

Accurate Training Data for Occupancy Map Prediction in Automated Driving Using Evidence Theory

Jonas Kälble[†] Sascha Wirges[†] Maxim Tatarchenko[†] Eddy Ilg[‡]
[†]Bosch Center for Artificial Intelligence, Germany
[‡]Saarland University, Germany
 {firstname.lastname}@de.bosch.com

Abstract

Automated driving fundamentally requires knowledge about the surrounding geometry of the scene. Modern approaches use only captured images to predict occupancy maps that represent the geometry. Training these approaches requires accurate data that may be acquired with the help of LiDAR scanners. We show that the techniques used for current benchmarks and training datasets to convert LiDAR scans into occupancy grid maps yield very low quality, and subsequently present a novel approach using evidence theory that yields more accurate reconstructions. We demonstrate that these are superior by a large margin, both qualitatively and quantitatively, and that we additionally obtain meaningful uncertainty estimates. When converting the occupancy maps back to depth estimates and comparing them with the raw LiDAR measurements, our method yields a MAE improvement of 30% to 52% on *nuScenes* and 53% on *Waymo* over other occupancy ground-truth data. Finally, we use the improved occupancy maps to train a state-of-the-art occupancy prediction method and demonstrate that it improves the MAE by 25% on *nuScenes*.

1. Introduction

The safe operation of automated vehicles requires accurate environment models to reason about geometry and semantics. The most popular geometry model today is an occupancy grid map that tessellates the space into a regular grid of volume elements called voxels [5]. Such occupancy grid maps allow for compute-efficient free-space estimation (e.g., via ray casting) and data storage via Octrees [21, 25].

For several years, automated driving has mainly relied on LiDAR due to its low and distance-independent measurement error. However, recent approaches at-

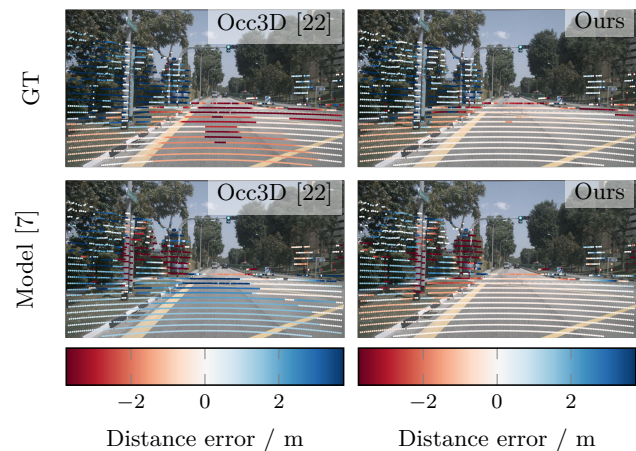


Figure 1. **Comparison of Depth Errors Between Occupancy Grid Maps and LiDAR Measurements.** We compare our grid mapping approach (right) to Occ3D [22] (left). Top: Depth errors between LiDAR scan and ground-truth occupancy map. Bottom: Depth errors between LiDAR scan and model predictions [7].

tempt to reconstruct the 3D environment from images [8, 14, 23, 30] as they argue LiDAR is sparse, expensive, and redundant¹.

Enabling camera-based approaches for occupancy estimation requires data from real-world traffic scenes with ground-truth (GT) geometry, where LiDAR is still the most reliable source of information due to its low measurement errors. First, refinement steps are applied to densify the sparse point clouds [22, 26]. Then the GT occupancy is reconstructed from the LiDAR measurements and used to train and test occupancy prediction methods. This is done by aggregating point clouds over multiple frames and compensating the movement of the ego vehicle and dynamic objects. Voxels are considered

¹<https://techcrunch.com/2019/04/22/anyone-relying-on-lidar-is-doomed-elon-musk-says/>

occupied if they contain LiDAR reflections and free if they are traversed by LiDAR rays between sensor origin and surface. Errors in the reconstruction may occur due to multiple reasons: 1) noisy laser measurements, 2) few measurements available for a voxel, 3) incorrectly applied densification of points, and 4) incorrect motion compensation.

We show that the currently used techniques [22, 23, 26, 27] yield very low-quality occupancy maps and present a new approach by modeling the sensor characteristics and using evidence theory to estimate the degree of belief in voxels being free, occupied, or uncertain. Using this framework enables us to explicitly assign missing measurements, occlusions, or conflicts due to measurement, label or pose estimation errors to an “*I don’t know*” hypothesis. The results show that our approach based on evidence theory is able to compensate for the reconstruction errors significantly better than existing methods, as it yields substantially more accurate occupancy maps when comparing them to the original LiDAR measurements. Fig. 1 shows an example of our method compared to Occ3D [22].

Furthermore, to train methods to predict occupancy from images, camera visibility masks are required in addition to the occupancy GT. We show that these masks significantly underestimate visibility and as a result prevent supervision in critical regions of the scene, manifesting in overestimated object boundaries and incorrect ground surface estimation. Excluding all the regions that are incorrectly masked as invisible leads to methods being incapable of performing reliable estimates in these areas and eventually renders them unusable for safety-critical applications. To this end, we propose a new loss weighting which makes use of our estimated uncertainty for supervision. We show that using our data and loss weighting, the predictions from state-of-the-art (SOTA) methods too can be substantially improved when comparing them to the original LiDAR measurements. Our contributions are threefold:

1. We provide a simple, multi-frame 3D LiDAR occupancy grid mapping method based on evidence theory, taking into account multiple input frames (see Sec. 3). Our approach captures continuous evidence on voxels being occupied, free and uncertain due to measurement errors or missing observations from LiDAR measurements.
2. We show that the data provided by SOTA benchmarks [22, 23] is inaccurate w.r.t. LiDAR measurements and generate more accurate, evidential occupancy grid maps that can be used as a substitution for binary occupancy prediction of much higher quality. Given our generated data, we demonstrate an occupancy estimation improvement of 30–53% MAE

- compared to the SOTA benchmark data (see Tab. 1).
3. Derived from our data, we propose a simple yet effective observability-based loss weighting to train SOTA models, producing 24% lower MAEs between occupancy predictions and LiDAR measurements (see Sec. 4.4).

The code to generate the data is available via GitHub².

2. Related Work

3D Occupancy Datasets. Datasets for automated driving, such as KITTI [6], nuScenes [2], Waymo [20] and Agroverse2 [28] record traffic scenes using multiple sensors like LiDARs and cameras. They provide additional data annotations to enable the training of models for downstream tasks such as 3D object detection [11, 13, 16, 17]. 3D occupancy prediction benchmarks for automated driving [12, 22, 23, 26, 27, 27] use the recorded sensor data to generate dense 3D occupancy grids to represent arbitrary geometry of the scene surrounding the vehicle.

A basic processing pipeline to obtain dense occupancy grids consists of three steps: First, LiDAR measurements are aggregated over multiple time steps by considering static and dynamic objects separately. Second, the points are partitioned into a dense voxel grid. Every voxel that contains a point is labeled as occupied. Finally, occlusion reasoning is needed for non-occupied voxels to evaluate whether a voxel is free or unobserved. Occ3D [22] casts rays from the LiDAR origins to each observed reflection and sets traversed voxels as free. Additionally, for the tasks of vision-based occupancy prediction, Tian et al. calculate a camera visibility mask by casting rays from the camera origins to the centers of occupied voxels. The first occupied voxel along each ray is marked as visible, any voxel behind it is labeled as unobserved due to occlusion.

Various techniques are used to improve the final voxel representation quality. SurroundOcc [27] uses Poisson surface reconstruction [10] to make the point clouds denser. Occ3D [22] proposes to fill the holes in the point cloud by mesh reconstruction through VDBFusion [24] and use image-guided voxel refinement to enforce 2D-to-3D consistency between camera images and the occupancy grid. OpenOccupancy [26] makes use of their Augment and Purifying pipeline together with additional human annotation effort to improve the label quality. In contrast, our method proposes to embed information about free and unobserved space much earlier in the pipeline. Therefore, we make use of the number of transmissions and reflections per voxel and resolve conflicting measurements using evidence theory (ET).

²<https://github.com/boschresearch/evidential-occupancy>

3D Occupancy Prediction. The goal of occupancy prediction is to estimate the state of each voxel in a scene to be either *occupied*, *free* or *unobserved* based on current and past camera images. Semantic occupancy prediction additionally includes the goal of predicting a specific class for each occupied voxel. The *CVPR 2023 Occupancy Prediction Challenge* hosted by Tong et al. and Tian et al. focuses on the task of semantic occupancy prediction from camera images only. The winning method FB-OCC by Li et al. [14, 15] builds upon BEVDet [8] and introduces forward-backward view transformations. TPVFormer [9] uses three perpendicular planes to represent the 3D space efficiently. BEVDepth [11] introduces additional depth supervision to guide the process of lifting 2D image features into 3D space. BEVDet4D [7] exploits temporal information to improve 3D object detection. Specifically, we choose BEVDet4D-Occ-R50 as our occupancy prediction baseline since it’s widely adopted in the domain of vision based occupancy prediction and serves as baseline for many comparisons.

Occupancy Grid Mapping. Bayesian static occupancy grid mapping is a well-established concept initially developed by Elfes [4] and Moravec [18] who model occupancy as a Bernoulli-distributed random variable. These approaches, however, do not consider the amount of information collected to estimate the cell state, neither do they explicitly handle conflicts between different measurements. Yang and Aitken [29] resolve the above-mentioned issues using evidence theory. Since then, more recent research focuses on 3D mapping, within dynamic environments and/or including semantics. For instance, Richter et al. [19] model the dynamic traffic scene by two evidential semantic top-view grid maps, representing the ground and everything on top by two top-view grid maps. Baur et al. [1] determine scene flow from LiDAR measurements to improve mapping in dynamic environments. In contrast to [1, 19], we focus on 3D evidential occupancy grid mapping, with instance flow determined by labeled object shapes and motion.

3. Method

Our method starts by mapping LiDAR measurements into a spherical voxel grid by adding up the reflections and transmissions for each voxel. In the next step, we warp the spherical voxel grids from different time steps to a common Cartesian voxel grid for the current time step and aggregate the number of reflections and transmissions. Finally, we assign evidences to a voxel being *occupied*, *free* or *uncertain*. In the following, we describe each step in detail.

3.1. Spherical Sensor Mapping

We start by mapping the LiDAR measurements from every frame into two spherical grid maps denoted as *reflections* and *transmissions*. Every voxel in the reflections grid map encodes how many reflections originated from there, which is usually high for voxels capturing object surfaces. Every voxel in the transmissions grid map, on the other hand, encodes how many LiDAR rays freely passed this voxel from the LiDAR origin to their reflection positions, which is usually high for voxels representing free space.

Given the LiDAR field of view and its characteristics, we map measured points into spherical coordinates with radial distance (ρ), polar angle (θ), and azimuth angle (ϕ), which we then quantize with $\Delta\rho, \Delta\theta, \Delta\phi$. The spherical voxel size corresponds to the measurement errors $\pm \frac{\Delta\rho}{2}$, $\pm \frac{\Delta\theta}{2}$ and $\pm \frac{\Delta\phi}{2}$ that we obtain from the sensor characteristics.

Reflections Grid Map. As illustrated in the top of Fig. 3, to obtain the reflection value for each voxel, we define a cube centered at the current measurement point with the same size as the grid cells. We then use the volume overlap to distribute the probability mass of the measurement into the neighboring voxels. This step can be efficiently implemented as weighted point scattering. As a result, each spherical voxel (ρ, ϕ, θ) is assigned a reflection value r .

Transmissions Grid Map. Given the reflections, we can now compute the transmissions grid map. For a given (ρ, ϕ, θ) bin, its transmission

$$q(\rho, \phi, \theta) = \sum_{\rho' > \rho} r(\rho', \phi, \theta) \tag{1}$$

is determined as the sum of all reflection values that lie within the same ϕ and θ bin but have a distance larger than ρ . Intuitively, this means that the more rays have landed farther than the current bin and therefore must have passed through the current bin, the higher the likelihood that this bin is free. Fig. 3 illustrates an example of this computation step which can be efficiently implemented via the exclusive cumulative sum of reflections from the maximum to the minimum radial distance.

Note that generally r and q are not probability values, but the number of events associated to a voxel, taking any positive real value due to the weighted assignment.

3.2. Multi-Frame Aggregation

In this step, we aggregate the reflections and transmissions from multiple time steps into a single Cartesian

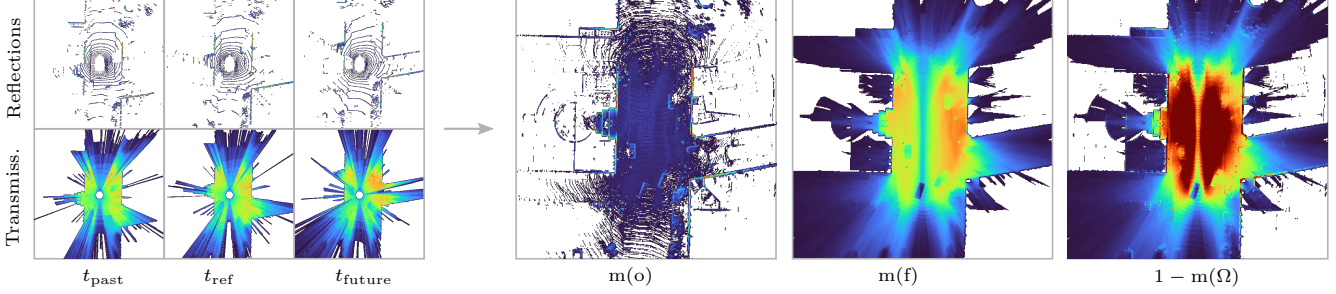


Figure 2. **Conversion of Transmissions and Reflections to Beliefs.** We start by calculating reflections and transmissions in a spherical coordinate system for each time step individually (left). Then we aggregate the number of reflections and transmissions from past, reference, and future frames into one common Cartesian grid at the reference time t_{ref} . During this process, we compensate for object and ego vehicle motion. We use the basic belief assignment described in Sec. 3.3 to obtain the occupied belief $m(o)$, the free belief $m(f)$, and the observability $1 - m(\Omega)$.

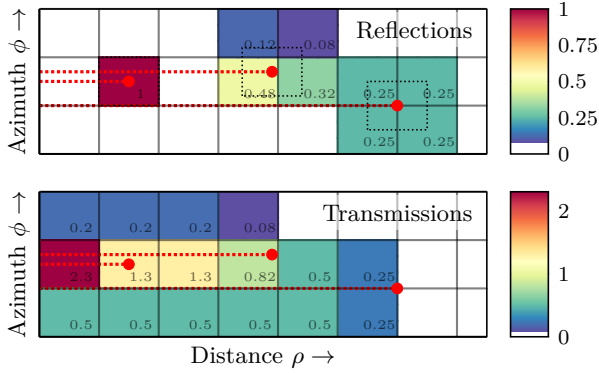


Figure 3. **Spherical Reflection and Transmission Grid Mapping.** Top: Weighted scattering of measurements to spherical voxels. Bottom: Transmission value computation through cumulative sum starting from maximum and ending in minimum distance. For illustration, we omit the polar angle θ . Colors denote the value assigned to a grid cell. The bottom shows the corresponding transmissions that are derived from the reflections.

voxel grid and warp information from past and future frames into the reference frame using the known per-object motion. Here, we consider all frames within a viewpoint displacement due to ego motion of less than δ_{max} from the reference frame at time t_{ref} , limited to at most N_{max} frames evenly sampled in time.

We start by defining a Cartesian volume aligned to the current vehicle pose in the reference frame at time t_{ref} . We consider the coordinate system of this frame to be the target coordinate system, into which all other frames (at different t) are warped. We assume that for every annotated object o from the reference frame we know its scene flow $f_o(t)$ with every past/future frame t . $f_o(t)$ allows us to transform every object point p from

the reference frame t_{ref} into frame t as follows:

$$p(\tilde{t}) = f_o(\tilde{t}) + p(t). \quad (2)$$

We now iterate over every object in frame t_{ref} and transform every of its i -th voxel centers $p_i(t_{\text{ref}})$ into all future and past frames, yielding a set of T new coordinates $\{p_i(t)\}_{t=1}^T$. Note that here t_{ref} is included among the T time steps, with all flow vectors being zero. Next, we transform each coordinate $p_i(t) = [x_i(t), y_i(t), z_i(t)]$ into the spherical coordinates

$$\rho_i(t) = \sqrt{x_i(t)^2 + y_i(t)^2 + z_i(t)^2} \quad \text{with} \quad (3)$$

$$\theta_i(t) = \arccos\left(\frac{z_i(t)}{\rho_i(t)}\right) \quad \text{and} \quad \phi_i(t) = \arctan\left(\frac{y_i(t)}{x_i(t)}\right).$$

Finally, we retrieve the number of reflections and transmissions via trilinear interpolation from the corresponding precomputed maps. In addition, we normalize sampled values by multiplying with the voxel volume ratio

$$s = \frac{V_{\text{cartesian}}}{V_{\text{spherical}}} \quad \text{with} \quad (4)$$

$$V_{\text{cartesian}} = \Delta x \Delta y \Delta z \quad \text{and} \quad (5)$$

$$V_{\text{spherical}} = \int_{\rho - \frac{\Delta \rho}{2}}^{\rho + \frac{\Delta \rho}{2}} \int_{\theta - \frac{\Delta \theta}{2}}^{\theta + \frac{\Delta \theta}{2}} \int_{\phi - \frac{\Delta \phi}{2}}^{\phi + \frac{\Delta \phi}{2}} \tilde{\rho}^2 \sin(\tilde{\theta}) \, d\tilde{\rho} \, d\tilde{\theta} \, d\tilde{\phi}, \quad (6)$$

between Cartesian and spherical cells. Finally, we aggregate the number of reflections and transmissions

$$r_i = \frac{1}{T} \sum_{t=1}^T r_i(t) \quad \text{and} \quad q_i = \frac{1}{T} \sum_{t=1}^T q_i(t) \quad (7)$$

from all $T \leq N_{\text{max}}$ frames within δ_{max} , normalized by the number of considered frames.

3.3. Evidential Mapping

Each voxel in the final grid map should not only represent information about occupancy and free space, but also uncertainty. To determine this information, we use the framework of ET (also referred to as Dempster-Shafer theory), which can be interpreted as a generalization of Bayesian theory [3]. It allows assigning evidence from different information sources not only to a single hypothesis but to all possible combinations. Let Ω be the set of hypotheses under consideration. Then, the basic belief assignment (BBA)

$$m: 2^\Omega \rightarrow [0, 1], \quad m(\emptyset) = 0, \quad \sum_{X \in 2^\Omega} m(X) = 1 \quad (8)$$

assigns evidence mass to each element of its power set 2^Ω . Here, we define

$$\Omega = \{o, f\}, \quad (9)$$

containing the hypotheses *occupied* (o) and *free* (f). Note that the hypothesis Ω contains both o and f and corresponds to being uncertain. As suggested in [19], we use the BBA

$$m(\omega|q, r) = \begin{cases} p_{\text{FN}}^q(1 - p_{\text{FP}}^r) & \text{for } \omega = o \\ p_{\text{FP}}^r(1 - p_{\text{FN}}^q) & \text{for } \omega = f \\ 1 - m(o) - m(f) & \text{for } \omega = \Omega \end{cases} \quad (10)$$

to determine each voxel’s evidence towards every possible hypothesis ω .

In this formulation, the false negative probability p_{FN} and false positive probability p_{FP} are parameters characterizing the sensor. p_{FN} is the probability of a voxel being estimated as free when it should in fact be occupied, p_{FP} is the probability of predicting the state of a voxel to be occupied when it should be free. These values are fixed for the entire dataset, and are determined heuristically as described in section Sec. 4.3. We show an example of input reflections and transmissions and our estimated beliefs in Fig. 2.

4. Experiments

4.1. Evaluation

We have observed that occupancy grids and especially visibility masks, which are used for model training, are of low quality. Therefore, we propose to quantify the agreement between the raw LiDAR measurements and the reconstructed occupancies by rendering depth maps and comparing them. We argue that comparing to the raw LiDAR measurements provides the most reliable available reference.

BBA to Binary Occupancy. To compare our continuous occupancy maps to other occupancy prediction methods, we need a binary occupancy value per voxel. Therefore, we compare the occupied belief $m_i(o)$ and the free belief $m_i(f)$ and assign a binary occupancy value o_i to each voxel i

$$o_i = \begin{cases} 1 & \text{if } m_i(o) > m_i(f) \\ 0 & \text{else} \end{cases} \quad (11)$$

This yields a binary occupancy map, which we use to compare against other occupancy methods.

Depth Evaluation. For each existing GT LiDAR point i , we march along the corresponding ray and extract the predicted depth value d_i^{pr} as the first intersection point of this ray with a voxel surface predicted to be occupied. We then compare the predicted value d_i^{pr} with the GT LiDAR depth d_i^{gt} . We report several standard metrics from the depth estimation literature: MAE, root mean squared error (RMSE) and its log-variant (RMSE_{log}), and the percentage of points that have a relative depth deviation δ smaller than a specified threshold. Since the dense occupancy grid has a limited lower and upper bound, we only consider LiDAR measurements within the enclosed volume. A complete overview of used parameters is given in Tab. 2.

4.2. Occupancy Grid Comparison

We compare our generated occupancy grids to those of Scene as Occupancy [23], SurroundOcc [27], Occ3D [22] and OpenOccupancy [26] on nuScenes [2] and to those of Occ3D [22] on Waymo [20]. Especially, Occ3D provides the semantic occupancy labels for the *CVPR 2023 3D Occupancy Prediction Challenge* built on top of the nuScenes dataset [2].

We use the described evaluation protocol from Sec. 4.1 and compare rendered depths in the occupancy grids to the LiDAR measurements. We summarize the quantitative results in Tab. 1 and show a qualitative result of the rendered depth error in Fig. 1. Fig. 4 depicts qualitatively the improvement of our method. The ground surface is significantly more accurate and objects like the streetlights have sharper boundaries. Flying artifacts behind moving objects are reduced. The supplementary material contains additional qualitative results. Our method outperforms all baseline methods across all metrics by a large margin. Therefore, our generated occupancy maps matches the LiDAR measurements better than previous methods according to all evaluated metrics. This also holds for different voxel sizes and across different datasets: We outperform other methods [22, 23, 26, 27] with comparable voxel size on nuScenes [2] and Waymo [20].




Figure 4. **Comparison of Occupancy Maps.** We show the occupancy map by OpenOccupancy [26] (left) and ours (right) for a nuScenes [2] scene. Both methods use a voxel size of 0.2 m. We encode the color depending on the voxel’s z -coordinate. The bottom row depicts the three front facing cameras of the vehicle. Our method shows improved quality on the ground surface and thin objects like the streetlights. It also reduces flying particles behind moving vehicles and pedestrians.

Method @ Voxel size		MAE \blacktriangle	RMSE \blacktriangle	RMSE _{log} \blacktriangle	$\delta < 1.25$ \blacktriangle	$\delta < 1.25^2$ \blacktriangle	$\delta < 1.25^3$ \blacktriangle
		in m	in m	in log m	in %	in %	in %
nuScenes	Scene as Occupancy [23] @ 0.5 m	2.44	4.99	0.57	64.4	81.7	87.4
	SurroundOcc [27] @ 0.5 m	2.36	5.11	0.46	72.0	83.6	89.0
	Occ3D [22] @ 0.4 m	1.80	4.43	0.37	81.4	89.7	93.3
	Ours @ 0.4 m	1.25	3.42	0.28	87.9	94.6	96.8
	OpenOccupancy [26] @ 0.2 m	1.89	4.39	0.51	75.0	85.5	90.5
	Ours @ 0.2 m	0.91	2.99	0.24	92.6	96.2	97.7
Waymo	Occ3D [22] @ 0.4 m	3.23	6.52	0.56	75.6	83.1	87.6
	Ours @ 0.4 m	1.51	3.98	0.23	90.6	95.1	97.1
	Ours @ 0.2 m	1.15	3.96	0.21	92.8	95.6	97.3

Table 1. **Comparison of Depth Errors Between Occupancy Grid Maps and LiDAR Measurements.** We use raw LiDAR measurements as a common ground truth to compare between different occupancy grid maps. Therefore, we calculate the first intersection of a LiDAR ray and an occupied voxel to determine a depth value as described in Sec. 4.1. We use established depth estimation metrics to evaluate the quality of the generated occupancy maps. For [22, 23, 26, 27], we map all semantic classes to be occupied. We use our BBA to determine occupied and free voxel cells by comparing the occupied belief $m(o)$ and free belief $m(f)$. Apart from the different spherical grid extents due to different FOVs of the LiDAR sensors we keep all hyperparameters constant for both datasets. We outperform existing methods with comparable voxel size across all metrics by a large margin both on nuScenes [2] and on Waymo [20].

4.3. Hyperparameters

As a first step, we determine the values for the false negative and false positive probabilities p_{FN} and p_{FP} from Eq. (10). These probabilities depend on multiple factors, including sensor characteristics and hyperparameters like the Cartesian voxel cell size. Therefore, we

do a grid search over possible combinations of p_{FN} and p_{FP} on a small portion of the nuScenes dataset. We evaluate rendered depths and compare them to the ground truth LiDAR measurements as described in Sec. 4.1. We then choose the best values for p_{FN} and p_{FP} according to Fig. 5. Note that we run the grid search on a small

Parameters	Value
Spherical grid extent	$\rho = 2.5$ m to 60 m $\theta = 75^\circ$ to 125° $\phi = -180^\circ$ to 180°
Spherical cell size	$\Delta\rho = 0.1$ m, $\Delta\theta = \Delta\phi = 0.5^\circ$
Cartesian grid extent	$x = -40$ m to 40 m $y = -40$ m to 40 m $z = -1$ m to 5.4 m
Cartesian cell size	$\Delta x = \Delta y = \Delta z = 0.2$ m or 0.4 m
False neg. probability	$p_{FN} = 0.8$ or 0.9
False pos. probability	$p_{FP} = 0.2$ or 0.1
Max. number of frames	$N_{\max} = 50$
Max. displacement	$\delta_{\max} = 20$ m

Table 2. **Parameter Overview.** If not noted otherwise, we use the parameters listed here. We choose the spherical grid extent such that it encloses most of the LiDAR measurements. The Cartesian grid extent matches the one of [22].

subset of the nuScenes dataset, but use the same hyperparameters for the evaluation on the Waymo dataset. This highlights the method’s applicability for different sensor setups.

4.4. Training

We use our generated GT data to train existing occupancy prediction methods, with only minor modifications of the original approaches. Throughout the experiments we keep all model specific parameters fixed and only change the last layer. All models are trained and evaluated using the nuScenes dataset. We choose Occ3D’s occupancy GT [22] as reference and compare it to training with our occupancy GT. In the following, we introduce our baseline, as well as two modes of supervision to make use of our training data.

Baseline. As detailed in the related work, BEVDet4D-Occ with ResNet50 as backbone by Huang et al. [7, 8] serves as our main baseline. We start from a checkpoint pre-trained on the data of the *CVPR 2023 Occupancy Prediction Challenge* [22, 23]. Since the original model is trained for semantic occupancy prediction, we fine-tune it for the task of occupancy prediction without semantics. Therefore, we map all occupied classes to a single binary occupancy target. We use the sigmoid activation function and the binary cross-entropy as a loss function. Similar to the pre-training procedure, we use the camera visibility mask provided in the challenge to only supervised voxels, which are labeled as visible. We use a small learning rate of 1×10^{-5} and fine-tune for three epochs. All subsequent experiments follow the same fine-tuning procedure. We only modify the last layer, the loss function and the GT per-voxel supervision.




Figure 5. **Determining the False Negative and False Positive Probabilities.** To select reasonable values for p_{FP} and p_{FN} , we run a grid search on the nuScenes mini split and evaluate against the raw LiDAR data. We plot MAE (▲, left) and relative depth deviation accuracy $\delta < 1.25$ (▲, right) for a voxel size of 20 cm and 40 cm (top and bottom row). Based on the data, we set $p_{FP} = 0.2$ and $p_{FN} = 0.8$ for a voxel size of 20 cm, and $p_{FP} = 0.1$ and $p_{FN} = 0.9$ for a voxel size of 40 cm.

Binary cross-entropy (BCE). The first way to use our training data is to model a binary decision between the occupied belief $m_i(o)$ and free belief $m_i(f)$, which is weighted with the uncertainty term $(1 - m(\Omega))$. We calculate the target probability as:

$$p_i^{\text{tgt}} = \frac{m_i(o)}{m_i(o) + m_i(f)}. \quad (12)$$

The total loss is the weighted sum of BCE values over all voxels i :

$$\mathcal{L} = \sum_i (1 - m_i(\Omega)) \mathcal{L}_{\text{BCE}}(p_i^{\text{tgt}}, p_i^{\text{pred}}). \quad (13)$$

This setup exactly matches our baseline, except that we exchange the visibility mask with our uncertainty weighting and use our per-voxel target probabilities.

Cross-entropy (CE). Alternatively, we can directly estimate the voxel BBA by predicting the mass values of its *occupied*, *free* and *uncertain* hypotheses. Since we model a distribution over the three hypotheses, we choose the softmax activation function, which automatically enforces their sum to be one. Following the

Name	Ground-truth	Predicted classes	Loss	Loss weighting
Baseline w/o fine-tuning	Occ3D [22]	16 occupied + 1 free class	CE	Camera visibility
Baseline	Occ3D [22]	{f, o}	BCE	Camera visibility
Ours BCE	Our BBA	{f, o}	BCE	$1 - m_i(\Omega)$
Ours CE	Our BBA	{f, o, Ω }	CE	$1 - m_i(\Omega)$

Table 3. **Overview of Parameters Used for Training Occupancy Prediction.** We choose BEVDet4D R50 [8] as our baseline. We fine-tune the baseline for the task of binary occupancy prediction using the occupancy maps from Occ3D [22]. Therefore, we replace the last layer to predict a single occupancy value and apply binary cross-entropy as loss function. “Ours BCE” is fine-tuned the exact same way, we only change the supervision signal by using our data. “Ours CE” additionally predicts the uncertainty mass of our BBA and is therefore trained with the cross-entropy loss. For all models, we only change the last layer and the supervision signal, but keep all other parameters the same.

Method	MAE \downarrow in m	RMSE \downarrow in m	RMSE _{log} \downarrow in log m	$\delta < 1.25$ \downarrow in %	$\delta < 1.25^2$ \downarrow in %	$\delta < 1.25^3$ \downarrow in %
Baseline without fine-tuning [7]	1.93	4.30	0.31	74.7	91.5	95.0
Baseline [7]	1.89	4.26	0.30	76.4	92.3	95.4
Ours BCE	1.44	3.60	0.25	88.0	94.0	96.7
Ours CE	1.42	3.47	0.24	88.5	94.5	97.0

Table 4. **Comparison of Depth Errors Between Occupancy Predictions and LiDAR Measurements.** We report errors on rendered depth values from occupancy grids vs. raw LiDAR measurements. We use depth evaluation metrics to quantify the degree of agreement between occupancy map and raw LiDAR measurement, which are calculated on the validation set of nuScenes dataset [2]. We only consider rays which are fully enclosed in the volume specified in Tab. 2.

standard multi-class classification setup, we train the network with the CE loss and additionally weight each term with an uncertainty term ($1 - m(\Omega)$):

$$\mathcal{L} = \sum_i (1 - m_i(\Omega)) \cdot \mathcal{L}_{CE}(\mathbf{p}_i^{\text{tgt}}, \mathbf{p}_i^{\text{pred}}). \quad (14)$$

Note that $\mathbf{p}_i^{\text{tgt}}$ and $\mathbf{p}_i^{\text{pred}}$ are now distributions over *occupied*, *free* and *uncertain* hypotheses rather than single values. A quick overview of the different training modes and supervision signals is given in Tab. 3.

4.5. Occupancy Prediction

We compare the predicted occupancy maps with the LiDAR measurements using the evaluation protocol described in Sec. 4.1. We report quantitative results in Tab. 4. Models trained with our data outperform the baselines across all metrics with respect to the agreement with the LiDAR scans. We notice that fine-tuning for binary occupancy estimation using the occupancy maps from [22] only yields small improvements, whereas fine-tuning with our occupancy maps improves performance significantly. We attribute this performance gain to our improved occupancy maps, since we only changed the supervision signal. The prediction of *occupied*, *free*

and *uncertain* hypotheses instead of a binary decision improves performance even further, but the improvement is small compared to the improvement due to better GT data. Please see the supplementary material for an additional evaluation of our uncertainty estimates for the GT data and model predictions.

5. Conclusion

In this paper we presented a novel approach for ground-truth occupancy data generation from LiDAR measurements. Based on the evidence theory, our approach maps LiDAR point clouds onto discrete 3D grids encoding belief masses of each voxel belonging to one of three possible states: *occupied*, *free* or *uncertain*. We compare our generated ground-truth to those of popular other methods, and demonstrated that our data much better agrees with the LiDAR measurements than that from existing preprocessing strategies. We also showed that occupancy estimation neural networks trained on our data yield superior prediction results without any modifications on the model architecture. An important future research direction includes the incorporation of semantic information into our preprocessing and training pipelines.

References

- [1] Stefan Andreas Baur, David Josef Emmerichs, Frank Moosmann, Peter Pinggera, Björn Ommer, and Andreas Geiger. SLIM: Self-supervised LiDAR scene flow and motion segmentation. In *ICCV*, 2021. 3
- [2] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yuxin Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for autonomous driving. *CVPR*, 2020. 2, 5, 6, 8
- [3] A. P. Dempster. A generalization of bayesian inference. *Journal of the Royal Statistical Society: Series B (Methodological)*, 30(2):205–232, 1968. 5
- [4] Alberto Elfes. Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6):46–57, 1989. 3
- [5] James D Foley. *Computer graphics: principles and practice*. Addison-Wesley Professional, 1996. 1
- [6] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *CVPR*, 2012. 2
- [7] Junjie Huang and Guan Huang. BEVDet4D: Exploit temporal cues in multi-camera 3D object detection. *arXiv preprint arXiv:2203.17054*, 2022. 1, 3, 7, 8
- [8] Junjie Huang, Guan Huang, Zheng Zhu, and Dalong Du. BEVDet: High-performance multi-camera 3D object detection in bird-eye-view. *CoRR*, abs/2112.11790, 2021. 1, 3, 7, 8
- [9] Yuanhui Huang, Wenzhao Zheng, Yunpeng Zhang, Jie Zhou, and Jiwen Lu. Tri-perspective view for vision-based 3D semantic occupancy prediction. In *CVPR*, 2023. 3
- [10] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, page 0, 2006. 2
- [11] Yin hao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and Zeming Li. BEVDepth: Acquisition of reliable depth for multi-view 3D object detection. In *AAAI*, pages 1477–1485, 2023. 2, 3
- [12] Yiming Li, Sihang Li, Xinhao Liu, Moonjun Gong, Kenan Li, Nuo Chen, Zijun Wang, Zhiheng Li, Tao Jiang, Fisher Yu, Yue Wang, Hang Zhao, Zhiding Yu, and Chen Feng. SSCBench: Monocular 3D semantic scene completion benchmark in street views. *arXiv preprint arXiv:2306.09001*, 2023. 2
- [13] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. BEVFormer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *ECCV*, Cham, 2022. 2
- [14] Zhiqi Li, Zhiding Yu, David Austin, Mingsheng Fang, Shiyi Lan, Jan Kautz, and Jose M. Alvarez. FB-OCC: 3D occupancy prediction based on forward-backward view transformation, 2023. 1, 3
- [15] Zhiqi Li, Zhiding Yu, Wenhai Wang, Anima Anandkumar, Tong Lu, and Jose M. Alvarez. FB-BEV: BEV representation from forward-backward view transformations. In *ICCV*, pages 6919–6928, 2023. 3
- [16] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. PETR: Position embedding transformation for multi-view 3D object detection. In *ECCV*, pages 531–548. Springer, 2022. 2
- [17] Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela L Rus, and Song Han. BEV-Fusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2774–2781. IEEE, 2023. 2
- [18] H. P. Moravec. Sensor fusion in certainty grids for mobile robots. *Sensor Devices and Systems for Robotics*, 9(2):253–276, 1989. 3
- [19] Sven Richter, Frank Bieder, Sascha Wirges, and Christoph Stiller. A dual evidential top-view representation to model the semantic environment of automated vehicles. *IEEE Transactions on Intelligent Vehicles*, pages 1–13, 2023. 3, 5
- [20] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott M. Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhiheng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. *CVPR*, 2019. 2, 5, 6
- [21] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3D outputs. In *ICCV*, 2017. 1
- [22] Xiaoyu Tian, Tao Jiang, Longfei Yun, Yue Wang, Yilun Wang, and Hang Zhao. Occ3D: A large-scale 3D occupancy prediction benchmark for autonomous driving. *arXiv preprint arXiv:2304.14365*, 2023. 1, 2, 3, 5, 6, 7, 8
- [23] Wenwen Tong, Chonghao Sima, Tai Wang, Li Chen, Silei Wu, Hanming Deng, Yi Gu, Lewei Lu, Ping Luo, Dahua Lin, and Hongyang Li. Scene as occupancy. In *ICCV*, 2023. 1, 2, 3, 5, 6, 7
- [24] Ignacio Vizzo, Tiziano Guadagnino, Jens Behley, and Cyrill Stachniss. VDBFusion: Flexible and efficient TSDF integration of range sensor data. *Sensors*, 22(3):1296, 2022. 2
- [25] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-CNN: Octree-based convolutional neural networks for 3D shape analysis. *ACM Transactions On Graphics (TOG)*, 36(4):1–11, 2017. 1
- [26] Xiaofeng Wang, Zheng Zhu, Wenbo Xu, Yunpeng Zhang, Yi Wei, Xu Chi, Yun Ye, Dalong Du, Jiwen Lu, and Xingang Wang. OpenOccupancy: A large scale benchmark for surrounding semantic occupancy perception. In *ICCV*, 2023. 1, 2, 5, 6

- [27] Yi Wei, Linqing Zhao, Wenzhao Zheng, Zheng Zhu, Jie Zhou, and Jiwen Lu. SurroundOcc: Multi-camera 3D occupancy prediction for autonomous driving. In *ICCV*, 2023. 2, 5, 6
- [28] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaese-model Pontes, Deva Ramanan, Peter Carr, and James Hays. Argoverse 2: Next generation datasets for self-driving perception and forecasting. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*. Curran, 2021. 2
- [29] T. Yang and V. Aitken. Evidential mapping for mobile robots with range sensors. *IEEE Transactions on Instrumentation and Measurement*, 55:1422–1429, 2006. 3
- [30] Yunpeng Zhang, Zheng Zhu, and Dalong Du. OccFormer: Dual-path transformer for vision-based 3D semantic occupancy prediction. In *ICCV*, 2023. 1