

Deep Inverse Sensor Models as Priors for evidential Occupancy Mapping

Daniel Bauer¹, Lars Kuhnert¹ and Lutz Eckstein²

Abstract—With the recent boost in autonomous driving, increased attention has been paid on radars as an input for occupancy mapping. Besides their many benefits, the inference of occupied space based on radar detections is notoriously difficult because of the data sparsity and the environment dependent noise (e.g. multipath reflections). Recently, deep learning-based inverse sensor models, from here on called deep ISMs, have been shown to improve over their geometric counterparts in retrieving occupancy information [1], [2], [3]. Nevertheless, these methods perform a data-driven interpolation which has to be verified later on in the presence of measurements. In this work, we describe a novel approach to integrate deep ISMs together with geometric ISMs into the evidential occupancy mapping framework. Our method leverages both the capabilities of the data-driven approach to initialize cells not yet observable for the geometric model effectively enhancing the perception field and convergence speed, while at the same time use the precision of the geometric ISM to converge to sharp boundaries. We further define a lower limit on the deep ISM estimate’s certainty together with analytical proofs of convergence which we use to distinguish cells that are solely allocated by the deep ISM from cells already verified using the geometric approach.

I. INTRODUCTION

For robotic applications, occupancy maps are a long established form of environment representation [4], [5], [6]. Because of their robustness towards environmental conditions, radars are heavily relied on to obtain the occupancy information for automotive applications [7]. However, radar data often comes in the form of sparse, noisy detections, based on which it is non-trivial to infer the occupancy state [8]. Methods to infer the environment based on the measurement are referred to as inverse sensor models (ISMs) [6] which can be divided into the categories of data-driven and geometric approaches. The geometric ISMs use the sensor’s measurement principle to define a geometric correlation between the sensor position and the detection. On the other hand, data-driven models define a machine learning model and use the measurements themselves to tune the ISM. In recent years, many deep learning approaches, from here on referred to as deep ISMs, have been proposed to handle the sparseness of radar data [1], [2], [3]. For the fusion of predictions over time, it is important to know that geometric ISMs only infer the state of cells which are in a geometric correlation between sensor and detection. Hence, the information provided by two separate measurements is largely independent. In contrast to that, deep ISMs do also infer the occupancy state of areas

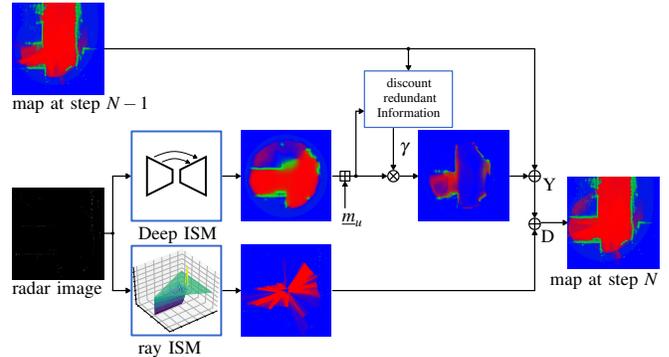


Fig. 1. Overview of approach to integrate deep ISM estimates into the evidential occupancy mapping framework, by accounting for redundant information (rgb \rightarrow {free, occupied, unknown} || best viewed in color with zoom).

far away from actual detections. This leads to situations in which two independent measurements can provide mostly the same information about a cell’s state. This results in a high amount of redundant information between time steps which, if not accounted for, accumulates and leads to wrong convergence. To deal with this redundancy, we propose the following. First, limit the certainty of the deep ISM estimate to a tuneable threshold m_u . Afterwards, discount the deep ISM prediction using a discount factor γ depending on the amount of redundant information between the current map and the deep ISM. Eventually, combine the non-redundant part of the deep ISM with the map and the geometric ISM’s prediction. This procedure is depicted in Fig. 1.

In summary, our contributions are first, the identification of redundant information as the cause of wrong convergence when integrating deep ISM predictions in the same way as geometric estimates. And secondly, an analytically verified procedure to account for the redundant information in a way that cell states solely allocated by a deep ISM have a lower bound m_u on their certainty which makes them distinguishable from pixels verified by the geometric ISM. Additionally, only new information is added in each time step, limiting the accumulation of small evidences. Using this integration scheme, we are able to leverage both the capabilities of deep ISMs to enhance the perception field and of geometric ISMs to converge to sharp boundaries.

II. RELATED WORK

A. Occupancy Mapping

Occupancy maps are a common representation of a robots environment [9], [10], [4], where the space around the

¹Daniel Bauer and Lars Kuhnert are with Ford Werke GmbH, Germany {dbauer31, lkuhnert}@ford.com

²Lutz Eckstein is with the Institute of Automotive Engineering, RWTH Aachen University, Germany office@ika.rwth-aachen.de

robot is discretized into a bird's-eye-view grid with each grid cell containing its occupancy information. The two dominant forms of encoding this information is as probabilities, proposed by Elfes [4], or in the Dempster-Shafer evidence framework [5], [11]. In contrast to the probabilistic formulation, the evidential representation uses an additional degree of freedom which allows to distinguish conflicting information between free and occupied from the state of missing information. This conflict can e.g. be used to indicate dynamic objects [12]. On the other hand, the amount of missing information correlates with the degree of exploration which we will exploit in this work to indicate whether pixels have only been allocated using the deep ISM or have also been verified using the geometric ISM.

B. Combination of independent Evidence

One central question in the evidential framework is how to combine two evidence masses. Under the assumption of informational independence, several combination rules have been proposed which fall into the categories of associative and non-associative rules [13]. The two dominant methods of both categories are the associative Dempster's rule [11] and the non-associative Yager's rule [14]. In our case, the non-associativity of Yager's rule is acceptable, since we always fuse information from the ISMs into the occupancy maps, thus giving us a clear order. In sec. IV-A, we analyse the pros and cons of both methods and explain in which scenario to take which rule.

C. Combination of dependent Evidence

In this work, we argue that the sources of evidence between two deep ISM predictions are not independent for each pixel. Hence, using the above mentioned combination rules results in unwanted effects, which will be elaborated on in sec. IV. However, a number of methods have been proposed to handle the combination of dependent evidences. These methods address the problem by either altering the combination rule itself or discounting the masses directly according to the amount of redundant information [15]. Here, we approach the problem by discounting the masses.

To quantize the amount of redundancy, several measures have been proposed ranging from the relative independence degree [16] to the Pearson [17] or Spearman rank correlation coefficient [18] through to using the mutual information between the two evidence masses [15]. Our approach is closely inspired by [15]. However, in sec. IV-D, we argue that using the unknown mass instead of the entropy as a measure of information fits better into the evidential occupancy mapping framework.

D. Geometric Inverse Sensor Models

To model the inverse relationship between the environment and the sensor, one approach is to use geometric information inherent in the measurement principle. In our case, we use lidar and radar sensors which both have a polar measurement scheme. This means that ideally, everything in a line between the sensor and the detection is free and the detection itself

is occupied. In case of the lidar, the data is accurate enough to apply the ideal sensor model as shown in Fig. 2. When, however, applied on sparse radar data, the ideal sensor model results in too sparse free space. To account for this, the ideal model can be convolved with e.g. a Gaussian noise [19], [3]. This not only enlarges the perceived free space but also accounts for detection inaccuracies. For our radar model, we build upon the work of [19] and adapt it in a way to produce sharper occupied space which is elaborated on in sec. V-B.

E. Data-driven Inverse Sensor Models

Recently, many deep learning approaches have been proposed to model the inverse radar characteristics. One half of the methods defines the inverse radar characteristic as a probabilistic model [1], [3] while the other half models it in the evidential framework [2], [3]. Since we choose to use the evidential occupancy mapping framework, we will only consider the second half. In this category, Sless et al. [2] proposed to train a network using accumulated radar images as input and inverse lidar models as target. The network uses a softmax with cross entropy loss to predict the three categories free, occupied and unknown. Another approach proposed by Wirges et al. [20] models the inverse lidar model using a multichannel input with detections, transmissions and intensity. As labels, they build occupancy maps and use patches around each vehicle pose to better capture the ground-truth occupancy state. They also use a softmax, however, train on L_1 and L_2 losses with a specific scaling factor to reduce the influence in presence of label uncertainty. Recently, we have proposed an evidential model [3] which is capable of learning the aleatoric uncertainty and encode it into the unknown mass. We compared it against a model similar to the two mentioned before and could demonstrate some improvements with regards of robustness and segmentation sharpness. Hence, we propose to use this model also in this work. In sec. V-C, we will provide an overview and show some exemplary prediction results.

III. PRELIMINARIES

A. Evidential Occupancy Mapping States

In this section, we present a brief review of the evidential occupancy mapping framework [21], [5]. For occupancy mapping, the state of a cell can either be free f or occupied o . Hence, the set of possible states called frame of discernment u , which will be treated as the unknown state, and its power set 2^u can be defined as

$$u = \{f, o\}, s^u = \{\emptyset, f, o, u\} \quad (1)$$

Evidences are represented in the form of mass functions \mathbf{m} that define the evidence mass for each state as follows

$$\mathbf{m} = [m_f, m_o, m_u]^\top \quad (2)$$

$$m_\emptyset = 0, m_k \in [0, 1] \text{ with } k \in \{f, o, u\} \quad (3)$$

$$\sum_k m_k = 1 \quad (4)$$

with the evidence mass for free m_f , occupied m_o and unknown m_u .

B. Evidential combination Rules

In this work, we only focus on the application of Dempster's and Yager's combination rules. For occupancy mapping, these can be written as follows

Dempster's Rule

$$\mathbf{m}_1 \oplus_D \mathbf{m}_2 = \begin{bmatrix} (m_{f1}m_{f2} + m_{f1}m_{u2} + m_{u1}m_{f2})/(1-K) \\ (m_{o1}m_{o2} + m_{o1}m_{u2} + m_{u1}m_{o2})/(1-K) \\ m_{u1}m_{u2}/(1-K) \end{bmatrix} \quad (5)$$

Yager's Rule

$$\mathbf{m}_1 \oplus_Y \mathbf{m}_2 = \begin{bmatrix} m_{f1}m_{f2} + m_{f1}m_{u2} + m_{u1}m_{f2} \\ m_{o1}m_{o2} + m_{o1}m_{u2} + m_{u1}m_{o2} \\ m_{u1}m_{u2} + K \end{bmatrix} \quad (6)$$

with the conflict K being defined as

$$K = m_{f1}m_{o2} + m_{o1}m_{f2} \quad (7)$$

C. Discounting Evidence

Given a discounting factor $\gamma \in [0, 1]$, the discounting rule can be formulated as follows

$$\gamma \otimes \mathbf{m} = [\gamma m_f, \gamma m_o, 1 - \gamma + \gamma m_u]^\top \quad (8)$$

D. Subjective Logic Transformation

Josang [22] proposed a framework which collects evidences $e \in \mathbb{R}_{>0}$ for all but the unknown class which are not equal to evidence masses. These evidences can be used to either define a probabilistic Dirichlet or an evidential mass function. He also includes transformations between the two representations. The formulas relevant for occupancy mapping are summarized below

Evidence Vector

$$\mathbf{e} = [e_f, e_o]^\top \quad (9)$$

$$S = 2 + e_f + e_o \quad (10)$$

Probabilistic View

$$\boldsymbol{\alpha} = \mathbf{e} + I \quad (11)$$

$$\mathbf{p} = \begin{bmatrix} p_f \\ p_o \end{bmatrix} = \mathbb{E}[\text{Dir}(\boldsymbol{\alpha})] = \frac{\mathbf{e} + I}{S} \quad (12)$$

$$p_o = \frac{m_u}{2} + m_o = \frac{1 - m_f + m_o}{2} \quad (13)$$

Evidential View

$$\mathbf{m} = [e_f/S, e_o/S, 2/S]^\top \quad (14)$$

E. Lower Limit on Unknown Mass

We use the operator \boxplus to ensure that a mass function \mathbf{m} has at least \underline{m}_u unknown mass assigned as follows

$$\mathbf{m} \boxplus \underline{m}_u = \begin{bmatrix} (1 - \alpha)m_f \\ (1 - \alpha)m_o \\ m_u + \Delta m_u \end{bmatrix} \quad (15)$$

$$\text{with } \Delta m_u = \max(0, \underline{m}_u - m_u), \alpha = \frac{\Delta m_u}{m_f + m_o} \quad (16)$$

IV. METHOD

A. Choice of Combination Rule

To chose an adequate combination rule, we first have to examine the properties of Dempster's and Yager's rules. Dempster's rule is defined in a way that the fused unknown mass can only be lesser or equal to the separate unknown masses

$$\max(m_{u1}, m_{u2}) \geq m_{u12} \quad (17)$$

Since Dempster's rule is associative, we only need to show this property with respect to one of the in-going masses

$$m_{u1} \geq m_{u12} = m_{u1} \frac{m_{u2}}{1-K} \quad |\text{with (5)} \quad (18)$$

$$\Leftrightarrow 1 \geq \frac{1 - m_{f2} - m_{o2}}{1 - m_{f1}m_{o2} - m_{o1}m_{f2}} \quad |\text{with (7), (4)} \quad (19)$$

$$\Leftrightarrow 0 \geq m_{f2} \underbrace{(m_{o1} - 1)}_{\leq 0} + m_{o2} \underbrace{(m_{f1} - 1)}_{\leq 0} \quad \blacksquare \quad (20)$$

On the other hand, Yager's rule adds the conflict K to the fused mass, hence allowing to recuperate unknown mass (see (6)).

Therefore, we propose to use Dempster's combination rule to fuse mass from the ray ISM since it ensures that, once the unknown mass has fallen below the threshold, it remains beneath it. This prohibits further changes through the deep ISM. On the other hand, we propose to use Yager's rule to fuse mass from deep ISMs, since it recuperates unknown mass in case of conflicting predictions. This allows us for falsely initialized cells to first move the falsely assigned mass back to the unknown state and from there shift it to the correct state without violating the lower bound \underline{m}_u .

B. Informational Interpretation of deep ISM Predictions

In each step, we want each cell to be allocated with its true evidence mass $\hat{\mathbf{m}}$ given all the measurements z observed so far. We train a neural network to predict $\tilde{\mathbf{m}}$, which approximates the true occupancy state at time step N . As input for the network, we transform the last L measurements z_{N-L}, \dots, z_N into the frame at time step N to obtain the accumulated input $z_{N-L:N}$. By doing so, we propose the following two approximations

$$\mathbf{m}(z_{N-L}, \dots, z_N) \approx \mathbf{m}(z_0, \dots, z_N) \quad (21)$$

$$\mathbf{m}(z_{N-L:N}) \approx \mathbf{m}(z_{N-L}, \dots, z_N) \quad (22)$$

First, we argue in (21) that all time steps earlier than $N-L$ do not contain additional relevant information for the grid cells being updated at step N . Secondly, by performing the approximation in (22), we loose the information of where the detections have been recorded, thus, arguing that the detection positions contain the main amount of measurement information. Therefore, in theory, the following should hold

$$\hat{\mathbf{m}}(z_0, \dots, z_n) \approx \tilde{\mathbf{m}}(z_{N-L:N}) \quad (23)$$

and we could update the map e.g. by replacing the cell state whenever the predictions unknown mass is smaller than the current cell state. With this approach, however,

we would always overwrite the geometric ISM's predictions while the cell state is still below \underline{m}_u , thus prohibiting a proper information fusion. Moreover, the approach removes the filtering aspect of occupancy mapping which results in outliers with small unknown masses being directly allocated to cells, as can be seen in Fig. 4. This falsely assigned mass has then to be reallocated by the geometric method which in turn decreases the convergence speed. To reduce the influence of outliers, we propose an iterative update scheme in the following sections which both reduces falsely and increases correctly assigned mass, as shown in Tab. I.

C. Identification of non-redundant Information

Instead of using the neural network predictions to replace the cell states, we can also use the predictions to update it. However, there is a lot of redundant information between $\tilde{\mathbf{m}}(z_{N-L:N})$ and $\mathbf{m}(z_0, \dots, z_{N-1})$. Even, if we only use one measurement z_N as network input, there would still be redundancy since the last couple of inputs already measured approximately the same environment. Using the accumulated input $z_{N-L:N}$, however, provides us with the possibility to assess the amount of non-redundant information $I(z_N|z_0, \dots, z_{N-1})$ of the latest measurement z_N , as follows

$$I(z_N|z_0, \dots, z_{N-1}) = I(z_0, \dots, z_N) - I(z_0, \dots, z_{N-1}) \quad (24)$$

$$I(z_N|z_0, \dots, z_{N-1}) \approx I(z_{N-L:N}) - I(z_0, \dots, z_{N-1}) \quad (25)$$

To account for the redundancy, we can use the identified amount of additional information to discount the network predictions with a discount factor $\gamma(I(z_N|z_0, \dots, z_{N-1}))$. The only remaining question is which quantity should be used as a measure of information $I(\dots)$ and how to define the discount factor based on that measure.

D. Choice of Information Measure and Discount Factor

For this work, we use the unknown mass as a measure of information content. The benefit over the usually used entropy [15] is that the case of total conflict, where all mass is equally distributed into free and occupied mass, is not viewed as total lack of information. Hence, we can still use the conflict case to represent dynamic objects with varying degrees of certainty based on the unknown mass. Therefore, we will write the non-redundant part of the prediction information as

$$\Delta m_u^N = \tilde{m}_u^N - m_u^{N-1} \quad (26)$$

With regards to the discount factor, we have two conditions. First, we want to reduce the redundant information in $\tilde{\mathbf{m}}$ so that the following holds

$$\gamma = \begin{cases} 0, \Delta m_u < 0 \\ 0, \Delta m_u \rightarrow 0 \\ 1, \Delta m_u \rightarrow 1 \end{cases} \quad (27)$$

To realize this relationship, we propose to use

$$\gamma = \tanh(\alpha \max(0, \Delta m_u)) \quad (28)$$

which allows to tune the convergence speed using the parameter α . For our experiments, we choose $\alpha = 10$. If we first limit the network prediction's unknown mass to \underline{m}_u using (15) and afterwards discount $\tilde{\mathbf{m}}$ using (28), we ensure that, after the cell state reached the unknown mass limit, it cannot be altered through $\tilde{\mathbf{m}}$. This, however, does not ensure that the cell state's unknown mass cannot be pushed below \underline{m}_u before the limit is reached. Thus, as a second condition, we want to choose the discount factor γ in a way that the fused unknown mass cannot fall beneath the given threshold \underline{m}_u , which can be written as follows

$$\mathbf{m}^N = \mathbf{m}^{N-1} \oplus_{\gamma} (\gamma^N \otimes \tilde{\mathbf{m}}^N) \quad (29)$$

$$\underline{m}_u \leq m_u^N \quad (30)$$

Using Yager's update rule and (8), we can define a lower bound on γ

$$\underline{m}_u \leq m_u^{N-1} (1 - \gamma^N + \gamma m_2^N) + K^N \quad (31)$$

In case (30) holds, the following is also given

$$\underline{m}_u \leq m_u^{N-1} \quad (32)$$

$$\text{and for } \tilde{m}_u^N = 1 \rightarrow m_u^N = m_u^{N-1} + K^N \geq \underline{m}_u \quad (33)$$

Under these conditions, we can rewrite (31) as follows

$$\gamma \leq \frac{m_u^{N-1} + K - \underline{m}_u}{m_u^{N-1} (1 - \tilde{m}_u^N)} \quad (34)$$

which provides an additional upper bound on γ . Eventually, the final discount factor is defined as

$$\gamma^N = \min \left(\frac{m_u^{N-1} + K - \underline{m}_u}{m_u^{N-1} (1 - \tilde{m}_u^N)}, \tanh(\alpha \max(0, m_u^{N-1} - \tilde{m}_u^N)) \right) \quad (35)$$

V. EXPERIMENTAL SETUP

For our experiments, we use the publicly available nuScenes dataset [7] from which we use the data of the 360°, roof mounted, 32 beam lidar, 5 radars mounted at the corners and one central at the front and dead reckoning lidar odometry. For all mapping, we first projected the inputs into a 512 × 512 bird's-eye-view grid corresponding to an area of 40 × 40 meters. Afterwards, one of the ISMs is applied on the detection image followed by a transformation, using the lidar odometry, back into the first measurements coordinate frame. For training and validation, we used the same data split as proposed by the nuScenes team, while removing scenes in which the ego vehicle does not move further than 1 meter. This results in 133708 training and 29846 validation samples.

A. Inverse Lidar Model

To compute the ray casting-based inverse lidar model, from now on referred to as ray ILM, several steps are performed. First, we remove all detections below 0.3m and above 3.0m. The remaining points are then projected on the ground plane and discretized into the image grid. Eventually, rays are cast in all directions starting at the lidar coordinate frame's center with a resolution of 0.2° up until a detection is hit or the maximal range (15 meters in our experiments) is reached. Finally, detection positions are assigned $m_o = 0.5$ and the ray points $m_f = 0.05$, as shown in Fig. 2.

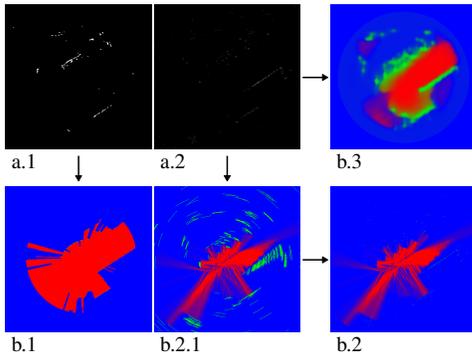


Fig. 2. Illustration of the lidar detections a.1 and the corresponding ray ILM b.1, the radar detections a.2, the intermediate and final ray IRM b.2.1, b.2 and the deep IRM b.3 (best viewed in color with zoom)

B. Inverse Radar Model

For the ray casting-based inverse radar model (ray IRM), the last $N = 10$ radar point clouds from each radar are projected into the current ego vehicle frame. Afterwards, cones with an opening angle of 2° are cast from the respective temporal sensor positions towards the corresponding detections. As an intermediate step, the segment of each cone at max range is marked as occupied. Next, the same cones are cast again towards the detections, marking each pixel $m_f = 0.3$. This time, however, the cones are restricted by the occupied segments, as shown in Fig. 2, b.2.1. Eventually, the occupied segments are removed and the last N detections are marked with $m_o = 0.5$ (Fig. 2, b.2).

C. Deep Inverse Radar Model

As a deep inverse radar model (deep IRM), we train Ev-Net as described in [3] to predict the ray ILM based on accumulated radar detection images. We adapt the architecture slightly to cope with the high resolution images by max pooling the inputs and bilinear upsampling the output. Moreover, we do not use occupancy map patches to train the model, since the results indicate only minor benefit. Ev-Net can be used to predict strictly positive evidences for free and occupied space $\tilde{\mathbf{e}} = [\tilde{e}_f, \tilde{e}_o]^\top$ and use them to either obtain a probabilistic or evidential occupancy state prediction (see (9) - (14)). The network, whose architecture is depicted in Fig. 3, can be trained using the following loss

$$\mathcal{L}_k = \sum_k (\hat{m}_k - \tilde{p}_k)^2 \cdot \frac{\tilde{p}_k(1 - \tilde{p}_k)}{S + 1}, \quad \mathcal{L}_u = (1 - \tilde{m}_u)^2 \quad (36)$$

To account for the class imbalance between the amount of free, occupied and unknown space, the expectation over each class is computed separately and afterwards accumulated to obtain the overall loss

$$\mathcal{L} = \mathbb{E}[\mathcal{L}_f] + \mathbb{E}[\mathcal{L}_o] + \mathbb{E}[\mathcal{L}_u] \quad (37)$$

An exemplary prediction of the network is shown in Fig. 2, b.3. The network runs without further optimizations with over 10 Hz on one core of an i7 8th gen intel CPU.

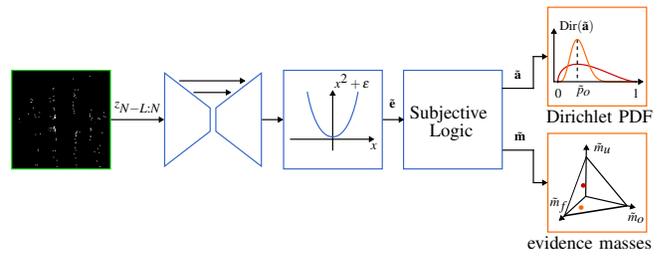


Fig. 3. Architecture of Ev-Net as proposed in [3]

VI. EXPERIMENTAL RESULTS

In this section, we compare occupancy maps which are created by accumulating

- 1) ray ILM with Dempster's rule as ground-truth reference
- 2) deep IRM with Dempster's rule (deep IRM (acc.))
- 3) deep IRM with lower limit \underline{m}_u and replacement of cells with lower unknown mass (deep IRM (replace))
- 4) deep IRM with lower limit \underline{m}_u and discounted redundant information using Yager's rule (deep IRM (disc.))
- 5) ray IRM with Dempster's rule
- 6) fusion of deep IRM (disc.) and ray IRM

A. Interpretation of qualitative Results

The exemplary occupancy mapping result using deep IRM (acc.) in Fig. 4 reflects the problem that small biases accumulate over time in occluded areas (e.g. box 2)). In our case, the model has a small bias towards the occupied state which can be seen by the slightly green circle in Fig. 2, b.3. This problem is solved in our proposed framework by only accumulating the non-redundant part of the information, as can be seen by the mapping results for fused IRM.

Moreover, box 1) illustrates an area where correctly allocated areas close to data can be later on falsely overwritten by purely interpolation-based estimates. To prevent this behavior, we set a lower limit on the deep IRM's unknown mass and only allow the ray IRM to decrease the unknown mass further. Since we only allow the deep IRM to allocate non-redundant information, we therefore prohibit any contributions of the deep IRM once this lower unknown threshold is reached. Hence, areas cannot be falsely overwritten later on, as can be seen in the occupancy map of fused IRM box 1).

Box 3) shows that the accuracy at object boundaries for fused IRM is similar to ray IRM's as opposed to the results for deep IRM (acc.).

Finally, box 4) shows an area which is already allocated for the fused IRM's occupancy map but only partially observed by the ray IRM's. This illustrates our frameworks capability to use the deep IRM as a prior to initialize areas based on data-driven interpolation. Additionally, the influence of the lower unknown limit on deep IRM can be seen by the lighter red in the regions solely allocated by the deep IRM.

B. Interpretation of quantitative Results

The per class mean intersection over union scores in Tab. I verify the problems when applying the deep IRM directly

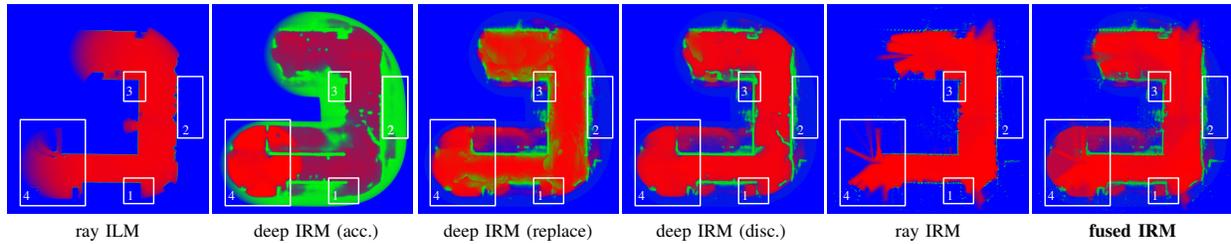


Fig. 4. Qualitative comparison of occupancy mapping results for the investigated fusion approaches with their respective ISMs. Box 1) shows an exemplary area where the deep IRM (acc.) maps were correctly initialized in the beginning but later on falsely overwritten. Box 2) illustrates an area where deep IRM (acc.) accumulates small biases in occluded areas. Both of these problems are solved by our framework in fused IRM. Additionally, box 3) shows that the fused IRM’s accuracy at object boundaries is similar to the ray IRM’s as opposed to deep IRM (acc.). Eventually, box 4) marks an area which is already initialized by the deep IRM but only partially observed by the ray IRM.

TABLE I

PER CLASS MEAN INTERSECTION OVER UNION SCORES			
ISM used for Mapping	fr. mIoU	oc. mIoU	un. mIoU
deep IRM (acc.)	61.78	2.07	11.05
deep IRM (replace)	69.63	3.43	45.96
deep IRM (disc.)	71.64	4.75	46.23
ray IRM	76.68	6.02	63.55
fused IRM	76.88	5.44	45.69

in the standard evidential occupancy mapping framework by iteratively fusing it into the map using Dempster’s rule of combination. Additionally, the scores show that our proposed fusion scheme converges to the ray IRM’s performance in the presence of sufficient measurements.

VII. CONCLUSIONS

We propose an integration procedure which only uses the non-redundant information of deep ISMs to update occupancy maps. The update is performed in a way that, when solely using the deep ISM, the map state’s certainty cannot fall beneath a given threshold. This enables us to initialize the map in areas not observed by the geometric ISM and, thus, increase the overall perceptive field and convergence speed. Moreover, we are still able to distinguish cells solely allocated by the deep ISM and not yet verified using the geometric ISM by the tuneable certainty threshold. Additionally, in presence of enough measurements, the map converges to the geometric ISM which results in sharper borders between free and occupied space.

REFERENCES

- [1] R. Weston, S. Cen, P. Newman, and I. Posner, “Probably unknown: Deep inverse sensor modelling in radar,” *arXiv preprint arXiv:1810.08151*, 2018.
- [2] L. Sless, B. El Shlomo, G. Cohen, and S. Oron, “Road scene understanding by occupancy grid learning from sparse radar clusters using semantic segmentation,” in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2019, pp. 0–0.
- [3] D. Bauer, L. Kuhnert, and L. Eckstein, “Deep, spatially coherent inverse sensor models with uncertainty incorporation using the evidential framework,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 2490–2495.
- [4] A. Elfes, “Using occupancy grids for mobile robot perception and navigation,” *Computer*, no. 6, pp. 46–57, 1989.
- [5] D. Pagac, E. M. Nebot, and H. Durrant-Whyte, “An evidential approach to map-building for autonomous vehicles,” *IEEE Transactions on Robotics and Automation*, vol. 14, no. 4, pp. 623–629, 1998.
- [6] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.
- [7] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nuscenes: A multimodal dataset for autonomous driving,” *arXiv preprint arXiv:1903.11027*, 2019.
- [8] M. A. Richards, J. Scheer, W. A. Holm, and W. L. Melvin, *Principles of modern radar*. Citeseer, 2010.
- [9] H. Carrillo, P. Dames, V. Kumar, and J. A. Castellanos, “Autonomous robotic exploration using occupancy grid maps and graph slam based on shannon and rényi entropy,” in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 487–494.
- [10] F. Fleuret, J. Berclaz, R. Lengagne, and P. Fua, “Multicamera people tracking with a probabilistic occupancy map,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 30, no. 2, pp. 267–282, 2008.
- [11] A. P. Dempster, “A generalization of bayesian inference,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 30, no. 2, pp. 205–232, 1968.
- [12] J. Moras, V. Cherfaoui, and P. Bonnifait, “Moving objects detection by conflict analysis in evidential grids,” in *Intelligent Vehicles Symposium (IV), 2011 IEEE*. IEEE, 2011, pp. 1122–1127.
- [13] H. Bostrom, R. Johansson, and A. Karlsson, “On evidential combination rules for ensemble classifiers,” in *2008 11th international conference on information fusion*. IEEE, 2008, pp. 1–8.
- [14] R. R. Yager, “On the dempster-shafer framework and new combination rules,” *Information sciences*, vol. 41, no. 2, pp. 93–137, 1987.
- [15] X. Su, L. Li, F. Shi, and H. Qian, “Research on the fusion of dependent evidence based on mutual information,” *IEEE Access*, vol. 6, pp. 71 839–71 845, 2018.
- [16] R. R. Yager, “On the fusion of non-independent belief structures,” *International journal of general systems*, vol. 38, no. 5, pp. 505–531, 2009.
- [17] X. Su, P. Xu, S. Mahadevan, and Y. Deng, “On consideration of dependence and reliability of evidence in dempster-shafer theory,” *JOURNAL OF INFORMATION & COMPUTATIONAL SCIENCE*, vol. 11, no. 14, pp. 4901–4910, 2014.
- [18] F. Shi, X. Su, H. Qian, N. Yang, and W. Han, “Research on the fusion of dependent evidence based on rank correlation coefficient,” *Sensors*, vol. 17, no. 10, p. 2362, 2017.
- [19] K. Werber, M. Rapp, J. Klappstein, M. Hahn, J. Dickmann, K. Dietmayer, and C. Waldschmidt, “Automotive radar gridmap representations,” in *Microwaves for Intelligent Mobility (ICMIM), 2015 IEEE MTT-S International Conference on*. IEEE, 2015, pp. 1–4.
- [20] S. Wirges, C. Stiller, and F. Hartenbach, “Evidential occupancy grid map augmentation using deep learning,” in *2018 IEEE intelligent vehicles symposium (IV)*. IEEE, 2018, pp. 668–673.
- [21] G. Shafer, “Dempster-shafer theory,” *Encyclopedia of artificial intelligence*, pp. 330–331, 1992.
- [22] A. Jøsang, “Subjective logic: A formalism for reasoning under uncertainty, ser,” *Artificial Intelligence: Foundations, Theory and Algorithms*. Springer International Publishing Switzerland, 2016.