

GroundNet: Monocular Ground Plane Normal Estimation with Geometric Consistency

Yunze Man^{1*}, Xinshuo Weng², Xi Li¹, Kris Kitani²
 Dept. of Computer Science and Technology, Zhejiang University, Zhejiang, China¹
 Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA²
 {yzman,xilizju}@zju.edu.cn, {xinshuow, kkitani}@cs.cmu.edu




Figure 1: Results of our approach: (a) input image; (b)(c) geometrically-consistent depth and surface normal estimation; (d) ground segmentation; (e) ground plane normal estimation; (f) computed horizon line from the estimated ground plane normal.

ABSTRACT

We focus on estimating the 3D orientation of the ground plane from a single image. We formulate the problem as an inter-mingled multi-task prediction problem by jointly optimizing for pixel-wise surface normal direction, ground plane segmentation, and depth estimates. Specifically, our proposed model, GroundNet, first estimates the depth and surface normal in two separate streams, from which two ground plane normals are then computed deterministically. To leverage the geometric correlation between depth and normal, we propose to add a consistency loss on top of the computed ground plane normals. In addition, a ground segmentation stream is used to isolate the ground regions so that we can selectively back-propagate parameter updates through only the ground regions in the image. Our method achieves the top-ranked performance on ground plane normal estimation and horizon line detection on the real-world outdoor datasets of ApolloScape and KITTI, improving the performance of previous art by up to 17.7% relatively.

CCS CONCEPTS

• **Computing methodologies** → **Scene Understanding**: *Camera Calibration; Neural Networks.*

* Work done when Yunze was an intern at CMU.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '19, October 21–25, 2019, Nice, France

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6889-6/19/10...\$15.00

<https://doi.org/10.1145/3343031.3351068>

KEYWORDS

ground plane normal estimation; horizon line detection; monocular vision; self-supervised learning via consistency; multitask learning

ACM Reference Format:

Yunze Man^{1*}, Xinshuo Weng², Xi Li¹, Kris Kitani². 2019. GroundNet: Monocular Ground Plane Normal Estimation with Geometric Consistency. In *Proceedings of the 27th ACM International Conference on Multimedia (MM '19)*, October 21–25, 2019, Nice, France. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3343031.3351068>

1 INTRODUCTION

Estimating the 3D orientation of the ground plane is an important pre-processing step for ground robots [27], wearable camera systems [52] and autonomous driving [46]. An accurate estimate of the ground plane can serve as an important prior for many perception and planning tasks, e.g., 3D object detection [7, 23, 51], 3D object tracking [21, 50], 3D semantic segmentation [1], image data synthesis [24], free space estimation [34], camera placement estimation [59], 3D reconstruction [18], and scene analysis [16, 17]. While many sensors can be used to directly estimate the ground plane (e.g., depth camera, stereo camera, laser scanner), we are primarily interested in innovating ground plane estimation algorithms for mobile platforms that are equipped with only a single RGB camera.

Perhaps the most classical approach to ground plane estimation makes use of multi-view geometry or motion cues to first triangulate points in 3D or directly obtain the 3D point cloud using depth sensors like LIDAR. Then a large plane is fitted to the 3D points using a robust model fitting algorithm like RANSAC [28]. When only a single image is available, parallel lines detected on the ground plane can be used to estimate vanishing points and the horizon line [15]. While these geometry-based approaches are

exact with noiseless input, their performance is highly dependent on the reliability of low-level computer vision algorithms to extract corner points or line segments. Not to mention that outdoor scenes typically lacks 3D architectural prior (i.e., the vertical relationships between walls, floor, and ceiling), making existing geometry based method insufficient for the estimation.

An alternative approach that can be used to estimate the ground plane is the use of monocular surface normal estimation [2]. The basic idea of this approach is to use small image patches to estimate the distribution of surface normals using local visual information. Most prior work in this area are pixel-wise normal estimation models. They are tailored to indoor scenes and are typically not designed to deal with the outdoor scenes which involve heavy noise, including diverse objects (e.g., vehicles, pedestrians) and natural ‘stuff’ (e.g., vegetation, sky) with dynamic motion. For this reason, the estimated normals often suffer from heavy local noise.

To leverage the advantages of both approaches, we propose the GroundNet, which computes the ground plane normal using the depth and surface normal cues in two separate streams: (1) in the surface normal estimation stream, we obtain a pixel-wise surface normal map of the ground region and then take the average over pixels as the ground plane normal; (2) in the depth estimation stream, we lift the ground regions in the input image to a point cloud using the estimated depth. Then a ground plane is fitted to the point cloud using the differentiable RANSAC (DSAC) [4].

In order to remove the effect caused by diverse objects on or nearby the ground in the outdoor scenes, a ground segmentation stream is used to isolate the ground regions. We then selectively back-propagate parameter updates only through the ground regions in the image. We share the backbone encoder network for all three networks: ground region segmentation, depth estimation and pixel-wise normal estimation.

In addition, our insight is that the estimated depth and surface normal should be geometrically consistent, meaning that the computed ground plane normals from the depth and surface normal streams should be nearly the same. However, this is not necessarily true in reality. To introduce this geometric consistency between the depth and surface normal streams, we propose to add a consistency loss to minimize the angular difference between the computed ground plane normals in two streams. This supervision signal is flown back to the surface normal and depth estimation streams such that our GroundNet can implicitly learn to predict the geometrically-consistent depth and surface normal in two streams. We argue that this consistency constraint can resolve the geometrical ambiguity when estimating the depth or normal alone from a single image.

The proposed GroundNet is shown in Figure 2. To show the effectiveness of our method, we evaluate on the real-world outdoor datasets of ApolloScape and KITTI, which involves heavy noise from objects. As the images in these two datasets lack of variety of the ground plane normal, we augment the images by adding random roll and pitch so that estimating the ground plane normal becomes challenging. Our method achieves the state-of-the-art performance on the datasets either with or without the augmentation.

Our contributions are summarized as follows: (1) we propose a novel network for end-to-end outdoor ground plane normal estimation; (2) we introduce a ground segmentation stream to isolate the ground regions so as to avoid the noise caused by irrelevant objects

in outside scenarios; (3) we propose to learn the essential correlation between depth and normal by geometric consistency loss, allowing better learning of 3D information through multi-modality interaction and refinement. (4) though extensive evaluation on the real-world outdoor datasets of KITTI and ApolloScape, the proposed GroundNet achieves the state-of-the-art performance on the task of ground plane normal estimation and horizon line detection, improving the previous state-of-the-art by up to 17.7% relatively.

2 RELATED WORK

Ground Plane Normal Estimation. Existing methods for ground plane estimation can be classified into geometry-based methods and learning-based methods. Geometry-based methods often extract the 3D scene structure (e.g., using multi-view cues, motion cues or depth sensors) and then the ground plane is fitted to the 3D points using a robust model fitting algorithm like RANSAC. [28] identifies the ground using the 3D point cloud from LIDAR. [31] obtains the video frame rate depth maps from the time-of-flight (TOF) cameras and exploits 4D spatiotemporal RANSAC for ground plane estimation. [39] generates the 3D point cloud under a stereo setup and then estimates the ground plane by the disparity. Assuming the scene is static, simultaneous localizing and mapping (SLAM) and structure from motion (SFM) approaches can also be used for extracting the 3D scene structure [29, 33, 38, 41, 57, 63], making ground plane estimation possible.

When only a single image is available, parallel lines detected on the ground plane can be used to estimate vanishing points and the horizon line [15]. Also, [30] shows how grouping detected line segments into quadrilaterals can be used to find orthogonal planes. However, these geometry-based methods are highly dependent on the reliability of low-level computer vision algorithms (e.g., planar homography estimation, line segment detection), especially when given only a single image.

The secondary category methods focus on applying machine learning technique to estimate the ground plane normal either directly or indirectly from related tasks. There are only a few prior works are direct methods. [13, 14] learn a classifier to classify local planar image patches and their orientations first. Then a Markov random field (MRF) [54] model is learned to segment the image into dominant plane segments. [32] achieves the ground plane recognition by learning the lighting-invariant texture feature using regularized logistic regression model. However, these methods make use of shallow learning model and do not benefit from the recent significant progress in deep learning. Also, they do not model the geometric relationship existing in the image. To the best of our knowledge, GroundNet is the first work of direct ground plane estimation which leverages the strong capacity of the deep neural network and the geometry consistency.

Also, one can estimate the ground plane implicitly by solving a related task. One such example of the task is 3D surface layout recovery [18, 60, 64]. These methods are capable of creating a simple indoor layout reconstruction from a single image and then the ground plane normal is able to be estimated. Another example task could be monocular surface normal estimation [2, 3, 6, 11, 25, 37, 45, 47], which usually formulate the problem as a dense pixel-wise prediction problem and learn a feed-forward deep neural network

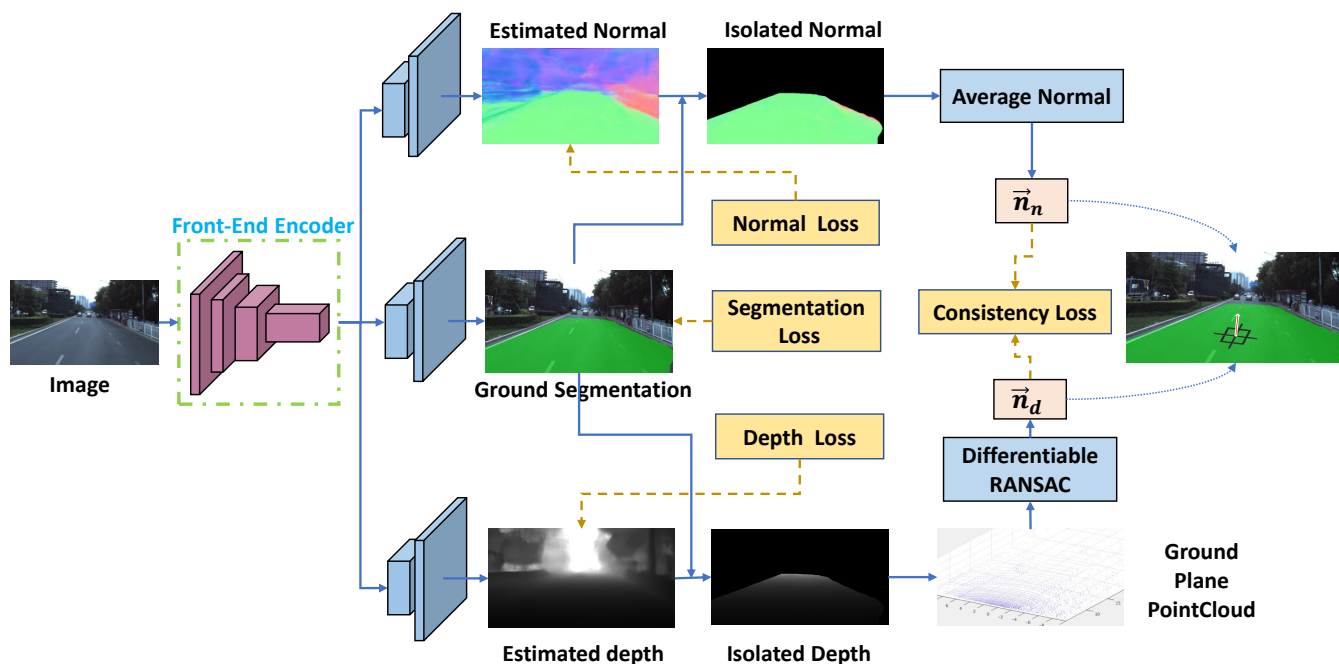


Figure 2: GroundNet: a single image is fed into the front-end encoder and then outputs three streams, namely surface normal estimation, depth estimation and ground segmentation. (1) The ground segmentation is used to isolate the ground regions so that we can selectively back-propagate parameter updates through only the ground regions in the image. (2) Depth stream: we convert the ground region in the image into a point cloud using the estimated depth, and then compute the ground plane normal \vec{n}_d by fitting a plane to the point cloud using the Differentiable RANSAC module. (3) Normal stream: we calculate the ground plane normal \vec{n}_n by taking the average of surface normal vectors of all pixels in the ground region. The final output of our ground plane normal is the average of the ground plane normals computed from the depth and normal streams.

classification model. On top of the estimated pixel-wise normals, one can group the pixels with similar normals into the dominant planes and then compute the average normal for each plane. Although these methods achieve significant progress, most of them are tailed only for the indoor scene while our method is not limited to the scene and can work significantly well on the real-world outdoor datasets. In addition, since these prior works do not explicitly estimate the plane normal, the fact that the ground plane is often a flat and smooth surface is ignored. In contrast, we parameterize the output of our methods to be a planar surface normal explicitly and in the meantime leverage the successful architecture design from the surface normal estimation methods.

Self-Supervised Learning via Geometric Consistency. Geometric consistency is proved to be a useful and free supervision signal in many tasks. [62] learns to predict dense flow field between different instances of the same category object consistently across the synthetic and real domain. [44, 48, 49] propose the cycle consistency loss as a free supervision to learn a tracker. [36] jointly optimizes the 3D surface point positions and normals to be consistent with the observed light refraction effect. [9] learn the keypoint detector, which is consistent across either different viewpoints or adjacent frames. In order to force the predicted 3D bounding box to be consistent with its 2D proposal, [51] proposes a 2D-3D bounding

box consistency loss. [5] learns a depth estimator under a stereo setting by enforcing estimated depth in two views to be consistent with the disparity. Also, the consistent camera pose estimator and 3D shape predictor are learned via the multi-view projection consistency loss in [42]. [8] learns a 3D geometrically-consistent feature map for reconstruction, segmentation and classification from multi-view observations. [26, 61] enforces the estimated ego-motion consistent with the computed motion using 3D (iterative closest point) ICP on the estimated depth between two frames.

Perhaps [35, 55, 56] are closest to our work in the aspect of predicting geometrically consistent depth and surface normal. It has been shown in prior works that the depth and normal are complementary and thus jointly optimizing the two with consistency constraints can improve the performance on both tasks. In our proposed method, we also propose to leverage this consistency loss. Different from prior works, the consistency loss is applied to the estimated ground plane normals from two streams (i.e., the normal estimation stream and depth estimation stream).

3 GEOMETRIC DEFINITION

Parameterization. Given a single RGB image as input, our goal is to estimate the 3D orientation of the ground plane if existing in the image, which is usually represented by an normal vector

$\mathbf{n} = [n_x, n_y, n_z]$ in the world coordinate. In addition, we can also represent the normal vector \mathbf{n} as (θ, ψ) , where θ and ψ are the roll and pitch with respect to the up-axis. Furthermore, we can convert the (θ, ψ) representation into a horizon line representation (θ, ρ) , where θ is the angle that the horizon line makes with the horizontal axis, and ρ is the perpendicular distance from the principal point to the horizon line. This allows us to compare our method against other horizon line estimation methods (see Sec. 5.3).

Pinhole Camera Model. In this paper, we assume the images are captured by a pinhole camera. Suppose $[u_i, v_i]^T$ is the location of a 2D point P_i on the image plane and $Q_i = [x_i, y_i, z_i]^T$ is the corresponding 3D location in the camera coordinate, we have

$$\begin{aligned} x_i &= (u_i - c_x) * z_i / f_x, \\ y_i &= (v_i - c_y) * z_i / f_y, \end{aligned} \quad (1)$$

where $[c_x, c_y]^T$ represents the principle point on the image plane and z_i is the depth of the point. f_x and f_y are the focal length along the x and y directions respectively. It can be also written in homogeneous coordinate as

$$P_i = \begin{bmatrix} \lambda u_i \\ \lambda y_i \\ \lambda \end{bmatrix} = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_i \\ y_i \\ z_i \end{bmatrix} = K_c Q_i \quad (2)$$

where K_c is intrinsic matrix. From Eq. 1 and 2, we know that given a 2D point $P_i = [u_i, v_i]$, the depth in camera coordinate z_i and camera intrinsic matrix K_c , we can calculate its 3D location in the camera coordinate $Q_i = [x_i, y_i, z_i]$. We will use this formulation later in depth stream to convert the image into a point cloud.

4 METHOD

4.1 Approach Overview

Figure 2 depicts the network architecture of the proposed GroundNet. Given a single input image, a front-end fully convolutional encoder first outputs three streams, including a surface normal estimation stream, depth estimation stream and a ground segmentation stream. (See Sec. 5.4 for detailed encoder and stream design.) The segmented ground is used to isolate the ground region in the estimated depth and surface normal, getting rid of the irrelevant objects and noises. We then compute the ground plane normals from the isolated surface normal and depth estimates.

So far, the problem has been formulated as a multi-task prediction problem, including depth estimation, normal estimation and an auxiliary ground segmentation task. However, a vital information is left unused – the geometric correlation between depth and normal. While multi-task learning provides us with informative features and enables cross-modality interactions, geometric consistency is the core constraint that allows the multi-modality information to mutually and correctly refine the prediction. To enforce explicit geometric consistency during training, we add a consistency loss between two ground plane normals computed from the depth and normal streams. During inference, the final output of our ground plane normal is the average of the ground plane normals computed from two streams, although they are already very close after training with the geometric consistency loss.

4.2 Depth Estimation Stream

Given an image I , we first estimate the pixel-wise depth $D(I)$ on the entire image. Then, a ground mask $M(I)$ is used to remove the non-ground region from the estimated depth $D(I)$,

$$\hat{D}(I) = D(I) \odot M(I) \quad (3)$$

As a result, $\hat{D}(I)$ contains only depth of the ground region, which we call the isolated depth. Following Eq. 1 and Eq. 2, we can project every pixel i in $\hat{D}(I)$ to 3D space as $Q(x_i, y_i, z_i)$ given the depth $\hat{D}(i)$ to obtain a 3D point cloud C for points on the ground. We then use a plane fitting algorithm $f(\cdot)$ to get the ground plane normal $\mathbf{n}_d = f(C)$ from the point cloud, where function $f(\cdot)$ can be either least square module, as used in [35], or RANSAC-based module.

Least Square (LS) Module. When using the LS module, it is typically assuming that all points in the point cloud C lie on the same plane without any outlier point. Therefore,

$$C\mathbf{n} = \mathbf{b} \quad (4)$$

where \mathbf{n} is the normal vector of the plane, $\mathbf{b} \in \mathcal{R}^{n \times 3}$ is a constant vector. In the least square equation, we want to minimize $\|C\mathbf{n} - \mathbf{b}\|^2$, which leads to a close-form solution as below

$$\mathbf{n} = \frac{(C^T C)^{-1} C^T \mathbf{1}}{\|(C^T C)^{-1} C^T \mathbf{1}\|} \quad (5)$$

where $\mathbf{1} \in \mathcal{R}^K$ is a constant vector with all 1 elements. Eq. 5 is differentiable and thus can be used in neural network for gradient propagation.

RANSAC-based module. To get rid of the adverse effect from the outliers, it is more suitable to use RANSAC than the LS. As the original RANSAC method is non-differentiable and cannot be used in neural network, we use the differentiable RANSAC (DSAC) proposed in [4]. Instead of using deterministic hypothesis selection (*argmax*), DSAC proposes to use a probabilistic selection, *i.e.*, the probability of selecting a hypothesis is higher when the hypothesis contains more inliers and has a higher score. We will show in Table 1 of the experiment section that the differentiable RANSAC module out-performs the least square module as the differentiable RANSAC module is robust to outliers.

4.3 Surface Normal Estimation Stream

In general, surface normal estimation methods suffer from diverse objects when applied to outdoor scenes. Therefore, similar to the depth stream, a ground mask $M(I)$ is used to isolate the ground region from the entire normal map $N(I)$, which we estimate from the input image I .

$$\hat{N}(I) = N(I) \odot M(I) \quad (6)$$

As a result, $\hat{N}(I)$ contains the normal only for pixels on the ground, which we call the isolated normal map. Then, to compute the ground plane normal (*i.e.*, a single normal vector representing the ground plane) from the pixel-wise isolated normal map, we simply take the average of normal vectors of all pixels on the ground in $\hat{N}(I)$:

$$\mathbf{n}_n = \frac{1}{M} \sum_{p_i \in M(I)} \hat{N}(p_i) \quad (7)$$

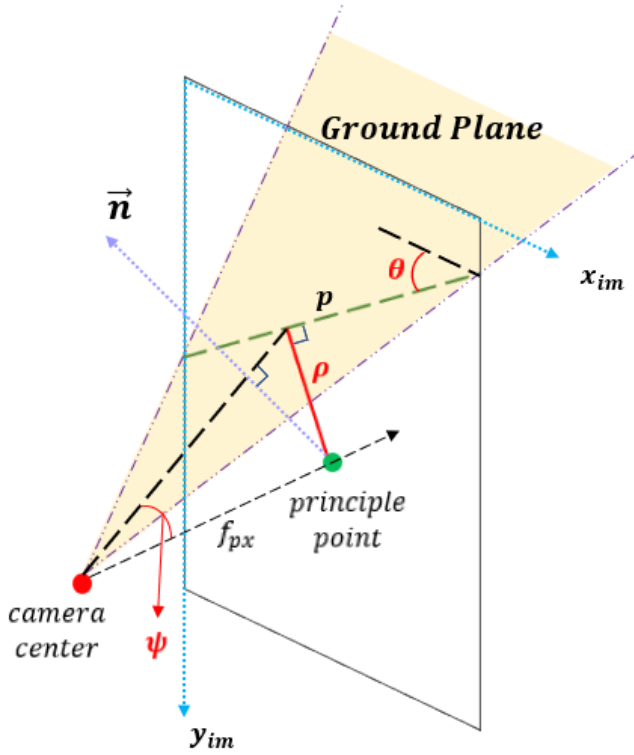


Figure 3: Illustration of the camera model. Ground plane normal $\bar{\mathbf{n}}$ can be transformed into image horizon line, parameterizing as roll θ and pitch ρ , where $\rho = f_{px} \cdot \tan\psi$.

where M is the total number of pixels on the ground in $\hat{N}(I)$ and \mathbf{n}_n is the final output of the ground plane normal computed from the normal estimation stream.

4.4 Geometric Consistency

From the geometric perspective, depth and surface normal are strongly correlated, which follows local linear orthogonality. Formally, for each pixel p_i , such a correlation can be written as a quadratic minimization for a set of linear equations:

$$\mathbb{E} \left(\hat{D}(p_i), \hat{N}(p_i) \right) = \sum_{p_j \in \Omega(p_i)} \left\| (Q_i - Q_j)^T \hat{N}(p_i) \right\|^2 \quad (8)$$

where $\Omega(p_i)$ is the set of neighbors of p_i , and Q_i, Q_j are 3D points projected from 2D points p_i, p_j given their depth in \hat{D} . In addition, in this work, we assume that the ground is a single flat plane. Therefore, Eq. 8 will lead to the same result for any pixel p_j on the ground. In other words, the equation holds even when the neighborhood is the entire ground region, i.e. $\Omega(p_i) = M(I)$. In this way, $\hat{N}(p)$ turns to \mathbf{n}_n . Also, \mathbf{n}_d represents the plane fitted over all 3D points (Eq. 5). Therefore, minimizing Eq. 8 is equivalent to minimize the difference between \mathbf{n}_d and \mathbf{n}_n .

4.5 Loss Functions

We now explain the loss functions of GroundNet. For pixel p_i , we denote the isolated depth map as $\hat{D}(p_i)$ and ground truth depth for the ground region as $D^{gt}(p_i)$ respectively. Similarly, we denote

the isolated surface normal map as $\hat{N}(p_i)$ and ground truth surface normal for the ground region as $N^{gt}(p_i)$. Then, for the depth estimation, the loss \mathcal{L}_{depth} is expressed as

$$\mathcal{L}_{depth} = \frac{1}{M} \sum_{p_i \in M(I)} \left\| \hat{D}(p_i) - D^{gt}(p_i) \right\|_2^2 \quad (9)$$

where M is the total number of pixels in isolated depth map. Similarly, the surface normal estimation loss \mathcal{L}_{normal} is:

$$\mathcal{L}_{normal} = \frac{1}{M} \sum_{p_i \in M(I)} \left\| \hat{N}(p_i) - N^{gt}(p_i) \right\|_2^2 \quad (10)$$

In addition, instead of predicting the depth and surface normal map independently, we enforce the geometric consistency between them by adding a consistency loss \mathcal{L}_{con} on top of the ground plane normals computed from depth and surface normal streams:

$$\mathcal{L}_{con} = \arccos \frac{\mathbf{n}_d \cdot \mathbf{n}_n}{\|\mathbf{n}_d\| \cdot \|\mathbf{n}_n\|} \quad (11)$$

Then, the overall loss function for the GroundNet is the weighted sum of all loss terms,

$$\mathcal{L}_{GroundNet} = \mathcal{L}_{depth} + \mathcal{L}_{normal} + \eta \mathcal{L}_{seg} + \lambda \mathcal{L}_{con} \quad (12)$$

where η and λ are hyper-parameters to controls the relative importance of the loss terms. For ground segmentation stream, we use the common softmax cross-entropy loss as \mathcal{L}_{seg} . Our GroundNet can be trained in an end-to-end manner.

5 EXPERIMENTS

In this section, we evaluate the performance of ground plane estimation network by conducting extensive evaluations (see Sec. 5.5 and 5.6) on two augmented public benchmark datasets, KITTI and ApolloScape (see Sec. 5.1 and 5.2). A thorough derivation of our evaluation metrics is provided in Sec. 5.3. Furthermore, we perform an in-depth ablation study (see Sec. 5.7) to evaluate each component of our method. Additional details about our network architecture and training procedure are reported in Sec. 5.4.

5.1 Datasets

KITTI is a famous and popular outdoor autonomous driving dataset [43], in which disparity depth and road semantic labels are provided for a subset of the dataset. We adopt the split scheme proposed by Eigen et al [12]. The ground truth of the ground plane normal is calculated from the given extrinsic matrix.

ApolloScape is a big autonomous driving dataset for scene parsing [20], instance segmentation and self localization. It contains a great number of image frames with complete depth information and scene labels. We use 40,963 images for training and 8330 images for validation.

5.2 Data Augmentation

For both KITTI and ApolloScape, the cameras are fixed to the car, which means the ground plane orientation does not change significantly over time. As a result, the ground plane normals lack variety in the dataset. Thus, we introduce random rolls and pitches to the dataset by performing rotation and vertical translation to the images, adding variety in training and evaluation process.

Table 1: Ground plane normal evaluation results on KITTI and ApolloScape. (5°, 0.05) means the evaluation dataset is augmented by random adding roll and pitch within 5 degree and 0.05 image units, respectively. LS stands for least square module, and DSAC stands for differentiable RANSAC module. *: Note that the result of [10] is obtained from the original published paper without any dataset augmentation, it only reports evaluation results on KITTI.

	error / deg						
	KITTI				ApolloScape		
	0°, 0	5°, 0.05	10°, 0.1	15°, 0.15	5°, 0.05	10°, 0.1	15°, 0.15
Marr SkipNet [3]	-	4.95	7.23	8.89	5.16	7.33	8.10
GeoNet [35]	-	2.98	4.45	6.57	2.83	4.37	5.72
HMM [10]*	4.10	-			-		
GroundNet (LS)	0.71	2.74	4.06	5.93	2.68	4.02	5.23
GroundNet (DSAC)	0.70	2.65	3.84	5.41	2.49	3.72	4.87



Figure 4: Qualitative results for the ground plane normal estimation. Red: groundtruth. Blue: GroundNet. Yellow: Marr SkipNet. We show that our proposed GroundNet consistently outperforms all other baselines.

Specifically, for roll, we randomly rotate the images around their principal points within a certain limit. For pitch, we randomly translate the images up or down within a certain limit. Afterward, the images are cropped according to the principle point to get rid of black margins. This method allows us to increase the plane normal variety without introducing significant distortion or inaccuracy to the datasets.

Moreover, we set three different rotation limits (5°, 10° and 15°) in order to show the influence of this augmentation. Similarly, the limit of vertical translation is set as 0.05, 0.1 and 0.15 image units.

Note that the training and testing data are both augmented. Therefore, although the data is collected from a car, the testing data will possess a wide variety of orientation, significantly increasing the evaluation difficulty.

5.3 Evaluation Metrics

We compare our method with not only normal estimation methods, but also horizon line estimation methods, because horizon estimation models usually focus on outdoor scenarios. In this section, we provide two evaluation metrics with these two kinds of methods. Particularly, we give proof that under a certain assumption, the ground plane normal is equivalent to the horizon line of an image.

Horizon Line Estimation. To compare with horizon line estimation methods, we change our parameterization from normal vector

$\mathbf{n} = [n_x, n_y, n_z]$ into (θ, ρ) and report the error of these two parameters, as mention in Sec. 3.

As shown in Figure 3, the ground plane can either be represented by its normal \mathbf{n} or be represented by its roll θ and pitch ψ to the horizontal plane, given the camera center and principle point (or given the intrinsic matrix K_c), we can transform our estimated normal into horizon line without losing any information.

From Eq. 2 we know, a point in camera coordinate, Q_i , is related to a point in image coordinate P_i as follows:

$$P_i = [\lambda u_i, \lambda v_i, \lambda]^T = K_c Q_i \quad (13)$$

where K_c is the intrinsic matrix. Therefore, $K_c^{-1} P_i = Q_i$. Let P and Q be the sets of points on the horizon line in camera and image coordinate, respectively. Then, as the normal of a plane is perpendicular to all in-plane vectors. we have $\mathbf{n} \perp Q$, or $Q^T \mathbf{n} = 0$. Therefore, from equation 13:

$$P^T K_c^{-T} \mathbf{n} = 0 \quad (14)$$

Let $K_c^{-T} \mathbf{n} = A = [a, b, c]^T$, the horizon line in image coordinate can be represented as:

$$P^T A = 0 \quad (15)$$

$$ax + by + cz = 0 \quad (16)$$

Let $z = 1, a/c = a_c, b/c = b_c$, and assume that the positive x -direction is to the right, the positive y -direction is down, principle



Figure 5: Qualitative results for the horizon line estimation. Red: groundtruth. Green: GroundNet (DSAC). Cyan: Perceptual. Yellow: DeepHorizon. Gray: Zhai et al. We show that our proposed GroundNet consistently outperforms all other baselines.

Table 2: Horizon line evaluation results on KITTI. (5°, 0.05) means augmentation limits as mentioned in Table 1. The unit for θ and ρ are degree and 10^{-2} image unit.

	5°, 0.05		10°, 0.1		15°, 0.15	
	θ	ρ	θ	ρ	θ	ρ
Zhai et al. [58]	3.36	3.9	4.37	4.9	5.99	6.2
DeepHorizon [53]	2.26	3.7	4.12	4.7	5.93	5.6
Perceptual [19]	1.98	3.0	2.94	3.6	4.66	4.5
GroundNet	1.94	2.6	2.65	3.1	4.17	3.8

point is (W_p, H_p) . Then we can calculate the horizon line (θ, ρ) parameters as:

$$\theta = \arctan -\frac{a_c}{b_c} \tag{17}$$

$$\rho = -\frac{|a_c W_p + b_c H_p + 1|}{\sqrt{a_c^2 + b_c^2}} \tag{18}$$

By these derivations, the horizon line can be derived from the ground plane normal only when the intrinsic K_c is available. In many use cases, such as autonomous cars or mobile robots, it is reasonable to assume that one has access to the intrinsic matrix of the camera. In both KITTI and ApolloScape datasets the camera intrinsic matrices are provided. Therefore, the estimated ground plane normal $\mathbf{n} = [n_x, n_y, n_z]^T$ can be re-parameterized as (θ, ρ) representing the horizon line in image coordinate.

Ground Plane Normal Estimation. To compare with other methods for ground plane normal estimation, we evaluate the angular error between the estimated ground plane normal \mathbf{n} with the ground truth ground plane normal \mathbf{n}_{gt} in terms of degree, as used in [10].

5.4 Implementation Details

We implement GroundNet using the publicly available Tensorflow framework. The front-end convolutional encoder of GroundNet

Table 3: Horizon line evaluation results on ApolloScape. (5°, 0.05) means augmentation limits as mentioned in Table 1. The unit for θ and ρ are degree and 10^{-2} image unit.

	5°, 0.05		10°, 0.1		15°, 0.15	
	θ	ρ	θ	ρ	θ	ρ
Zhai et al. [58]	2.58	2.9	3.94	3.7	5.79	5.1
DeepHorizon [53]	1.98	3.2	3.42	3.3	4.43	3.9
Perceptual [19]	1.77	2.9	2.62	3.0	3.68	3.8
GroundNet	1.92	2.6	2.58	2.8	3.59	3.5

uses VGG-16 [40] as backbone, plus dilated convolution and global pooling as stated in [35]. Using deep features from the front-end encoder, we perform separate deconvolutional operations to generate three sets of task-specific feature maps. Then the separate convolutional operations are performed to produce the prediction maps. The resolution of three maps are made to be the same as input data for the followed 2D to 3D geometric operation. The surface normal map has 3 channels, including 3-direction information, and the other two maps each has 1 channel.

We initialize the front-end encoder with network pre-trained on ImageNet. Note that as the ground segmentation task is simple, for now, we pre-train it with direct ground label supervision separately and fix its weights afterward. The weight of consistency loss λ is set to 0.05. Adam optimizer [22] is used to optimize the network on both datasets, (β_1, β_2) is set to (0.9, 0.999). For KITTI, the initial learning rate is set to 1e-4. For ApolloScape dataset, we initialize with model pre-trained on KITTI and fine-tune with learning rate set to 5e-5. We train GroundNet on one NVIDIA GTX1080Ti GPU, occupying 4GB of memory with a batch size of 4. The running time is 0.92s per image.

In practice, we set the valid depth threshold to 30m. This means only pixels with estimated depth value under 30m are used to fit the plane. This approximation prevents the far away pixels from introducing noise and biases into the estimation.

Table 4: Ablation results on KITTI and ApolloScape with dataset augmentation. We show the final joint training results and the stand-alone surface normal and depth streams without geometry consistency loss. LS stands for least square module.

	error / deg					
	KITTI			ApolloScape		
	5°, 0.05	10°, 0.1	15°, 0.15	5°, 0.05	10°, 0.1	15°, 0.15
surface normal stream	6.73	7.99	9.35	5.78	7.47	8.24
depth stream (LS)	3.01	4.47	6.52	2.97	4.39	5.66
depth stream (DSAC)	2.92	4.29	6.41	2.74	4.25	5.38
joint learning (LS)	2.74	4.06	5.93	2.68	4.02	5.23
joint learning (DSAC)	2.65	3.84	5.41	2.49	3.72	4.87

5.5 Ground Plane Normal Evaluation

The quantitative and qualitative results for the ground plane normal estimation are shown in Table 1 and Figure 4. We compare our methods with two state-of-the-art surface normal estimation methods: Marr-SkipNet [3] and GeoNet [35]; and a direct ground plane normal estimation method: HMM [10].

We re-train GeoNet and Marr SkipNet on KITTI and ApolloScape with the same augmentation strategy. For GeoNet, $\{\alpha, \beta, \gamma, \lambda, \eta\}$ are set to $\{0.95, 9, 0.05, 0.01, 0.5\}$, with initial learning rate set to $1e^{-4}$. For Marr-SkipNet, we train the caffe model with initial learning rate set to $5e^{-4}$. As these two methods provide dense surface normal maps instead of ground plane normal vector, the final result is calculated by taking the average normal value of all pixels on the estimated ground region.

From Table 1, it is clear that GroundNet significantly outperforms all three competitors under four dataset augmentation strategies (including no augmentation). HMM is a non-learning geometry based method, which is tailored to predict ground plane normal; while other two learning-based methods are designed to predict pixel-wise surface normal map without geometry constraints. We claim that GroundNet gets better results with the integration of learning and geometry consistency.

5.6 Horizon Line Evaluation

Following the derivation in Sec. 5.3, we are able to convert our estimated ground plane normal to the horizon line given the intrinsic matrix. We thus compare our method with three state-of-the-art horizon line estimation methods: Zhai et al. [58], DeepHorizon [53] and Perceptual Method [19]. Since [19] does not has the open source code, we implement a DenseNet based model according to their paper.

Quantitative results on KITTI and ApolloScape are shown in Table 2 and 3 respectively. We also plot the qualitative results in Figure 5. We find that our proposed GroundNet outperforms other methods on two datasets under all kinds of dataset augmentation. While the previous methods is able to capture the horizon line based on evident visual cues, GroundNet can get a robust estimation based on depth estimation, surface normal estimation and their intrinsic geometric relationship. Thus GroundNet gets better results when roll and pitch of the ground are steep, or when irrelevant objects (e.g., cars, trees) cover important visual cues like traffic paints and pavements.

Interestingly, the performance of [19] is close or sometimes slightly better than GroundNet when the dataset augmentation

is small (i.e., 5°, 0.05). However, when we increase the amplitude of the augmentation, it is clear to see that the performance of the GroundNet gets better. This shows that our proposed method is more robust to the diversity of the rolls and pitches in the input data. Also, we notice that the average performance of all methods on ApolloScape dataset is better than that KITTI. This is because the images from the ApolloScape dataset are collected on the main road and thus contains more consistent scenarios than KITTI.

5.7 Ablation Study

In order to justify the effect of different model components, we conduct ablation experiments for GroundNet. The results are shown in Table 4. We can see that the joint training model performs better than stand-alone surface normal or depth stream. Note that the results of surface normal and depth stream are not even comparable to GeoNet, as shown in Table 4. This result shows the effectiveness of our geometric consistency loss. At the same time, the differentiable RANSAC module tends to improve the results, compared with the least square module. Thus, we prove that every component of the model helps to get a better ground estimation, so the proposed GroundNet is self-consistent.

6 CONCLUSION

In this paper, we propose GroundNet for ground plane normal estimation from a single image. GroundNet leverages the advantages of both geometry-based and learning-based methods. A geometric consistency loss is applied to two ground plane normals computed from the surface normal stream and depth stream so that GroundNet is able to predict the consistent depth and surface normal. We tailored our model for outdoor scenarios by adding an assistive ground segmentation stream to get rid of irrelevant regions containing diverse objects and background. Experimental results on KITTI and ApolloScape datasets show that the proposed method outperforms previous state-of-the-art methods for both ground plane normal estimation and horizon line estimation in outdoor scenes.

REFERENCES

- [1] Jose M. Alvarez, Theo Gevers, Yann LeCun, and Antonio M. Lopez. 2012. Road Scene Segmentation from a Single Image. *ECCV* (2012).
- [2] Aayush Bansal, Xinlei Chen, and Bryan Russell. 2017. PixelNet: Representation of the Pixels, by the Pixels, and for the Pixels. *arXiv:1702.06506* (2017).
- [3] Aayush Bansal, Bryan Russell, and Abhinav Gupta. 2016. Marr Revisited: 2D-3D Alignment via Surface Normal Prediction. *CVPR* (2016).
- [4] Eric Brachmann, Alexander Krull, Sebastian Nowozin, Jamie Shotton, Frank Michel, Stefan Gumhold, and Carsten Rother. 2017. DSAC - Differentiable

- RANSAC for Camera Localization. *CVPR* (2017).
- [5] Gabriel J Brostow. 2017. Unsupervised Monocular Depth Estimation with Left-Right Consistency. *CVPR* (2017).
 - [6] Weifeng Chen, Donglai Xiang, and Jia Deng. 2017. Surface Normals in the Wild. *ICCV* (2017).
 - [7] Xiaozhi Chen and Yukun Zhu. 2015. 3D Object Proposals for Accurate Object Class Detection. *NIPS* (2015).
 - [8] Ricson Cheng, Ziyang Wang, and Katerina Fragkiadaki. 2018. Geometry-Aware Recurrent Neural Networks for Active Visual Recognition. *NIPS* (2018).
 - [9] Xuanyi Dong, Shou-i Yu, Xinhua Weng, Shih-en Wei, Yi Yang, and Yaser Sheikh. 2018. Supervision-by-Registration: An Unsupervised Approach to Improve the Precision of Facial Landmark Detectors. *CVPR* (2018).
 - [10] Ralf Dragon and Luc Van Gool. 2014. Ground Plane Estimation Using a Hidden Markov Model. *CVPR* (2014).
 - [11] David Eigen and Rob Fergus. 2015. Predicting Depth, Surface Normals and Semantic Labels with a Common Multi-Scale Convolutional Architecture. *ICCV* (2015).
 - [12] David Eigen, Christian Puhrsch, and Rob Fergus. 2014. Depth Map Prediction from a Single Image Using a Multi-Scale Deep Network. *NIPS* (2014).
 - [13] Osian Haines and Andrew Calway. 2012. Detecting Planes and Estimating Their Orientation from a Single Image. *BMVC* (2012).
 - [14] Osian Haines and Andrew Calway. 2015. Recognising Planes in a Single Image. *TPAMI* (2015).
 - [15] Richard Hartley and Andrew Zisserman. 2007. Multiple View Geometry in Computer Vision. *Cambridge University Press* (2007).
 - [16] Varsha Hedau, Derek Hoiem, and David Forsyth. 2010. Thinking Inside the Box: Using Appearance Models and Context Based on Room Geometry. *ECCV* (2010).
 - [17] Derek Hoiem, Alexei Efros, and Martial Hebert. 2006. Putting Objects in Perspective. *CVPR* (2006).
 - [18] Derek Hoiem, Alexei A. Efros, and Martial Hebert. 2007. Recovering Surface Layout from an Image. *IJCV* (2007).
 - [19] Yannick Hold-Geoffroy, Kalyan Sunkavalli, Jonathan Eisenmann, Matthew Fisher, Emiliano Gabbareto, Sunil Hadap, and Jean-François Lalonde. 2018. A Perceptual Measure for Deep Single Image Camera Calibration. *CVPR* (2018).
 - [20] Xinyu Huang, Xinjing Cheng, Qichuan Geng, Binbin Cao, Dingfu Zhou, Peng Wang, Yuanqing Lin, and Ruigang Yang. 2018. The ApolloScape Open Dataset for Autonomous Driving. *CVPRW* (2018).
 - [21] Saad M. Khan and Mubarak Shah. 2006. A Multiview Approach to Tracking People in Crowded Scenes Using a Planar Homography Constraint. *ECCV* (2006).
 - [22] Diederik P Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. *ICLR* (2015).
 - [23] Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven Waslander. 2018. Joint 3D Proposal Generation and Object Detection from View Aggregation. *IROS* (2018).
 - [24] Namhoon Lee, Xinhua Weng, Vishnu Naresh Boddeti, Yu Zhang, Fares Beainy, Kris Kitani, and Takeo Kanade. 2016. Visual Compiler: Synthesizing a Scene-Specific Pedestrian Detector and Pose Estimator. *arXiv:1612.05234* (2016). <http://arxiv.org/abs/1612.05234>
 - [25] Bo Li, Chunhua Shen, Yuchao Dai, Anton Van Den Hengel, and Mingyi He. 2015. Depth and Surface Normal Estimation from Monocular Images Using Regression on Deep Features and Hierarchical CRFs. *CVPR* (2015).
 - [26] Reza Mahjourian, Martin Wicke, and Anelia Angelova. 2018. Unsupervised Learning of Depth and Ego-Motion from Monocular Video Using 3D Geometric Constraints. *CVPR* (2018).
 - [27] Aashi Manglik, Xinhua Weng, Eshed Ohn-bar, and Kris M Kitani. 2019. Future Near-Collision Prediction from Monocular Video: Feasibility, Dataset, and Challenges. *arXiv:1903.09102* (2019). <http://arxiv.org/abs/1903.09102>
 - [28] Matthew W. McDaniel, Takayuki Nishihata, Christopher A. Brooks, and Karl Iagnemma. 2010. Ground Plane Identification Using LIDAR in Forested Environments. *ICRA* (2010).
 - [29] Branislav Mićušik and Jana Košecká. 2009. Piecewise Planar City 3D Modeling from Street View Panoramic Sequences. *CVPRW* (2009).
 - [30] Branislav Micusik, Horst Wildenauer, and Markus Vincze. 2008. Towards Detection of Orthogonal Planes in Monocular Images of Indoor Environments. *ICRA* (2008).
 - [31] Faisal Mufti, Robert Mahony, and Jochen Heinzmann. 2012. Robust Estimation of Planar Surfaces Using Spatio-Temporal RANSAC for Applications in Autonomous Vehicle Navigation. *Robotics and Autonomous Systems* (2012).
 - [32] J. A. De Jesus Osuna-Coutino, Claudia Cruz-Martinez, Jose Martinez-Carranza, Miguel Arias-Estrada, and Walterio Mayol-Cuevas. 2016. I Want to Change My Floor: Dominant Plane Recognition from a Single Image to Augment the Scene. *ISMAR* (2016).
 - [33] Hannes Ovrén and Per-Erik Forssén. 2018. Spline Error Weighting for Robust Visual-Inertial Fusion. *CVPR* (2018).
 - [34] David Pfeiffer and Uwe Franke. 2010. Efficient Representation of Traffic Scenes by Means of Dynamic Stixels. *IEEE Intelligent Vehicles Symposium* (2010).
 - [35] Xiaojuan Qi, Renjie Liao, Zhengzhe Liu, Raquel Urtasun, and Jiaya Jia. 2018. GeoNet: Geometric Neural Network for Joint Depth and Surface Normal Estimation. *CVPR* (2018).
 - [36] Yiming Qian, Minglun Gong, and Yee-Hong Yang. 2016. 3D Reconstruction of Transparent Objects with Position-Normal Consistency. *CVPR* (2016).
 - [37] Zhongzheng Ren and Yong Jae Lee. 2018. Cross-Domain Self-supervised Multi-task Feature Learning using Synthetic Imagery. *CVPR* (2018).
 - [38] Johannes Lutz Schönberger and Jan-Michael Frahm. 2016. Structure-from-Motion Revisited. *CVPR* (2016).
 - [39] Stephen Se and Michael Brady. 2002. Ground Plane Estimation, Error Analysis and Applications. *Robotics and Autonomous Systems* (2002).
 - [40] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. *ICLR* (2015).
 - [41] Jean Philippe Tardif, Yanis Pavlidis, and Kostas Daniilidis. 2008. Monocular Visual Odometry in Urban Environments Using an Omnidirectional Camera. *IROS* (2008).
 - [42] Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. 2018. Multi-View Consistency as Supervisory Signal for Learning Shape and Pose Prediction. *CVPR* (2018).
 - [43] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. 2017. Sparsity Invariant CNNs. *3DV* (2017).
 - [44] Ning Wang, Yibing Song, Wengang Zhou, Wei Liu, and Houqiang Li. 2019. Unsupervised Deep Tracking. *arXiv:1904.01828* (2019).
 - [45] Peng Wang, Xiaohui Shen, Bryan Russell, Scott Cohen, Brian Price, and Alan L. Yuille. 2016. SURGE: Surface Regularized Geometry Estimation from a Single Image. *NIPS* (2016).
 - [46] Sen Wang, Daoyuan Jia, and Xinhua Weng. 2018. Deep Reinforcement Learning for Autonomous Driving. *arXiv:1811.11329* (2018). <http://arxiv.org/abs/1811.11329>
 - [47] Xiaolong Wang, David F. Fouhey, and Abhinav Gupta. 2015. Designing Deep Networks for Surface Normal Estimation. *CVPR* (2015).
 - [48] Xiaolong Wang, Allan Jabri, and Alexei A. Efros. 2019. Learning Correspondence from the Cycle-Consistency of Time. *CVPR* (2019).
 - [49] Xinhua Weng and Wentao Han. 2018. CyLKs: Unsupervised Cycle Lucas-Kanade Network for Landmark Tracking. *arXiv:1811.11325* (2018). <http://arxiv.org/abs/1811.11325>
 - [50] Xinhua Weng and Kris Kitani. 2019. A Baseline for 3D Multi-Object Tracking. *arXiv:1907.03961* (2019). <https://arxiv.org/pdf/1907.03961.pdf>
 - [51] Xinhua Weng and Kris Kitani. 2019. Monocular 3D Object Detection with Pseudo-LiDAR Point Cloud. *arXiv:1903.09847* (2019). <http://arxiv.org/abs/1903.09847>
 - [52] Xinhua Weng, Shangxuan Wu, Fares Beainy, and Kris Kitani. 2018. Rotational Rectification Network: Enabling Pedestrian Detection for Mobile Vision. *WACV* (2018).
 - [53] Scott Workman, Menghua Zhai, and Nathan Jacobs. 2016. Horizon Lines in the Wild. *BMVC* (2016).
 - [54] Shangxuan Wu and Xinhua Weng. 2018. Image Labeling with Markov Random Fields and Conditional Random Fields. *arXiv:1811.11323* (2018). <http://arxiv.org/abs/1811.11323>
 - [55] Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, and Ram Nevatia. 2018. LEGO: Learning Edge with Geometry all at Once by Watching Videos. *CVPR* (2018).
 - [56] Zhenheng Yang, Peng Wang, Wei Xu, Liang Zhao, and Ramakant Nevatia. 2018. Unsupervised Learning of Geometry with Edge-aware Depth-Normal Consistency. *AAAI* (2018).
 - [57] Chang Yuan and Gérard Medioni. 2006. 3D Reconstruction of Background and Objects Moving on Ground Plane Viewed from a Moving Camera. *CVPR* (2006).
 - [58] Menghua Zhai, Scott Workman, and Nathan Jacobs. 2016. Detecting Vanishing Points Using Global Image Context in a Non-Manhattan World. *CVPR* (2016).
 - [59] Wei Zhang and Jana Kosecka. 2003. Extraction, Matching and Pose Recovery Based on Dominant Rectangular Structures. *ICCV* (2003).
 - [60] Yinda Zhang, Shuran Song, Ping Tan, and Jianxiang Xiao. 2014. PanoContext: A Whole-Room 3D Context Model for Panoramic Scene Understanding. *ECCV* (2014).
 - [61] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. 2017. Unsupervised Learning of Depth and Ego-Motion from Video. *CVPR* (2017).
 - [62] Tinghui Zhou, Philipp Krähenbühl, Mathieu Aubry, Qixing Huang, and Alexei A. Efros. 2016. Learning Dense Correspondence via 3D-Guided Cycle Consistency. *CVPR* (2016).
 - [63] Zihan Zhou. 2012. Robust Plane-Based Structure From Motion. *CVPR* (2012).
 - [64] Chuhan Zou, Alex Colburn, Qi Shan, and Derek Hoiem. 2018. LayoutNet: Reconstructing the 3D Room Layout from a Single RGB Image. *CVPR* (2018).