



HAL
open science

Flow-guided Motion Prediction with Semantics and Dynamic Occupancy Grid Maps

Rabbia Asghar, Wenqian Liu, Lukas Rummelhard, Anne Spalanzani, Christian Laugier

► To cite this version:

Rabbia Asghar, Wenqian Liu, Lukas Rummelhard, Anne Spalanzani, Christian Laugier. Flow-guided Motion Prediction with Semantics and Dynamic Occupancy Grid Maps. ITSC 2024 - 27th IEEE International Conference on Intelligent Transportation Systems, Sep 2024, Edmonton, Canada. pp.1-6, <10.48550/arXiv.2407.15675>. <hal-04885618>

HAL Id: hal-04885618

<https://hal.science/hal-04885618v1>

Submitted on 14 Jan 2025

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 CC BY 4.0 - Attribution - International License

Flow-guided Motion Prediction with Semantics and Dynamic Occupancy Grid Maps

Rabbia Asghar, Wenqian Liu, Lukas Rummelhard, Anne Spalanzani, Christian Laugier
Univ. Grenoble Alpes, Inria, 38000 Grenoble, France, email: FirstName.LastName@inria.fr

Abstract—Accurate prediction of driving scenes is essential for road safety and autonomous driving. Occupancy Grid Maps (OGMs) are commonly employed for scene prediction due to their structured spatial representation, flexibility across sensor modalities and integration of uncertainty. Recent studies have successfully combined OGMs with deep learning methods to predict the evolution of scene and learn complex behaviours. These methods, however, do not consider prediction of flow or velocity vectors in the scene. In this work, we propose a novel multi-task framework that leverages dynamic OGMs and semantic information to predict both future vehicle semantic grids and the future flow of the scene. This incorporation of semantic flow not only offers intermediate scene features but also enables the generation of warped semantic grids. Evaluation on the real-world NuScenes dataset demonstrates improved prediction capabilities and enhanced ability of the model to retain dynamic vehicles within the scene.

Index Terms—Scene Prediction, Motion Forecasting, Deep Learning, Autonomous Vehicles

I. INTRODUCTION

Predicting driving scenes holds significant importance in enhancing road safety, optimizing traffic flow, and advancing autonomous driving technologies. By accurately anticipating various elements within a driving environment such as pedestrian movements, and vehicle trajectories, predictive models enable proactive decision-making for both human drivers and autonomous systems. This foresight aids in mitigating risks, preventing accidents, and ultimately saving lives.

Dynamic Occupancy Grid Maps (DOGMs) represents static and dynamic elements within the environment in a bird’s-eye-view (BEV) grid. One of the key advantages lies in their flexibility regarding sensor dependency, as they can be generated from various types of sensors including Lidar, radars, and cameras. In our research, we leverage probabilistic DOGMs [1], which further enhances predictive capabilities by incorporating uncertainty estimation. Moreover, by integrating deep learning techniques with such probabilistic DOGMs, we can effectively learn complex interactions within the driving environment, tackle the intrinsic spatiotemporal problems, and lead to more robust and precise predictions.

In the domain of autonomous vehicles, Occupancy Grid Map (OGM) predictions are often ego-centric ([2], [3]), meaning they are fixed relative to the ego-vehicle frame. Dependent on the motion of the ego-vehicle, the future predicted grids suffer from the disappearance of static and dynamic scenes over time and loss of scene structure, especially during turns. To address these shortcoming, Asghar *et al.* [4] proposed making allo-centric DOGM predictions by

representing the scene evolution in a fixed reference frame. This approach successfully preserves scene integrity independent of ego-vehicle motion and focuses on the behavior and motion prediction of dynamic agents in the scene.

Asghar *et al.* [5] extended this work to consider availability of semantic information with the input and to predict agents motion in the scene as sequential BEV semantic grids. The framework demonstrated improved abilities to predict dynamic agents as well as the possibility to evaluate with the ground truth.

In this work, we build further on the existing approaches to formulate a multi-task framework by learning the flow of the scene to improve future semantic grids. For our input we consider both occupancy probabilities and estimated velocities of the dynamic scene components from DOGMs, along with the semantic information from the camera images. Our future predictions are allo-centric, fixed with respect to the latest ego position. We generate predictions for both the future semantic grids and the flows. The predicted flows also serve as intermediate features, capturing to the spatial and temporal changes in the semantic information within the scene. By forecasting the flows, we anticipate how the semantic content of the scene will evolve over time. Subsequently, we iteratively warp the current semantic grid with the future flows to obtain the sequence of future semantic grids. A brief overview of this work can be accessed online[†].

The main contribution of this paper is a novel multi-task framework that combines DOGM and semantics to predict multi-step vehicle semantic grids, as well as the multi-step flow of the complete scene. Evaluation on real-world NuScenes dataset [6] shows how the proposed framework improves prediction capabilities and learns to recognize vehicle behaviors within the scene without prior maps.

II. RELATED WORK

A. Occupancy Grid Prediction

In the state-of-the-art literature, the challenge of predicting ego-centric future OGM (Occupancy Grid Map) involves incorporating spatio-temporal deep-learning methods ([7], [8], [9], [10],). Asghar *et al.* [4] proposed a novel approach for future DOGM prediction within a fixed reference frame, preserving scene integrity particularly during ego-vehicle maneuvers.

Studies by Toyungyernsub *et al.* [2] and Schreiber *et al.* [11] proposed separate predictions of the static and dynamic

[†]<https://youtu.be/rhyC2Xo0B-g>

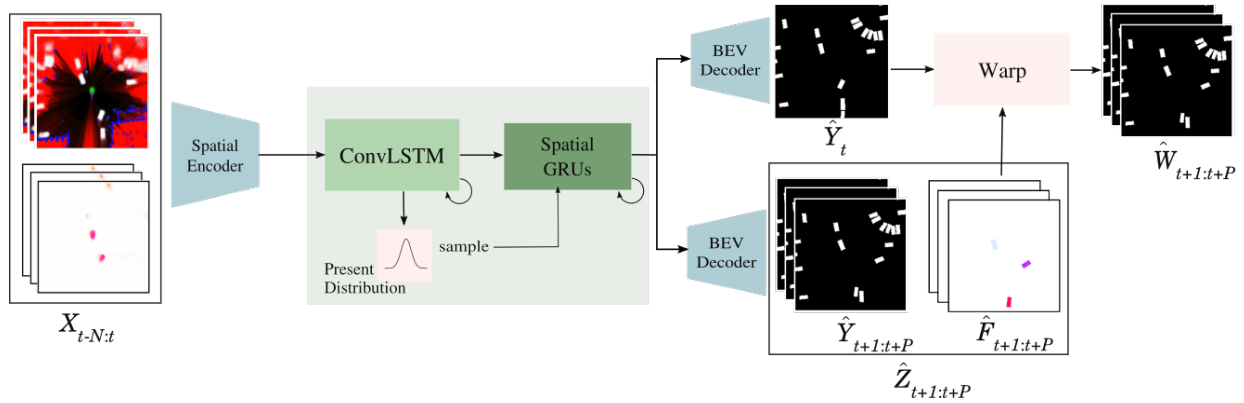


Fig. 1: An overview of our proposed network. Our network utilizes a sequence of DOGMs, occupancy state grids, associated velocity grids, and semantic grids $X_{t-N:t}$ as input to capture the scene evolution. Subsequently, it predicts vehicle semantic grid \hat{Y}_t at current time step and a sequence of future vehicle semantic grids $\hat{Y}_{t+1:t+P}$ and scene flows $\hat{F}_{t+1:t+P}$.

scene, yet they do not predict velocities or scene flows. Mann *et al.* [3] enhanced DOGMs with semantic vehicle labels, while [5] integrated DOGMs with additional map and vehicle semantics, thus improving prediction capabilities and accommodating diverse future scenarios through a probabilistic distributions module.

In this paper, contrary to the conventional OGM prediction methods, we forecast both the vehicle occupancy grids and the dynamics of the entire scene. We carry forward the work of Asghar *et al.* [5], incorporate velocity information, predict the scene flows in the output, generate warped semantic grids, and demonstrate how these enhancements can improve the prediction capabilities of vehicles in the scene.

B. Motion Forecasting via Occupancy Grids

In the realm of motion forecasting, the literature tends to predominantly focus on approaches utilizing detection, tracking, and trajectory prediction of agents. Nevertheless, recent efforts have emerged to explore the efficacy of occupancy grids for motion forecasting, albeit to a lesser extent. Notably, Rules of the Road [12] presented a comparative analysis between trajectory methods and occupancy grids, proposing a dynamic program for decoding likely trajectories from occupancy grids under a simple motion model. MotionNet [13] introduced a joint perception and motion prediction network that outputs BEV maps, with each grid cell encoding object category and motion information. In response to the demand for real-time prediction, Terwilliger *et al.* [14] proposed semantic forecasting by predicting current frame segmentation and future optical flow with RNN-based networks. MP3 [15] advances a concept akin to Flow Fields, termed Motion Fields, to predict forward motion vectors and associated probabilities per grid cell, facilitating planning tasks. Following this line, Fiery [16] employed instance segmentation to predict future occupancy grids and motion flows of dynamic agents. Waymo [17] presented a model employing backward flows to predict both occupancy and flow in a spatio-temporal grid. Additionally, Waymo launched a challenge [18] addressing the same research problem. This challenge, however, considers availability of tracked agents as well as prior road map information.

Different from these methods, we employ probabilistic version of DOGMs [1] as our input source, which provide static and dynamic occupancy states of the scene in a BEV grid. Our input DOGMs not only offer probabilistic information even when only partial data is available, but also facilitate seamless integration of various sensor inputs and configurations without retraining the entire network.

III. SYSTEM OVERVIEW

Our methodology builds upon the overall pipeline proposed in [5], which addresses vehicle motion forecasting into semantic grids from OGMs. This work specifically focuses on incorporating scene flow to further improve the vehicle motion predictions as semantic grids. We discuss in detail here our proposed approach, and the pipeline is summarized in Fig. 1.

A. Problem Formulation

We formally define the task of vehicle motion prediction, see Fig. 1. Let $X_t \in \mathbb{R}^{6 \times w \times h}$ be the t -th frame of the input grid where w and h denote the width and height respectively. X_t comprises of three occupancy state channels, two velocity channels, and one semantic grid channel. Let $Z_t \in \mathbb{R}^{3 \times w \times h}$ be the t -th frame of the output grid that contains a single vehicle semantic grid channel, $Y_t \in \mathbb{R}^{1 \times w \times h}$, and two flow grid channels, $F_t \in \mathbb{R}^{2 \times w \times h}$.

Given a set of input sequence $X_{t-N:t}$, the task of our proposed multi-head network is to predict 1) the vehicle grid \hat{Y}_t at the current time step, and 2) the sequence of the future grids $\hat{Z}_{t+1:t+P}$. This sequence encompasses both the future vehicle semantic grids $\hat{Y}_{t+1:t+P}$, and the future flow grids $\hat{F}_{t+1:t+P}$, with P representing the prediction horizon.

Next, we utilize the predicted sequence of flow grids $\hat{F}_{t+1:t+P}$ to warp the vehicle grid \hat{Y}_t recursively, and obtain an additional sequence of warped future vehicle semantic grids $\hat{W}_{t+1:t+P}$. This warping process can be expressed by Eq. 1,

$$\hat{W}_{t+1} = f_w(\hat{W}_t, \hat{F}_{t+1}) \quad (1)$$

where f_w is the warping operation and $\hat{W}_t = \hat{Y}_t$ initially.

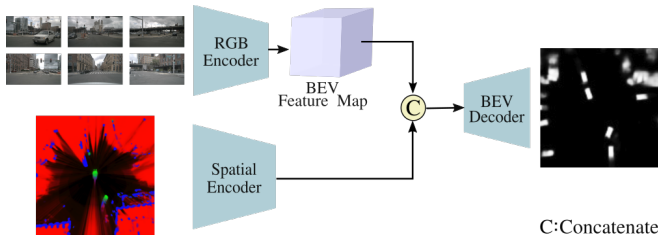


Fig. 2: To obtain semantic information, we fuse BEV features encoded from camera images with occupancy state grid. The Red, Green, and Blue channels in the RGB grid denote the unknown state, dynamic and static occupied states respectively. The black areas within the occupancy state grid indicate free space.

B. Input Scene Representation

1) *Dynamic occupancy grid maps*: DOGMs provide a grid-based representation of the environment in a bird’s-eye view. Each cell in the grid is estimated in parallel within the system and contains information about its occupancy state, as well as associated dynamics. Unlike classic object detection-and-tracking methods, such approaches mitigates risks linked to uncautious thresholding.

We incorporate Bayesian dynamic occupancy grid filter [1] to generate DOGMs from LiDAR sensor data. This approach is used to estimate probabilities associated with four possible occupancy states for each cell in the grid: i) free, ii) occupied and static, iii) occupied and dynamic, and iv) unknown occupancy. Additionally, the framework provides velocity estimates for each cell, represented by mixtures of grids and particle sets.

In the proposed network, the input $X_t \in \mathbb{R}^{6 \times w \times h}$ comprises six channels, with five channels sourced from the output of a DOGM. These encompass information pertaining to three DOGM occupancy states (the definitive visualization represented in Fig. 2) and two-channel DOGM velocities along the x and y -axis. Figure 1 showcases the DOGM state grids and the velocity grids depicted below as two-dimensional flow illustrations in the input. An additional layer of semantic information is overlaid on top of the DOGM state grids, depicted in white. The inclusion of semantic information is discussed in the section below.

2) *Semantic information*: A semantic grid contains the probability assigned to each cell for being occupied by a specific semantic category, such as vehicles in this case. The approach outlined in the Lift, Splat, Shoot (LSS) [19] is repurposed in this work to prepare the semantic information for input into our proposed network. LSS extracts feature maps from multiple input camera images in BEV and decodes them into vehicle occupancy prediction represented as a semantic grid. Our adaptation involves concatenating corresponding DOGMs alongside the feature maps extracted from camera images and decoding the combined information into semantic grids. This fusion of DOGMs with BEV features results in a more comprehensive understanding of the scene. This modified pipeline is illustrated in Fig 2.

Note that the semantic information can be incorporated in this framework via alternative methods as well. For example, vehicle semantic information from camera images

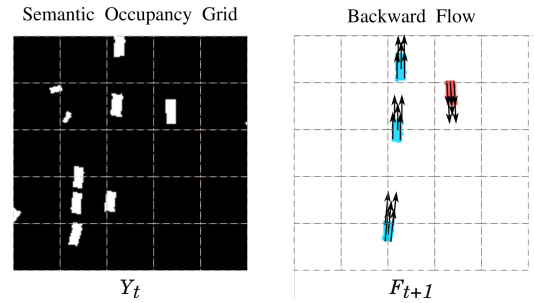


Fig. 3: In F_{t+1} , backward flow vectors indicate the motion of vehicles and point to the origin of their respective occupancy in Y_t . Red and blue represent dynamic vehicles moving upward and downward within the grid, respectively.

can be integrated into the Bayesian perception framework [1], allowing for probabilistic estimation alongside the DOGM.

C. Flow-guided Predictions

The network is equipped with two decoder heads: one is dedicated to predicting a semantic grid for current time step, while the other forecasts sequential future grids and flows simultaneously.

The flow is represented as a two-channel grid that models the velocity of vehicles along the x and y dimensions. We utilize backward flow to capture vehicle motion from their future positions back to their current positions. At each time step, the backward flow contains velocity vectors for each cell, indicating the change in occupancy from the previous time step. Refer to Fig. 3 for a visual illustration.

To generate a sequence of flow-guided semantic grid, we start with the current semantic grid \hat{Y}_t . This grid is initially warped using the flow \hat{F}_{t+1} to obtain a warped semantic grid \hat{W}_{t+1} . Subsequently, \hat{W}_{t+1} is further warped with the next flow \hat{F}_{t+2} . This process is repeated recursively for all the subsequent time steps to yield a sequence of warped vehicle semantic grids. Worth noting, our method not only predicts the flow of occupancy but also incorporates the potential flow in free spaces.

Alongside predicting the flow, the network also forecasts a sequence of future semantic grids. This simultaneous prediction aids the network in correlating the occupancy of the semantic grid with the corresponding flow vectors.

D. Model Architecture

The model combines spatio-temporal predictions with a conditional variational approach.

1) *Encoder*: Each frame in the input sequence, denoted as $X_{t-N:t}$, is passed through the spatial encoder. Then the generated features are sequentially fed into a ConvLSTM (Convolutional Long Short Term Memory) block, enabling the simultaneous capture of spatial and temporal information. To achieve this, we use 4 ConvLSTM units, each with 128 hidden features, ensuring the extraction of spatio-temporal features that encompass the past evolution of the scene.

2) *Probabilistic modeling*: To allow for multimodal predictions, we adopt a conditional variational approach [20]. This approach allows the network to learn two distributions that capture the evolution of the scene in latent space: i)

present distribution that is conditioned on present spatiotemporal features and ii) future distribution that is conditioned on present spatiotemporal as well as future ground truth. Both distributions are parameterized as diagonal Gaussian distributions, each consisting of 32 latent space dimensions.

3) *Decoder*: The spatial GRUs utilize spatio-temporal features to iteratively forecast future states, conditioned on samples from Gaussian distributions. Each timestep comprises of 3 units of convolutional GRUs. These multi-step future states are then fed to the two decoder heads, as discussed in III-C.

For more details on the model, we refer the reader to [5].

E. Losses

The network is trained to learn vehicle motion forecasting as semantic grids and their corresponding flow representations. All vehicle semantic grids are trained with binary cross-entropy loss \mathcal{L}_{BCE} with a positive sample weight of 5. This includes loss for \hat{Y}_t , $\hat{Y}_{t+1:t+P}$, and $\hat{W}_{t+1:t+P}$. The loss for warped semantic grids \hat{W} acts as self-supervised loss for the flow grid \hat{F} .

$$\mathcal{L}_{BCE_{sum}} = \lambda_d \mathcal{L}_{BCE_{\hat{Y}_t}} + \lambda_b \mathcal{L}_{BCE_{\hat{Y}_{t+1:t+P}}} + \lambda_w \mathcal{L}_{BCE_{\hat{W}_{t+1:t+P}}}$$

where λ_d , λ_b and λ_w are the loss weights for \hat{Y}_t , $\hat{Y}_{t+1:t+P}$, and $\hat{W}_{t+1:t+P}$ respectively.

For the supervised loss of the flow grid \hat{F} , we define an L1 loss, only for the cells that are occupied by the vehicles. Additionally, we introduce Kullback-Leibler divergence loss \mathcal{L}_{KL} to encourage alignment between the present and the future distributions, described in III-D.2.

The final loss is weighted sum of all these losses.

$$\mathcal{L} = \mathcal{L}_{BCE_{sum}} + \lambda_f \mathcal{L}_{flow} + \lambda_k \mathcal{L}_{KL}$$

where λ_f and λ_k are the respective loss weights.

IV. EXPERIMENTS

A. Dataset

We evaluate our proposed model on the real-world NuScenes dataset [6]. The original dataset consists of 700 and 150 scenes for training and validation respectively. Each scene has a duration of 20s, with lidar data available at 10Hz and annotated keyframes at frequency of 2Hz. The lidar based DOGMs are generated comprising of state grids and velocity grids, with a resolution of 0.1 m. For semantic grid prediction, all 6 camera images at each keyframes are considered along with respective DOGM state grids.

The ground truth annotations for semantic grids are prepared based on the vehicle bounding box annotations provided in the dataset. The velocity or speed of the vehicles is unavailable. Therefore, the ground truth flow of the vehicles is estimated by calculating the displacement of vehicle centroids on the grid, which is then normalized by the grid dimensions.

All grids are initially generated ego-centrally, with the ego-vehicle positioned at the center and facing upward. Future predictions in each sequence are made relative to the ego-vehicle frame of reference at the latest timestep. The past

and future grids are subsequently transformed to facilitate allo-centric predictions in a fixed frame, using the available odometry data from the dataset, and then cropped to cover an area of 60x60m.

In total, we have 23K training and 5K validation sequences respectively.

B. Training

The input sequence $X_{t-2:t}$ consists of 3 frames, spanning over 1.0s. The network is trained to make predictions for present \hat{Y}_t and 4 future frames $\hat{Y}_{t+1:t+4}$, for instances 0.5s to 2.0s, with steps of 0.5s. For training, the grid images are resized to 240x240 pixels, thus each pixel has a resolution of 0.25m. The loss coefficients are set to $\lambda_b = 1.0$, $\lambda_w = 1.0$, $\lambda_f = 0.05$ and $\lambda_k = 0.005$. Our network is implemented on Pytorch and trained on 8 Nvidia Tesla V100 GPUs. Adam Optimizer is employed at the learning rate of 3×10^{-4} and weight decay of 1×10^{-7} . Dataset is trained for 20 epochs with the batch size of 18. This takes approximately 5 hours.

V. RESULTS

A. Qualitative Evaluation

Figure 4 shows predicted flow and warped occupancies at 1.0s and 2.0s across three scenes. The two-dimensional flow vectors are represented as RGB images, with colors indicating various vector directions, while occupancies are depicted with distinct colors at each timestep. Multimodal predictions are generated using 100 random samples from a Gaussian distribution, with color intensity corresponding to the frequency of occupancy at each cell. Additionally, semantic channels are superimposed on the state grids in the input for visualization purposes. Notably, the network learns scene structure from DOGM occupancy information and the behavior of dynamic vehicles, as map or road information is unavailable in the input.

The network’s flow grid predictions focus on warping the semantic grid from one timestep to the next, rather than specifically capturing the flow of occupancy. This allows for predicting both the movement of dynamic vehicles and the resulting free space. In the flow images, white spaces dominate the static scene structure observed in the input state grids. Regions with no observed occupancies, whether in the input or the predicted horizon, contain random values in the flow grid as they do not impact the predicted grids.

Collision Avoidance: A notable attribute in the flow-warped semantic grids is their effective collision avoidance capability during predictions. The motion prediction of dynamic vehicles consistently demonstrates a tendency to navigate away from static occupancies within the scene, regardless of whether they are static vehicles or unrecognized categories. As illustrated in Fig. 4, multimodal predictions avoid the static obstacles within the scene, as highlighted in the magnified sections of grids in the rightmost column. For instance, in the first example, bus predictions exhibit cautious behavior in the face of static cells ahead. Similarly, in the third scenario, the multimodal predictions steer clear of small static obstacles within the space.

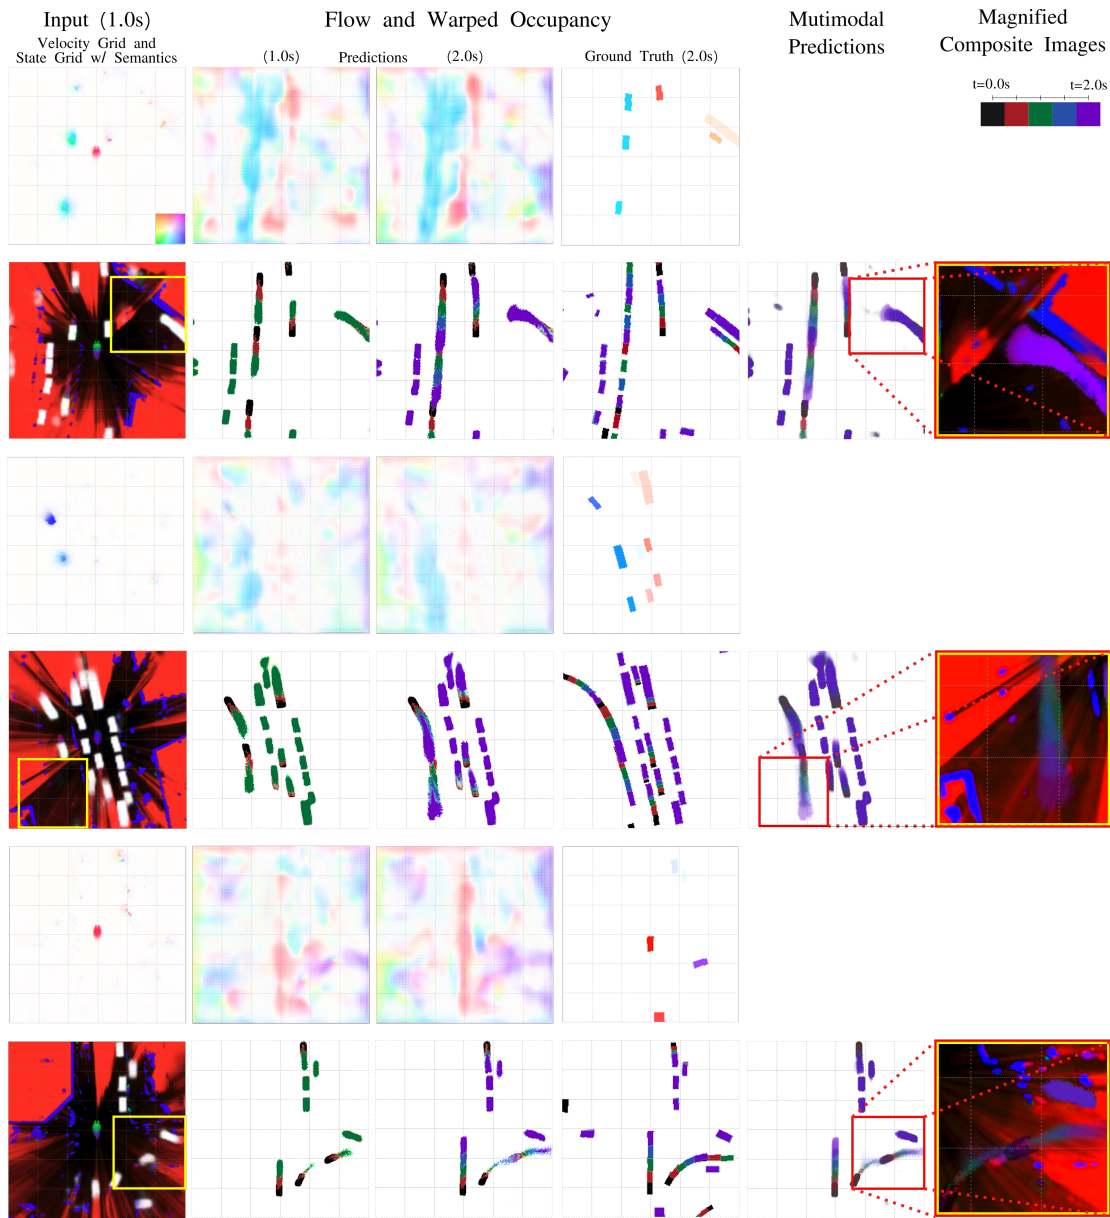


Fig. 4: Semantic flow and warped occupancy prediction examples are demonstrated on three scenes from the Nuscenes Dataset [6], covering an area of 60x60m. The first cloumn displays the DOGM input with semantic labels (in white) at the latest timestep. Flow and warped occupancies are shown for 1.0s and 2.0s. Occupancies are color-coded to visualize the range from 0.0s to 2.0s, transitioning from black to purple. The last column showcases magnified current static part of the scene, on which are displayed agent predictions. The magnified scene correspond to the yellow and red boxes, providing a zoomed-in view of the multimodal predictions overlaid on occupancies in the DOGM.

Furthermore, the predictions in dynamic spaces also show notable improvement. The semantic grid predictions without flow tend to get blurry when multiple dynamic vehicles are in close proximity due to uncertainty, for examples refer to [5]. However, with the inclusion of flow, predictions become sharper, resulting in clearer delineation even in dense traffic. This enhancement is particularly evident in the second example, especially among the top vehicles in the scene.

B. Quantitative Evaluation

1) *Ablation Study*: We conduct an ablation study with different configurations of the proposed network to investigate

the role of semantic flow. We consider 3 different variations for the future prediction decoder: i) only predict semantic grids $\hat{Y}_{t+1:t+P}$, ii) only predict flows $\hat{F}_{t+1:t+P}$, and iii) predict both together. These variations are trained with and without the velocity grid in the input.

We evaluate the performance of the semantic and warped grids using binary classification evaluation metrics: Intersection of Union (IoU) and Area under the precision-recall curve (AUC). IoU is computed with a threshold of 0.5, while AUC is calculated across 100 linearly spaced thresholds. Table I shows the average IoU and AUC values over a 2.0s prediction horizon for six cases. The IoUs for our

proposed architecture are similar to those of the method solely predicting semantic grids, see the first and the last row. However, incorporating flow predictions significantly enhances AUC, indicating improved precision and recall - the proportion of correct vehicle predictions and correctly predicted occupancies, respectively.

The integration of DOGM velocity grid and prediction of semantics and flow collectively result in optimal IoU and AUC, enhancing the network capacity for behavior learning. Interestingly, employing the same configuration without velocity input yields inferior scores, suggesting the crucial contribution of velocity inclusion in avoiding local minima during training.

TABLE I: Ablation study on input and output configuration of the architecture. We report mean values of IoU and AUC over a 2.0s prediction horizon. Bold indicates best.

| Input Velocity | Output Semantics/Flow | IoU(\uparrow) | | AUC(\uparrow) | |
|-------------------|--------------------------|-------------------|--------------|-------------------|--------------|
| | | Sem | Warped | Sem | Warped |
| | Semantics only | 0.493 | - | 0.613 | - |
| ✓ | Semantics only | 0.487 | - | 0.626 | - |
| | Flow only | - | 0.481 | - | 0.631 |
| ✓ | Flow only | - | 0.480 | - | 0.630 |
| | Semantics + Flow | 0.485 | 0.463 | 0.629 | 0.613 |
| ✓ | Semantics + Flow | 0.503 | 0.493 | 0.649 | 0.631 |

2) *Retention of vehicles*: We further evaluate the network ability to predict static and dynamic agents by tracking individual agents. In our observations, dynamic agent predictions tend to lose clarity with each time step and may eventually become indistinct. Similarly, predictions of vehicles in semantic grids, though improved, also gets blurrier or even disappears when the network is uncertain of vehicle behavior.

To assess network’s ability to retain vehicles in the scene, we track vehicles perceived in the current semantic grid \hat{Y}_t individually. If the prediction of the vehicle does not overlap with the ground truth semantic grid, we consider it *lost*. We define overlap as at least 10 cells with an occupancy probability higher than 0.3, covering an area of 0.625 m² (approximately a third of a vehicle). We study three cases: a) using four-channel semantic-OGM predictions with a video prediction network PredRNN [21], b) our proposed network without flow, and c) the complete proposed network.

TABLE II: Retention of dynamic and static vehicles lost in a 2.0s prediction horizon. Bold indicates best.

| Network | Vehicles Retained (%) (\uparrow) | |
|---------------|--------------------------------------|-------------|
| | Dynamic | Static |
| PredRNN [21] | 90.2 | 97.0 |
| Ours w/o flow | 95.8 | 99.4 |
| Ours | 97.6 | 99.7 |

We show the percentage of static and dynamic vehicles retained by the three methods in Table II, considering a total of 10K dynamic and 20K static vehicles in the validation sequences. Under these metrics, our method tends to retain more agents, static or dynamic, than mere classic video prediction approach, with the inclusion of flow information further improving the results.

VI. CONCLUSION

In this work, we introduce a novel forecasting approach that leverages DOGMs and semantic information to predict scene evolutions as flow-guided semantic grids. Our framework is capable of multi-modal predictions, learning patterns and behaviours based on the observable static scene and motion of dynamic agents. By predicting flow alongside semantic grids, this method tends to improve prediction results, especially in dynamic parts. While only vehicles were considered in the flow-guided agent prediction, future work should extend this approach to include other types of agents, and unidentified grid-level components.

REFERENCES

- [1] L. Rummelhard, J.-A. David, A. G. Moreno, and C. Laugier, “A cross-prediction, hidden-state-augmented approach for dynamic occupancy grid filtering,” in *ICARCV*, 2022.
- [2] M. Toyungyernsub, E. Yel, J. Li, and M. J. Kochenderfer, “Dynamics-aware spatiotemporal occupancy prediction in urban environments,” in *IROS*, 2022.
- [3] K. S. Mann, A. Tomy, A. Paigwar, A. Renzaglia, and C. Laugier, “Predicting future occupancy grids in dynamic environment with spatio-temporal learning,” in *IV*, 2022.
- [4] R. Asghar, L. Rummelhard, A. Spalanzani, and C. Laugier, “Allo-centric occupancy grid prediction for urban traffic scene using video prediction networks,” in *ICARCV*, 2022.
- [5] R. Asghar, M. Diaz-Zapata, L. Rummelhard, A. Spalanzani, and C. Laugier, “Vehicle motion forecasting using prior information and semantic-assisted occupancy grid maps,” in *IROS*, 2023.
- [6] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liang, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nuscenes: A multimodal dataset for autonomous driving,” in *CVPR*, 2020.
- [7] 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 *IV*, 2018.
- [8] J. Dequaire, P. Ondruška, D. Rao, D. Wang, and I. Posner, “Deep tracking in the wild: End-to-end tracking using recurrent neural networks,” in *IJRR*, 2017.
- [9] N. Mohajerin and M. Rohani, “Multi-step prediction of occupancy grid maps with recurrent neural networks,” in *CVPR*, 2019.
- [10] B. Lange, M. Itkina, and M. J. Kochenderfer, “Attention augmented convlstm for environment prediction,” in *IROS*, 2021.
- [11] M. Schreiber, S. Hoermann, and K. Dietmayer, “Long-term occupancy grid prediction using recurrent neural networks,” in *ICRA*, 2019.
- [12] J. Hong, B. Sapp, and J. Philbin, “Rules of the road: Predicting driving behavior with a convolutional model of semantic interactions,” in *CVPR*, 2019.
- [13] P. Wu, S. Chen, and D. N. Metaxas, “Motionnet: Joint perception and motion prediction for autonomous driving based on bird’s eye view maps,” in *CVPR*, 2020.
- [14] A. Terwilliger, G. Brazil, and X. Liu, “Recurrent flow-guided semantic forecasting,” in *WACV*, 2019.
- [15] S. Casas, A. Sadat, and R. Urtasun, “Mp3: A unified model to map, perceive, predict and plan,” in *CVPR*, 2021.
- [16] A. Hu, Z. Murez, N. Mohan, S. Dudas, J. Hawke, V. Badrinarayanan, R. Cipolla, and A. Kendall, “FIERY: Future instance segmentation in bird’s-eye view from surround monocular cameras,” in *ICCV*, 2021.
- [17] R. Mahjourian, J. Kim, Y. Chai, M. Tan, B. Sapp, and D. Anguelov, “Occupancy flow fields for motion forecasting in autonomous driving,” in *IEEE RAL*, 2022.
- [18] “Occupancy and flow prediction challenge,” <https://waymo.com/open/challenges/terms/?continue=%2Fopen%2Fchallenges%2F2024%2Foccupancy-flow-prediction%2F>, accessed: 2024-04-15.
- [19] J. Philion and S. Fidler, “Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d,” in *ECCV*, 2020.
- [20] A. Hu, F. Cotter, N. Mohan, C. Gurau, and A. Kendall, “Probabilistic future prediction for video scene understanding,” in *ECCV*, 2020.
- [21] Y. Wang, H. Wu, J. Zhang, Z. Gao, J. Wang, S. Y. Philip, and M. Long, “Predrnn: A recurrent neural network for spatiotemporal predictive learning,” in *TPAMI*, 2022.