

Supplementary Materials

Implicit Occupancy Flow Fields for Perception and Prediction in Self-Driving

Ben Agro*, Quinlan Sykora*, Sergio Casas*, Raquel Urtasun
Waabi, University of Toronto
{bagro, qsykora, sergio, urtasun}@waabi.ai

In this appendix, we first describe implementation details that are relevant to IMPLICITO, including the full computation graph of our implicit decoder architecture, the encoder architecture, and training details. Then, we provide implementation details for the baselines. Last but not least, we showcase several additional results:

- Introspection of IMPLICITO ($K = 4$) using visualizations of the attention offsets, cross attention weights, and qualitative results.
- Analysis of occupancy metrics over time for both object-based and object-free prior art as well as IMPLICITO.
- Analysis of how performance is affected by grid resolution for explicit object-free methods.
- Insights into the occupancy probability calibration of different methods via reliability diagrams.

1. Training Details

Initialization: We found that training was more consistent if we initialized the weights with a small standard deviation of 0.01 and a bias of 0 for the linear layer that predicts the offsets $\Delta\mathbf{q}$. This ensures that the initial offset predictions are small enough such that $\mathbf{r} = \mathbf{q} + \Delta\mathbf{q}$ still represents a point within the feature map \mathbf{Z} , but large enough such that the predicted offsets do not get stuck at $\Delta\mathbf{q} = \mathbf{0}$.

Loss hyperparameters: We use a flow regression loss weighting of $\lambda_f = 0.1$ (downweighted by a factor of 10 relative to occupancy loss).

Optimizer: We use an initial learning rate of 1.0×10^{-3} , which is multiplied by a factor of 0.25 every 6 training epochs. We train for a total of 20 epochs, using the AdamW optimizer [5] with a weight decay of 1.0×10^{-4} .

Query points used during training: We use $|\mathcal{Q}| = 220,000$ query points per example in the batch, sampled uniformly on the spatio-temporal volume. Future work may investigate more effective sampling methods for training.

2. Baseline Implementation Details

Object detector: For a fair comparison with object-based and hybrid baselines (MULTIPATH, LANEGCN, GORELA and OCCFLOW), we use the same state-of-the-art object detector, Two-stage PIXOR, in all of their implementations. This detector is a variant of the original PIXOR [8], with the following improvements:

- We use multi-scale deformable self-attention [9] instead of the upsampling deconvolutional layers after the ResNet backbone to aggregate information from different scales and enhance the feature extraction. Following PIXOR, we output the feature maps with $1/4$ resolution.

*Denotes equal contribution

- We use the output from the dense detector header as the first-stage results. The top 500 bounding boxes after non-maximum suppression (NMS) are used as region proposal for the second stage. The 2D IoU threshold for NMS is set to 0.7.
- We use RotatedROIAlign [7] to extract 3×3 ROI features from the feature map in the second-stage header, and apply two self-attention layers: one on features within each ROI and another between features of different ROIs.
- Two MLPs are used to predict the classification score and box refinement for the region proposals.
- We use DETR-like set-based loss [1] with bipartite matching for both stages. An additional IoU loss is used for bounding box regression besides the smooth L1 loss.

Trajectory prediction baselines: For the methods that perform trajectory prediction (MULTIPATH, LANE GCN, GoRELA), we train Two-stage PIXOR jointly with the trajectory prediction model in an end-to-end fashion, as proposed in SpAGNN [2].

OccFlow: This hybrid method performs object-based perception (object detection) and object-free prediction (occupancy-flow). Since it requires kinematic information in the raster passed from perception to prediction, we train the detector to additionally output longitudinal, lateral and angular velocity and acceleration. We then postprocess this information into a raster with five channels: occupancy probability based on detector confidence, velocity in x and y directions, and acceleration in x and y directions, for each pixel. Given this raster, we employ the same encoder architecture as our model for a fair comparison. Finally, a fully convolutional header decodes both occupancy and backward flow on a spatio-temporal grid.

MP3: This baseline also uses the same encoder as our model for a fair comparison. It then decodes initial occupancy and forward flow over time with a fully convolutional header, and the future occupancy is obtained by warping the initial occupancy with the temporal flow as described in the original paper.

3. Additional Results

3.1. Qualitative Results for IMPLICITO with $K = 4$ offsets

Fig. 1 presents the occupancy predictions over time, reverse flow predictions, attention offsets, and cross attention weights of IMPLICITO with $K = 4$ across four different scenes in AV2. The reverse flow predictions, attention offsets, and cross attention weights are all visualized at the final timestep in the prediction horizon. For clarification on the different flow visualizations used throughout main manuscript and supplementary, please see Fig. 3. The cross-attention weights are those used in the weighted sum of the K interpolated feature vectors in the feature aggregation module.

First, we notice that the offsets learn to look mainly backward and forward in the direction of the lanes. Second, we find that the decoder pays more attention to those offsets that look backwards along the lanes. This is particularly apparent in Offset 1; it is weighted heavily in the lane regions where it points against the flow of traffic. This aligns with our intuition that the offset sampling mechanism helps the model attend to the where occupancy was in the past — where the LiDAR evidence is — in order to predict the future.

We hypothesize that IMPLICITO with $K = 4$ does not outperform IMPLICITO with $K = 1$, because, assuming the offsets look backwards in time and space, $K = 1$ can already model multi-modal futures for occupancy, in the same way that the backwards-flow formulation does. Then, multiple offset predictions would only provide more information in areas where occupancy could have come from multiple directions in the past, like the intersection in the Scene 3 of Fig. 1. Indeed, focusing in on this example in Fig. 2, we see that IMPLICITO with $K = 4$ is qualitatively able to predict future occupancy from both the left turning car (A) and the cars entering the intersection (B), as highlighted in the blue box, while IMPLICITO with $K = 1$ only predicts occupancy from B. However, these examples are rare in the training and evaluation data, and $K > 1$ adds noise in cases where it is not needed, explaining the slightly worse metrics of $K = 4$ relative to $K = 1$.

3.2. Measuring Occupancy Predictions Over Time

This section presents an analysis of the occupancy prediction accuracy as a function of prediction time Δt . We want to investigate (1) how do the occupancy metrics of the object-free baselines from the literature and IMPLICITO evolve over time, and (2) how well do the explicit models (those that predict occupancy on a spatio-temporal grid) predict occupancy at timesteps that are not aligned with the grid.




Figure 1. Qualitative results of IMPLICITO with $K = 4$ on AV2. The grey scale images display the cross attention weights, which sum to 1 across the columns. The HSV color images are the offset predictions. The predicted reverse flow field is masked by the ground truth occupancy for visual comparison (see Fig. 3). The reverse flow, attention offsets, and cross attention weights are all visualized at the final timestep.

Expanding on (2): Due to memory constraints, all explicit models (e.g., MP3 and OCCFLOW) output an occupancy map with temporal dimension $(5/0.5) + 1 = 11$, where the $+1$ comes from the prediction at the initial observation time ($\Delta t = 0$). However, typically motion planners may need higher temporal resolution query points, because, for example, a highway vehicle can travel 15 m in 0.5 s. IMPLICITO by-passes this problem because it is trained on continuous query points at no additional cost. Thus, in this section we evaluate all models at a time discretization of 0.1 s (the period of LiDAR the sweeps). For the explicit models, this requires linear interpolation on the predicted occupancy map.




Figure 2. A side by side comparison of IMPLICITO with $K = 1$ and $K = 4$. The blue box highlights that $K = 4$ is qualitatively better able to predict occupancy in the intersection that could have come from two past locations, namely **A** and **B**.

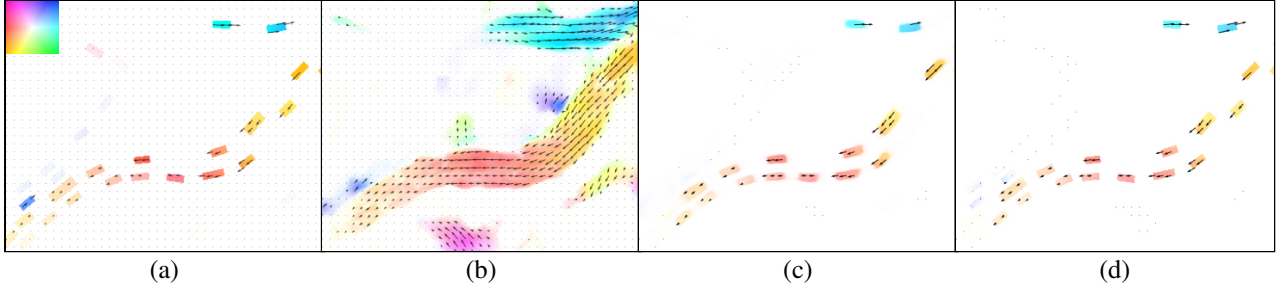


Figure 3. Visualizations of the reverse flow predictions of IMPLICITO with $K = 4$ at the initial timestep. (a) is the ground truth flow field, (b) is the predicted flow field, (c) is the predicted flow field masked by the occupancy predictions, and (d) is the predicted flow field masked by the ground truth occupancy.




Figure 4. mAP and Soft IoU as functions of prediction time for IMPLICITO and various object-free models from the literature on HwySim.

The results for HwySim are presented in Fig. 4. The first trend to note is that while all models have a similar mAP and Soft-IoU at the initial timestep, IMPLICITO has a higher mAP and Soft-IoU at later timesteps. We attribute this to its expressive attention offset mechanism, which increases the effective receptive field, allowing accurate occupancy predictions far away from the LiDAR evidence information in \mathbf{Z} . Secondly, we observe that at prediction times that require interpolation (e.g., those not at $\{0, 0.5, \dots, 5\}s$), the explicit models have worse mAP and Soft-IoU than IMPLICITO. This is because this interpolation inherently fails to model how occupancy moves; e.g., interpolation at a spatio-temporal point between grid timesteps can result in no predicted occupancy even if occupancy had passed through that point along its trajectory. IMPLICITO does not suffer from this problem because it can produce predictions at any continuous time. We observe this in




Figure 5. mAP and Soft IoU as functions of prediction time for various object-free models from the literature and IMPLICITO on AV2.

the smooth mAP and Soft-IoU curves over time of IMPLICITO. We notice that while MP3 has a more accurate occupancy ranking (mAP) over time than OCCFLOW, it has a worse Soft-IoU. The better occupancy ranking is attributed to the flow-warping mechanism of MP3, which imposes a realistic prior on the evolution of occupancy, and increases the effective receptive field. However, as described by Mahjourian et al. [6], MP3 suffers from lost occupancy over time, as reflected in its poor Soft-IoU.

We also include occupancy metrics over time on AV2, presented in Fig. 5. While the conclusions are similar here, they are less obvious, likely due to the fact that the vehicles move a smaller proportion of the ROI than in HwySim on average, making interpolation more accurate. To illustrate the advantages of IMPLICITO we compare it qualitatively against baselines on a single scene from AV2 in Fig. 6. GORELA predicts many agents entering the intersection at once, which is unrealistic. OCCFLOW has a detection with incorrect size and orientation of a vehicle at $\Delta t = 0.0$ s, and suffers from spreading / inaccurate occupancy predictions at later timesteps, likely due to the limited receptive field of its convolutional decoder. MP3 has disjoint-pixel occupancy predictions, which is caused by its forward-flow warping mechanism and was first observed by Mahjourian et al. [6]. In contrast, IMPLICITO exhibits accurate occupancy predictions and realistic multi-modality (vehicle going straight or turning left, and a lane-change).

3.3. Spatial Resolution Evaluation

In this experiment, we investigate how the performance of an explicit object-free model is affected by its output grid resolution. We train three MP3 models, each with a different output grid resolution; 0.2 m, 1.0 m, and 2.0 m. The rasterized input to the encoder is kept the same, and the models are evaluated with a spatial resolution of 0.2 m by bi-linearly interpolating the output grid. The results are presented in Tab. 1, along with the metrics of IMPLICITO for comparison. We notice that at an output resolution of 2.0 m, the performance of MP3 is much worse than with a 0.2 m output resolution. However, 1.0 m is only slightly worse. This is likely because the ground-truth labels are rectangular bounding boxes, so as long as the discretization resolution is less than the smallest bounding box dimension, bi-linear interpolation on the output grid will be accurate. Most vehicles have a smallest side length greater than 1.0 m, but less than 2.0 m, explaining the results.

	mAP \uparrow	Soft-IoU \uparrow	ECE \downarrow	EPE \downarrow	Flow Grounded	
					mAP \uparrow	Soft-IoU \uparrow
MP3 (2.0 m)	0.635	0.286	0.731	0.547	0.986	0.280
MP3 (1.0 m)	0.769	0.416	0.307	0.458	0.927	0.460
MP3 (0.2 m)	0.774	0.422	0.201	0.472	0.902	0.466
IMPLICITO	0.799	0.480	0.193	0.267	0.936	0.597

Table 1. Comparing MP3 [3] with different spatial output resolutions, evaluated at a spatial resolution of 0.2 m on AV2. The results for IMPLICITO are presented for reference.

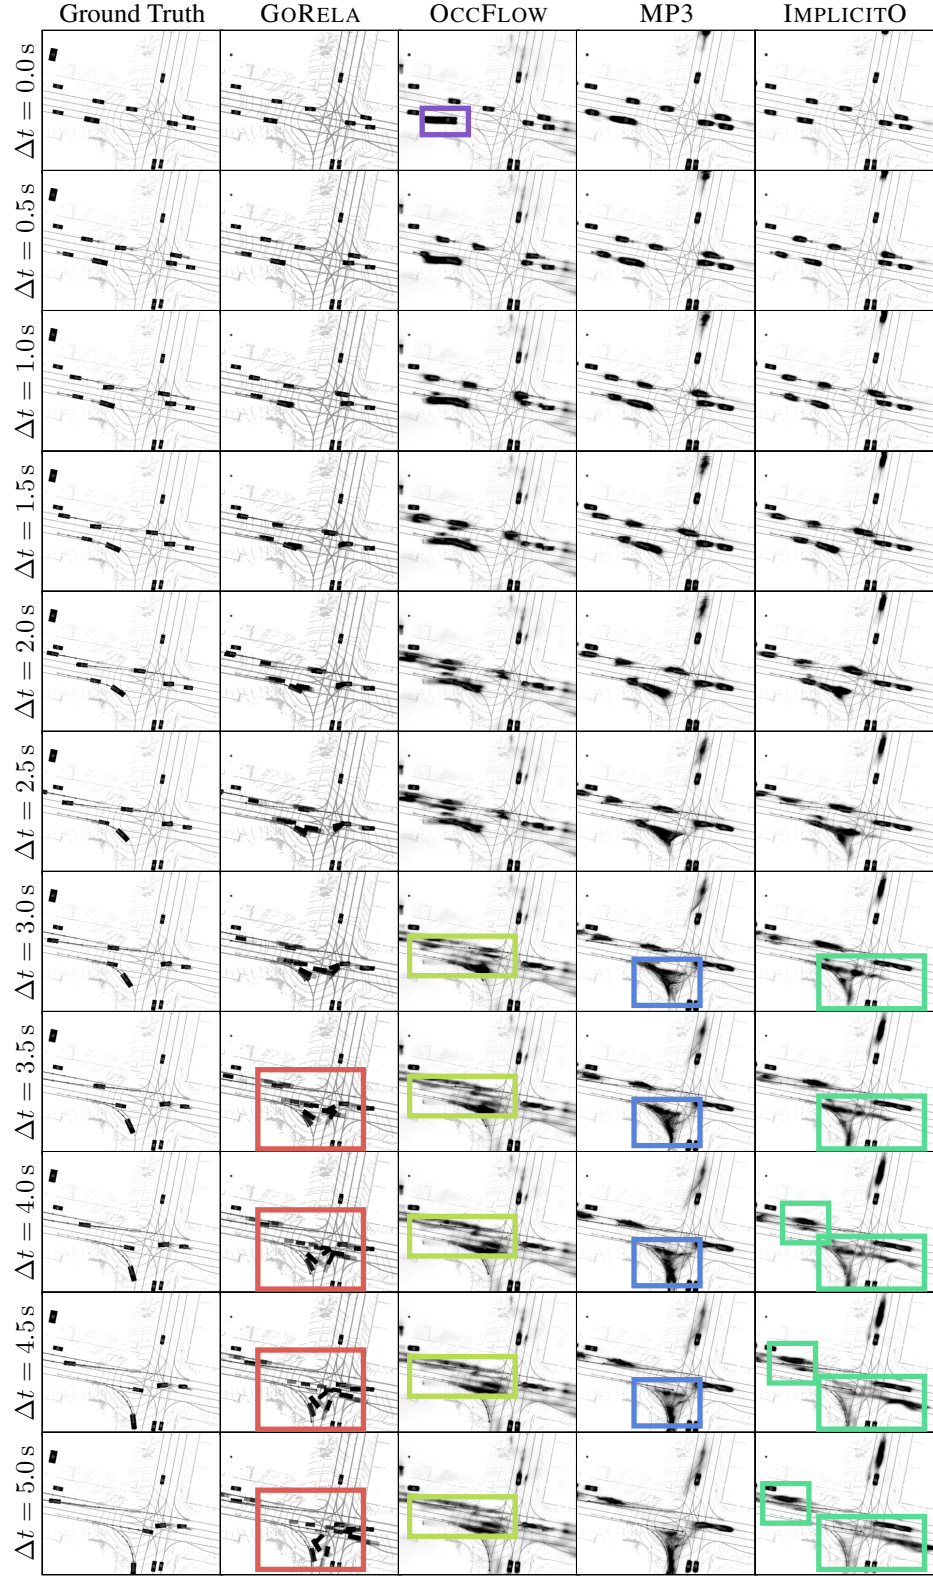


Figure 6. Occupancy predictions over time on a single scene from AV2 for various baseline models and IMPLICITO. The alpha channel denotes occupancy probability, colored in black. Observations are denoted with colored boxes: **inconsistent with actors**, **spreading occupancy predictions**, **incorrect detection size**, **disjoint-pixel occupancy predictions**, **realistic multi-modal predictions**,




Figure 7. Reliability diagrams on AV2.

3.4. Analyzing Occupancy Calibration with Reliability Diagrams

Fig. 7 displays the reliability plots for all baselines from the literature on AV2. All the methods in the top row are object-based, while those in the bottom row are object-free. We notice that the object-based methods are poorly calibrated on the task of occupancy prediction relative to the object-free methods. Object-based methods predict discrete multi-modal trajectories which poorly capture the uncertainty in occupancy, and result in overconfidence. On the contrary, object-free methods can capture non-parametric occupancy distributions without thresholding to produce instances, resulting in better calibration.

3.5. Query Point Training Procedure

As described in Sec. 3.3, we train IMPLICITO with a batch of continuous query points \mathcal{Q} sampled uniformly from the future spatio-temporal volume. This section ablates this querying procedure by training IMPLICITO on AV2 by querying points on a regular spatio-temporal grid with 0.5 s temporal resolution and varying spatial resolutions. Our evaluation procedure is the same as described in Sec. 4, and our results are presented in Tab. 2. We observe a slight disadvantage of training with points from a regular grid with respect to our proposed method.

	Spatial Resolution			
	0.2 m	0.5 m	1.25 m	2.0 m
mAP \uparrow	-0.19%	-0.25%	-0.83%	-0.90%
Soft-IoU \uparrow	-0.46%	-1.86%	-1.66%	-2.24%

Table 2. Comparing the performance of IMPLICITO when trained with query points sampled from a regular spatio-temporal grid against the proposed continuous query point training procedure. We report the relative change in evaluation metrics.

3.6. LiDAR History Ablation

In this section we measure the effect of LiDAR input frequency and history duration on the performance of IMPLICITO. First, we highlight that the LiDAR frequency and history duration as well as encoder architecture were held constant for all

models so changes in performance due to different LiDAR inputs are expected to be similar across all methods. When we train IMPLICITO with 10 past LiDAR frames at 10 Hz and see a minor relative improvement of 0.24% in mAP and 0.98% in Soft-IoU. For reference, MP3 improves similarly (0.45% mAP and 0.22% Soft-IoU).

3.7. Inferring Agent Identity

In this section, we provide a proof of concept for inferring agent identity (ID) with appropriate post-processing, similar to [4,6]. Inferring agent ID requires some form of instance prediction (like the “centerness” field from [4], or object detections in OCCFLOW [6]).

Without any additional training required, we pair IMPLICITO $K = 1$ with an object detector (the same one used by the object-based baselines). Then, for any query point \mathbf{q} , we add the attention offset $\Delta\mathbf{q}$ to find the reference point $\mathbf{r} = \mathbf{q} + \Delta\mathbf{q}$. Finally, we assign the ID of occupancy at \mathbf{q} to that of detection with the closest centroid to \mathbf{r} . Recall that despite the attention offset $\Delta\mathbf{q}$ being unsupervised, we found \mathbf{r} to point to the occupancy at the current time $\Delta t = 0$ (i.e., where the LiDAR evidence for that object is; see L721-737, Fig. 4 in the main paper).

Figure 8 visualizes the predicted occupancy colored by inferred instance ID (i.e., its associated detection). This simple approach illustrates that agent ID can be preserved in our implicit occupancy-flow method, and opens the door to future work on this area to improve this heuristic.




Figure 8. Qualitative examples of IMPLICITO inferring agent ID.

References

- [1] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 2
- [2] Sergio Casas, Cole Gulino, Renjie Liao, and Raquel Urtasun. Spatially-aware graph neural networks for relational behavior forecasting from sensor data. *arXiv preprint*, 2019. 2
- [3] Sergio Casas, Abbas Sadat, and Raquel Urtasun. Mp3: A unified model to map, perceive, predict and plan. In *CVPR*, 2021. 5
- [4] Anthony Hu, Zak Murez, Nikhil Mohan, Sofia Dudas, Jeffrey Hawke, Vijay Badrinarayanan, Roberto Cipolla, and Alex Kendall. Fiery: Future instance prediction in bird’s-eye view from surround monocular cameras. In *ICCV*, 2021. 8
- [5] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 1
- [6] Reza Mahjourian, Jinkyu Kim, Yuning Chai, Mingxing Tan, Ben Sapp, and Dragomir Anguelov. Occupancy flow fields for motion forecasting in autonomous driving. *RA-L*, 2022. 5, 8
- [7] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 2
- [8] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3d object detection from point clouds. In *CVPR*, 2018. 1
- [9] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv*, 2020. 1