



HAL
open science

Allo-centric Occupancy Grid Prediction for Urban Traffic Scene Using Video Prediction Networks

Rabbia Asghar, Lukas Rummelhard, Anne Spalanzani, Christian Laugier

► **To cite this version:**

Rabbia Asghar, Lukas Rummelhard, Anne Spalanzani, Christian Laugier. Allo-centric Occupancy Grid Prediction for Urban Traffic Scene Using Video Prediction Networks. ICARCV 2022 - 17th International Conference on Control, Automation, Robotics and Vision, Dec 2022, Singapore, Singapore. hal-03862095

HAL Id: hal-03862095

<https://inria.hal.science/hal-03862095>

Submitted on 21 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Allo-centric Occupancy Grid Prediction for Urban Traffic Scene Using Video Prediction Networks

Rabbia Asghar¹, Lukas Rummelhard¹, Anne Spalanzani¹, Christian Laugier¹

Abstract—Prediction of dynamic environment is crucial to safe navigation of an autonomous vehicle. Urban traffic scenes are particularly challenging to forecast due to complex interactions between various dynamic agents, such as vehicles and vulnerable road users. Previous approaches have used ego-centric occupancy grid maps to represent and predict dynamic environments. However, these predictions suffer from blurriness, loss of scene structure at turns, and vanishing of agents over longer prediction horizon. In this work, we propose a novel framework to make long-term predictions by representing the traffic scene in a fixed frame, referred as allo-centric occupancy grid. This allows for the static scene to remain fixed and to represent motion of the ego-vehicle on the grid like other agents'. We study the allo-centric grid prediction with different video prediction networks and validate the approach on the real-world Nuscenes dataset. The results demonstrate that the allo-centric grid representation significantly improves scene prediction, in comparison to the conventional ego-centric grid approach.

Index Terms— Scene Prediction, Deep Learning, Autonomous Vehicles

I. INTRODUCTION

Prediction of traffic scene evolution is essential to an autonomous vehicle for planning as well as detecting dangerous situations. In urban traffic scenarios, the vehicles not only interact with other vehicles, but also share space with vulnerable road users such as pedestrians and cyclists. Key challenges involve the uncertainty and multi-modality of the behaviour of agents in the environment, and complex multi-agents interactions [1]. While human drivers show superior ability to forecast the agents' behaviour and interactions in such traffic scenes, it remains a challenge for autonomous vehicles.

Data-driven methods provide powerful tools to solve prediction problems, particularly dealing with complex social interactions [2]. Most conventional approaches are object or agent-based and rely on heavily pre-processed data [3], [4]. Dynamic Occupancy Grip Maps (DOGMs), on the other hand, allow for end-to-end learning due to their discretized spatial representation, without higher-level segmentation [5]. Additionally, DOGMs are versatile in terms of sensor dependency, and can be generated from a variety of raw sensor data, such as Lidar or camera images. In our work, we use Bayesian-filter-based DOGM [6] that provide us with a spatially-dense model representation of static and dynamic space, as well as free and unknown space in the environment, as shown in Fig1.

As the DOGM is generated using data from the vehicle-mounted sensors, the grid is traditionally ego-centric, i.e. the position of ego-vehicle is fixed in the grid. While this is an effective method in scene representation, it complicates the long-term prediction problem. For a dynamic ego-vehicle, the complete scene translates and/or rotates around the ego-vehicle, even the static components in the scene. Therefore, the prediction network must transform every cell in the grid, leading to blurry and vanishing static scene at longer prediction time horizons.

To address this, we instead generate DOGMs with respect to a fixed reference frame, referred as **allo-centric grid**. While the observed space around the ego-vehicle remains the same, the static scene structure in the allo-centric grid remains fixed. This is illustrated in Fig. 1 where the ego-vehicle is encircled, the vehicle moves like other agents in the scene.

We approach the long-term multi-step predictions of allo-centric DOGM as a video prediction problem due to the inherent similarities between an image and an occupancy grid, and both being a spatio-temporal problem [7]. Results incorporating different video prediction networks are studied, including state-of-the-art recurrent neural networks and memory-augmented network approaches. We compare and evaluate the prediction results of allo-centric and ego-centric grids for identical scenes and demonstrate the superior performances of the allo-centric grid predictions.

The proposed approach is validated with the real-world NuScenes dataset [3] of urban traffic scenes. We show that allo-centric grids significantly improve the prediction results and demonstrate the ability to retain the scene structure and learn behaviours.

The paper is organized as follows. Section II discusses related work to video and scene predictions. Section III describes the system overview. Section IV and V present implementations, results and analysis. Finally conclusions are drawn in section VI.

II. RELATED WORK

A. Video Prediction

Spatio-temporal deep-learning methods have been effectively used for video prediction problems. Commonly, combinations of Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) are incorporated. CNNs are capable of extracting spatial information and capturing inter-dependencies of the surrounding pixels while RNNs, such as long short-term memory (LSTM) blocks,

¹ Univ. Grenoble Alpes, Inria, 38000 Grenoble, France, email: First-Name.LastName@inria.fr

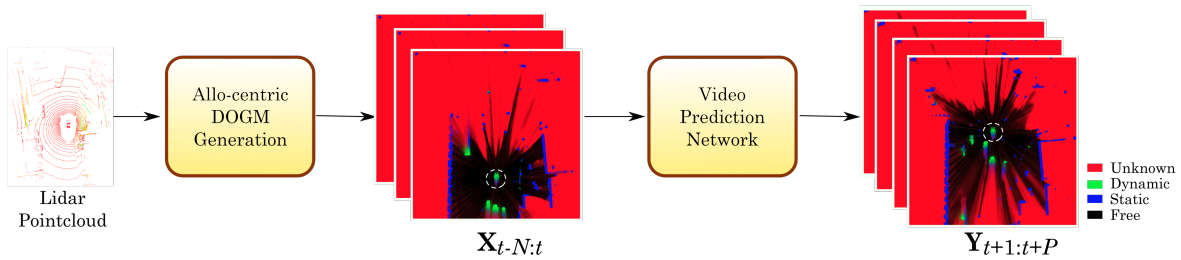


Fig. 1: Overview of our proposed approach. The allo-centric DOGM is represented as an image. Each channel red, green and blue represent unknown, dynamic and static cells respectively. The black space represents known free space. The ego-vehicle is circled in dotted line in both input and target output sequences.

capture the sequential or temporal dependencies. Lotter et al. proposed Predictive Coding Network (PredNet), a deep learning network architecture that comprises of vertically-stacked Convolutional LSTMs (ConvLSTMs) where the local error and the prediction signal are propagated bottom-up and top-down respectively [8]. Wang et al. addresses the video prediction challenges of capturing short-term and long-term dynamics with the PredRNN architecture [9]. Building on their original approach [10], they introduce memory-decoupled spatio-temporal LSTM (ST-LSTM) blocks, feature zigzag memory flow and a novel curriculum learning strategy to improve prediction results. Kim et al. takes inspiration from memory-augmented networks to use external memory (LMC-Memory) to learn and store long-term motion dynamics and propose a memory query decomposition to address the high-dimensionality of motions in video predictions [11].

B. Occupancy Grid Prediction

Jeon et al. proposed conventional ConvLSTM to predict interaction-intensive traffic scenes on occupancy grids [12]. The approach represents only vehicles in the occupancy grid, their states extracted from camera inputs. Desquaire et al. [13], proposed an end-to-end object-tracking approach by incorporating directly Lidar sensor data to predict the binary grid, using recurrent neural network. To incorporate ego-vehicle motion, they utilize a spatial transformer to allow internal memory of RNNs to learn environment of the state. Mohajerin et al. [14] suggested an RNN-based architecture with a difference learning method, and makes OGM prediction in the field of view of ego-vehicle front camera. Schreiber et al. [15] proposed an encoder-decoder network architecture, along with skip connections, to make long-term DOGM predictions. While they collect the sensor data from an autonomous vehicle, the vehicle remains stationary and only acts as the sensor collection point at different intersections. Itkina et al. proposed to use evidential occupancy grid and implement PredNet architecture for the prediction [16]. The approach is then carried forward to develop the double-pronged architecture [17] and attention-augmented ConvLSTM [18]. The latter work is able to make long-term predictions, however at turns the predictions still lose the scene structure. Mann et al. [19] addressed the problem of OGM prediction in urban scenes by incorporating vehicles

semantics in the environment. Their proposed method depends on the annotated vehicle data labels available in the dataset.

Contrary to the conventional Occupancy Grid Prediction, we present an allo-centric DOGM representation to predict the urban traffic scene with respect to a fixed reference frame. Apart from the conventional recurrent representation learning approaches, we also use memory-augmented learning-based video-prediction method, in relevance to learning long-term motion context of the dynamic agents.

III. SYSTEM OVERVIEW

We discuss here the overall proposed approach for allo-centric DOGM prediction, the pipeline is summarized in Fig. 1.

A. Dynamic Occupancy Grid Maps

Dynamic occupancy grid maps provide a discretized representation of environment in a bird's eye view, where every cell in the grid is independent and carries information about the associated occupancy and velocity.

To generate DOGMs, we incorporate the Conditional Monte Carlo Dense Occupancy Tracker (CMCDOT) [6]. This approach associates four occupancy states to the grid. Each cell carries the probabilities of the cell being i) occupied and static, ii) occupied and dynamic, iii) unoccupied or free and iv) if the occupancy is unknown. The probabilities of these four states sum to one. In our work, we make use of three of these states and represent the grid as an RGB image. The channels Red, Green and Blue represent the unknown state, dynamic state and static state respectively. The associated probabilities of the cell in the 3-channel DOGM grid are interpreted as the pixel values of the RGB images. The RGB grid images can be seen in Fig. 1-2. Low probabilities in all three channels leave the grid-image black, therefore, representing free space.

For allo-centric grid generation, we define the grid in the world frame, close to the initial position of ego-vehicle. The state probabilities are initially computed in an ego-centric grid, since we use the on-board sensor data. To ensure that we have cell information for the complete allo-centric grid dimensions when the vehicle is dynamic and moving away from the world frame origin, a much larger ego-centric

DOGM is computed. This information is then fused to update every cell states in the allo-centric grid in the world frame.

We compare the allo-centric and ego-centric grids at 4 time instants for the same scene and same grid dimensions in Figure 2. In the allo-centric grid, the ego-vehicle (illustrated in the pink box) can be seen moving with respect to the grid, while it remains fixed in the ego-centric grid. It is important to note that the observable space around the ego-vehicle remains the same for both grids. However, since they are defined in different frames, the two cover different spaces in the scene at a given time. We illustrate the common space covered by both grids since the start of the sequence, marked by yellow boundary.




Fig. 2: Visualization of allo-centric and ego-centric grids, generated for the same scene. The area marked by yellow lines is the common region covered by both grids up until the t -th sequence. The ego-vehicle is boxed in pink grid and the bus passing by is encircled in white.

B. Problem Formulation

We formally define the task of predicting the scene in allo-centric DOGM representation, as sequence-to-sequence learning, see Fig. 1. A sequence comprises of a set of sequential grid images that capture the evolution of a given scene. Let $X_t \in \mathbb{R}^{3 \times W \times H}$ and $Y_t \in \mathbb{R}^{3 \times W \times H}$ be the t -th frame of the 3-channel grid-image where W and H denote the width and height respectively. The input sequence for the grid-image is denoted by $X_{t-N:t}$, representing N consecutive frames. Given a set of input sequence, the task of the network is to predict future grid images, i.e. output sequence. The target and predicted output sequences are denoted by $Y_{t+1:t+P}$ and $\hat{Y}_{t+1:t+P}$ where P is the prediction horizon.

For training and testing data, the DOGMs can be generated for both the input and the target sequences, leaving behind no additional need for labelled data or human intervention. Since the input sequences, $X_{t-N:t}$, and output sequences, $Y_{t+1:t+P}$, are represented as images, this prediction task can be considered a video prediction problem.

C. Deep Learning Prediction Architectures

To study and compare the scene prediction with ego-centric and allo-centric grids, we train our datasets with

different video prediction networks. We consider 3 networks, briefly discussed in section II-A: PredNet, PredRNN, LMC-Memory with memory alignment learning (here on referred as LMC-Memory).

PredNet [8], inspired from predictive coding, makes predictions based on how the predicted frames deviate from the target [20]. The original work tests the network on vehicle mounted camera images from Kitti dataset [21] and demonstrates the ability to capture both egocentric motion as well as motion of objects in camera images. We consider PredRNN [9] and LMC-Memory architecture [11] as the state of the art video prediction networks that aim to capture long-term dependencies and motion context. PredRNN implements novel ST-LSTM units with a zigzag internal memory flow and proposes memory decoupling loss to discourage learning redundant features. LMC-Memory architecture, on the other hand, proposes an external memory block with its own parameters to store various motion contexts. The approach also offers an efficient computation method since the motion context for long-term multi-step predictions is computed only once for a given input sequence.

We study these networks capabilities to retain the occupancy of the static region, and the ability to predict motion of dynamic agents in DOGM.

D. Unknown Channel and Loss functions

In both ego-centric and allo-centric grids, a significant part of the scene remains unobserved, see Fig. 2 (unknown channel is represented in red). This is more pronounced in the initial frames of the allo-centric grid, where the Lidar is unable to detect the farthest area from the ego-vehicle.

While it is more relevant to learn the evolution of static and dynamic components in the scene, inclusion of unknown channel is useful for our prediction task. A Lidar based grid is often unable to capture the full shape of a vehicle. For example, we can see in Fig. 2 how the occupied cells by the bus vary in different time steps on the grid. It is only in the 2.0s time step that a rectangular shape is observed, otherwise different parts of the bus remain unknown. The unknown channel at different instants also carries spatial information of the agents with respect to the ego-vehicle. Thus, with the sequential frames and the unknown channel, we assist the network to be able to extract spatial information and learn scene representation.

The inclusion of unknown channel and emphasis on learning static and dynamic components is addressed in the loss function. Loss function L in the implemented video prediction networks is modified to carry the weighted sum of the RGB channels:

$$L = \alpha L_R + \beta(L_G + L_B) \quad (1)$$

where,

L_R , L_G and L_B represent the loss for unknown (red), dynamic (green) and static channels (blue) respectively. In order to encourage the network to learn and improve the prediction of the static and dynamic channels, α is always kept smaller than β .

IV. EXPERIMENTS

A. Dataset

We study the prediction performance on the real-world NuScenes dataset [3]. The original dataset consists of 850 scenes for training and 150 scenes for testing, each scene is approximately 20s long. We generate the DOGM grid-based on the Lidar pointcloud and available odometry. For allo-centric grid, we represent the scene with respect to a fixed reference frame and a grid dimension of 60 x 60m, with a resolution of 0.1m per cell. Each sequence starts with the ego-vehicle heading facing up, capturing the scene 10m behind and 50m ahead of it. The initial pose was selected to ensure that the ego-vehicle remains within the grid for the total sequence length, even when running at a high speed. For egocentric grid, we generate a grid of the same dimensions and resolution, and the ego-vehicle fixed in the center. Each sequence is comprised of 35 frames, a time duration of 3.5s with DOGM grid images generated every 0.1s. In total, we have 4,250 training and 750 testing sequences respectively.

B. Training

The input sequence $X_{t-9:t}$ consists of 10 frames (1.0s). Each network is trained to make predictions $\hat{Y}_{t+1:t+25}$ for 25 future frames (2.5s). Both the allo-centric and ego-centric datasets are trained with the original parameters of the respective video prediction network. For training with PredRNN and LMC Memory networks, both allo-centric and ego-centric grid images are resized to 192x192 pixels. **PredRNN** is trained with a batch size of 4 and a learning rate of 10^{-4} . The number of channels of each hidden state is set to 64. The loss function is the sum of L2 and decoupling loss, and the values of α and β in Eq. (1) are set to 0.2 and 0.8. **LMC-Memory** is trained with a learning rate of 2×10^{-4} , memory slot is set to 100 and ConvLSTM to 4 layers for frame predictions. The loss function is the sum of L1 and L2 losses. The values of α and β are set to 0.2 and 0.8. For training with **PredNet**, the grid images are resized to 160x160 pixels. The network is set to 4 hierarchical layers with an initial learning rate of 10^{-3} . The loss function is the L1 loss of only the first layer, the values of α and β are set to 0.05 and 0.8. All models are trained on Adam optimizer for 30 epochs.

V. EVALUATION

For evaluation, we are particularly interested in static and dynamic agents in the scene. We discussed in section III-D, the utility of unknown regions in learning scene representation. But the unknown region occupies a big portion of the grid and, thus, in evaluation, overshadows the performance of more interesting and relevant segments: static and dynamic regions. For this reason, we evaluate the dataset and network performances based on two channels of the predicted images, the blue and green channels representing static and dynamic components in the scene. We encourage the readers to refer to the video¹ for a better visualization of the results.

¹<https://youtu.be/z-0BVM93X8c>

A. Quantitative Evaluation

The allo-centric and ego-centric grids at any instant observe different parts of the scene, see Fig. 2. For fair comparison between them, we modify the test dataset and crop out the part of each t -frame that has not been observed until the t -th sequence by both grids. Thus, for example, the part of the grids outside of the yellow dotted boxes in Fig. 2 are blacked out for the input sequence frames $X_{t-N:t}$ as well as the target frames in the output sequence $Y_{t+1:t+P}$.

We measure the performances using three metrics: MSE (Mean Square Error), SSIM (Structured Similarity Indexing Method), and LPIPS (Learned Perceptual Image Patch Similarity) [22]. MSE is calculated by the pixel-wise difference between the ground truth and the predicted frame per channel and per cell. However, with MSE, the slightest error in predicted motion can result in large errors in the ego-centric grids dataset. The SSIM and LPIPS metrics evaluate the prediction results based on the structural similarity and perception similarity respectively. Lower values are better for MSE and LPIPS while higher values are better for SSIM.

Table I shows average results for the complete 2.5s prediction horizons. The MSE score of allo-centric grids is significantly lower compared to the one of ego-centric grids. Since the complete scene transforms with respect to the ego-vehicle, the MSE is always higher in the ego-centric grid. The SSIM and LPIPS scores are also significantly superior for the allo-centric grid, due to the tendency of ego-centric grids to get increasingly blurry for higher prediction horizons.

Network	MSE x 10^{-2} (↓)	SSIM(↑)	LPIPS(↓)
Allo-centric grid			
LMC-Memory	0.894	0.895	0.167
PredRNN	0.882	0.904	0.167
PredNet	0.905	0.888	0.172
Ego-centric grid			
LMC-Memory	1.302	0.856	0.217
PredRNN	1.138	0.845	0.234
PredNet	1.335	0.847	0.225

TABLE I: Average results with allo-centric and ego-centric grids for prediction horizon of 2.5s. The allocentric grid outperforms the other in all three video prediction networks.

In Fig. 3, we plot scores of the metrics for every 0.5s prediction step. The results with allo-centric grid (shown in blue) always perform better than the ego-centric grids. Among the three prediction networks, overall PredRNN performs the best with allo-centric grids. However, with the ego-centric grids (results shown in orange), PredRNN offers a good MSE score but the SSIM and LPIPS performances drop after 1.0s. This is because PredRNN tends to make blurry and diffused predictions in the output frames; this helps reduce the MSE but the scene loses its structures. This is further seen in the qualitative results discussed in section V-B and illustrated in Fig. 4.

B. Qualitative Evaluation

The prediction results between the allo-centric and ego-centric grids differ drastically when the ego-vehicle is turning




Fig. 3: Results with the MSE(\downarrow), SSIM(\uparrow) and LPIPS(\downarrow) metrics with allo-centric and ego-centric grids for input sequences of 1.0s and prediction horizon up to 2.5s. For fair comparison, all test sequence frames were modified to only contain the scene observable in both the allo-centric and ego-centric grids. The allo-centric grid (results plotted in blue) outperforms the other with all three video prediction networks.




Fig. 4: Qualitative results for the ego-vehicle leaving a roundabout on both allo-centric (4a) and ego-centric grids (4b). The input sequence consists of 10 frames (1.0s) and output predicted sequence of up to 25 frames (2.5s). The prediction results are shown at 0.5s, 1.5s and 2.5s instants and are magnified at the interesting spaces, marked by red box in the target(ground truth) frames. The best results can be observed with LMC-Memory network with the allo-centric grid that retains the scene structure and predicts the motion of the ego-vehicle best.

at an intersection or driving on a curved road. Figure 4 shows results for a sequence where the ego-vehicle is exiting a roundabout. In this scene, while there are no other dynamic agents, the network needs to predict the behaviour of the ego-vehicle, when it is driving along the curved static segment (alluding to road structure) and is headed towards static objects/obstacles. For the allo-centric grids, the challenge is to predict the ego-vehicle pose while the scene remains static. The best results are achieved with the LMC-Memory. The vehicle pose is well-predicted up to 2.5s, its orientation is

adjusted so that it does not hit the static components. For the same grid, the PredRNN fails to learn and predict the behaviour resulting in false prediction of collisions. The ego-vehicle, while getting more blurry, diffuses into the static obstacles on the road. With the PredNet, the ego vehicle is almost already lost at 1.5s prediction horizon. This is expected behaviour since PredNet is ideally not aimed at long-term video predictions. With all three networks, the ego-vehicle gets more blurry, however with PredRNN, the static scene also tends to get blurry at larger prediction horizon.

In the ego-centric grid, the whole scene rotates around the ego-vehicle. LMC-memory and PredNet significantly lose the static components ahead of the vehicle. The rotation results in increasing blurriness at every time step. PredRNN predictions are more diffused and faint blurry cells are still visible ahead of the vehicle, even at 2.5s prediction horizon. In context of planning and safe navigation, this high uncertainty in the environment structure renders the prediction results unreliable.

VI. DISCUSSION AND FUTURE WORK

In this work, we presented a novel allo-centric dynamic occupancy grid approach for long-term prediction of urban traffic scene, and compared it to the conventional ego-centric DOGM approach. We trained and tested various video prediction networks to show that allo-centric DOGM representation has superior ability to predict the same scene.

The most significant improvement is the allo-centric grid's ability to retain the static scene structure, especially when the vehicle turns. The ego-centric grid, on the other hand tends to lose the static scene, and hence the crucial information about whether the given space is occupied or free.

The results of allo-centric grids prediction with state-of-the-art PredRNN and LMC-Memory approaches have shown complementary benefits. PredRNN predictions, though diffuse and get more blurry, are capable of maintaining agents longer. We observe that LMC-memory shows better tendency at learning behaviours in comparison to the PredRNN.

It is pertinent to mention here that the two grids are still very similar. In both scenarios, the observable space updates relative to the position of the vehicle in the scene. Thus, in allo-centric grid while the grid is no more fixed to the ego-vehicle, the ego-vehicle bias remains.

All three video prediction networks tested in this work address the prediction problem as deterministic. However, the behaviour of agents in urban traffic scene tends to be multimodal. For future work, the addition of multimodal prediction capabilities in the network architecture would be interesting. Additionally, the incorporation of semantics in the occupancy grid such as agent type and offline road information could assist in learning behaviours and interactions.

REFERENCES

- [1] S. Mozaffari, O. Y. Al-Jarrah, M. Dianati, P. Jennings, and A. Mouzakitis, "Deep learning-based vehicle behavior prediction for autonomous driving applications: A review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 1, pp. 33–47, 2020.
- [2] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social lstm: Human trajectory prediction in crowded spaces," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 961–971.
- [3] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," *arXiv preprint arXiv:1903.11027*, 2019.
- [4] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. R. Qi, Y. Zhou *et al.*, "Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9710–9719.

- [5] A. Nègre, L. Rummelhard, and C. Laugier, "Hybrid sampling bayesian occupancy filter," in *2014 IEEE Intelligent Vehicles Symposium Proceedings*. IEEE, 2014, pp. 1307–1312.
- [6] L. Rummelhard, A. Nègre, and C. Laugier, "Conditional monte carlo dense occupancy tracker," in *2015 IEEE 18th International Conference on Intelligent Transportation Systems*. IEEE, 2015, pp. 2485–2490.
- [7] S. Oprea, P. Martinez-Gonzalez, A. Garcia-Garcia, J. A. Castro-Vargas, S. Orts-Escolano, J. Garcia-Rodriguez, and A. Argyros, "A review on deep learning techniques for video prediction," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. [Online]. Available: <https://arxiv.org/pdf/2004.05214.pdf>
- [8] W. Lotter, G. Kreiman, and D. Cox, "Deep predictive coding networks for video prediction and unsupervised learning," *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, pp. 1–18, 2017.
- [9] Y. Wang, H. Wu, J. Zhang, Z. Gao, J. Wang, P. Yu, and M. Long, "Predrnn: A recurrent neural network for spatiotemporal predictive learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2022.
- [10] Y. Wang, M. Long, J. Wang, Z. Gao, and P. S. Yu, "Predrnn: Recurrent neural networks for predictive learning using spatiotemporal lstms," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017.
- [11] S. Lee, H. G. Kim, D. H. Choi, H.-I. Kim, and Y. M. Ro, "Video prediction recalling long-term motion context via memory alignment learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 3054–3063.
- [12] H.-S. Jeon, D.-S. Kum, and W.-Y. Jeong, "Traffic Scene Prediction via Deep Learning: Introduction of Multi-Channel Occupancy Grid Map as a Scene Representation," in *2018 IEEE Intelligent Vehicles Symposium (IV)*, 2018, pp. 1496–1501.
- [13] J. Dequaire, P. Ondruška, D. Rao, D. Wang, and I. Posner, "Deep tracking in the wild: End-to-end tracking using recurrent neural networks," *The International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 492–512, jun 2017. [Online]. Available: <https://doi.org/10.1177/0278364917710543>
- [14] N. Mohajerin and M. Rohani, "Multi-step prediction of occupancy grid maps with recurrent neural networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10 600–10 608.
- [15] M. Schreiber, S. Hoermann, and K. Dietmayer, "Long-term occupancy grid prediction using recurrent neural networks," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 9299–9305.
- [16] M. Itkina, K. Driggs-Campbell, and M. J. Kochenderfer, "Dynamic environment prediction in urban scenes using recurrent representation learning," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 2052–2059. [Online]. Available: <https://arxiv.org/abs/1904.12374>
- [17] M. Toyungyernsub, M. Itkina, R. Senanayake, and M. J. Kochenderfer, "Double-Prong ConvLSTM for Spatiotemporal Occupancy Prediction in Dynamic Environments," *arXiv preprint arXiv:2011.09045*, 2020. [Online]. Available: <http://arxiv.org/abs/2011.09045>
- [18] B. Lange, M. Itkina, and M. J. Kochenderfer, "Attention Augmented ConvLSTM for Environment Prediction," *arXiv preprint arXiv:2010.09662*, 2020.
- [19] K. S. Mann, A. Tomy, A. Paigwar, A. Renzaglia, and C. Laugier, "Predicting future occupancy grids in dynamic environment with spatio-temporal learning," 2022. [Online]. Available: <https://arxiv.org/abs/2205.03212>
- [20] R. P. Rane, E. Szügyi, V. Saxena, A. Ofner, and S. Stober, "Prednet and predictive coding: A critical review," in *Proceedings of the 2020 international conference on multimedia retrieval*, 2020, pp. 233–241.
- [21] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *International Journal of Robotics Research (IJRR)*, 2013.
- [22] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *CVPR*, 2018.