pp. 1–13, 2017.
- [18] Y. Li and Y. Ruichek, "Occupancy grid mapping in urban environments from a moving on-board stereo-vision system," *Sensors*, vol. 14, no. 6, pp. 10454–10478, 2014.
- [19] S.-I. Oh and H.-B. Kang, "Fast occupancy grid filtering using grid cell clusters from LIDAR and stereo vision sensor data," *IEEE Sensors J.*, vol. 16, no. 19, pp. 7258–7266, Oct. 2016.
- [20] S. Zheng *et al.*, "Conditional random fields as recurrent neural networks," in *Proc. Int. Conf. Comput. Vis.*, 2015, pp. 1529–1537.
- [21] P. Luc, C. Couprie, S. Chintala, and J. Verbeek, "Semantic segmentation using adversarial networks," in *Proc. Adv. Neural Inf. Process. Syst. Workshop*, 2016.
- [22] S. Sengupta, E. Greveson, A. Shahrokni, and P. H. S. Torr, "Urban 3D semantic modelling using stereo vision," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2013, pp. 580–585.
- [23] A. Hermans, G. Floros, and B. Leibe, "Dense 3D semantic mapping of indoor scenes from RGB-D images," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2014, pp. 2631–2638.
- [24] V. Vineet *et al.*, "Incremental dense semantic stereo fusion for large-scale semantic scene reconstruction," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2015, pp. 75–82.
- [25] A. Kundu, Y. Li, F. Dellaert, F. Li, and J. M. Rehg, "Joint semantic segmentation and 3D reconstruction from monocular video," in *Proc. Euro. Conf. Comput. Vis.*, 2014, pp. 703–718.
- [26] Z. Zhao and X. Chen, "Building 3D semantic maps for mobile robots using RGB-D camera," *Intell. Service Robot.*, vol. 9, no. 4, pp. 297–309, 2016.
- [27] S. Yang, Y. Huang, and S. Scherer, "Semantic 3D occupancy mapping through efficient high order CRFs," in *Proc. Int. Conf. Intell. Robots Syst.*, 2017, pp. 590–597.
- [28] J. McCormac, A. Handa, A. Davison, and S. Leutenegger, "SemanticFusion: Dense 3D semantic mapping with convolutional neural networks," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 4628–4635.
- [29] K. Tateno, F. Tombari, I. Laina, and N. Navab, "CNN-SLAM: Real-time dense monocular SLAM with learned depth prediction," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 6565–6574.
- [30] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2366–2374.
- [31] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised learning of depth and ego-motion from video," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 6612–6619.
- [32] C. Godard, O. M. Aodha, and G. J. Brostow, "Unsupervised monocular depth estimation with left-right consistency," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 6602–6611.
- [33] N. Mayer *et al.*, "A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 4040–4048.
- [34] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, arXiv:1312.6114.
- [35] A. Paszke *et al.*, "Automatic differentiation in PyTorch," in *Proc. Adv. Neural Inf. Process. Syst. Workshop*, 2017.
- [36] O. Russakovsky *et al.*, "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, no. 3, pp. 211–252, 2015.
- [37] J. L. Schönberger, M. Pollefeys, A. Geiger, and T. Sattler, "Semantic visual localization," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 6896–6906.
- [38] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, arXiv:1412.6980.
- [39] H. Hirschmüller, "Stereo processing by semiglobal matching and mutual information," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 2, pp. 328–341, Feb. 2008.