

QuAD: Query-based Interpretable Neural Motion Planning for Autonomous Driving

Sourav Biswas^{*,1,2}, Sergio Casas^{*,1,2}, Quinlan Sykora^{1,2}, Ben Agro^{1,2}, Abbas Sadat¹, Raquel Urtasun^{1,2}

¹Waabi, ²University of Toronto

{sbiswas, sergio, qsykora, bagro, asadat, urtasun}@waabi.ai

Abstract—A self-driving vehicle must understand its environment to determine the appropriate action. Traditional autonomy systems rely on object detection to find the agents in the scene. However, object detection assumes a discrete set of objects and loses information about uncertainty, so any errors compound when predicting the future behavior of those agents. Alternatively, dense occupancy grid maps have been utilized to understand free-space. However, predicting a grid for the entire scene is wasteful since only certain spatio-temporal regions are reachable and relevant to the self-driving vehicle. We present a unified, interpretable, and efficient autonomy framework that moves away from cascading modules that first perceive, then predict, and finally plan. Instead, we shift the paradigm to have the planner query occupancy at relevant spatio-temporal points, restricting the computation to those regions of interest. Exploiting this representation, we evaluate a candidate trajectory around key factors such as collision avoidance, comfort, and progress for safety and interpretability. Our approach achieves better highway driving quality than the state-of-the-art on high-fidelity closed-loop simulations.

I. INTRODUCTION

Self-driving vehicles (SDVs) strive to reach their destinations safely and comfortably by analyzing their surroundings, envisioning potential future scenarios, and using this information to determine a plan of action to carry out. This is repeated with every new observation.

The majority of autonomy frameworks are *object-based* [9], [11], [29], [44], which implies detecting a discrete set of objects, typically obtained by thresholding output confidence scores from an object detector, predicting a small set of hypothetical future trajectories, and finally planning a safe trajectory. However, this approach loses information about the scene from thresholding and has limited representation of future object uncertainty. Furthermore, the expressivity of the future trajectory forecasts is limited, as keeping the number of hypotheses low is crucial for real-time inference.

Alternatively, *sensor-to-plan* imitation learning frameworks avoid reasoning about individual objects by learning to map sensor data directly to plans [7], [14], [27], [43]. These learned policies are typically brittle to distributional shift since supervision is only provided in states visited by the expert during training, so the policy never learns to recover from its own mistakes [32]. Moreover, the decisions are not easy to interpret or explain, which is important for system validation and verification.




Fig. 1: The set of potential plans for the ego vehicle, with time into the future colored from blue to green. Our neural motion planner, QUAD, builds upon two observations: (1) the plans’ reachable space is much smaller than the full spatio-temporal volume and (2) many ego states throughout the trajectories are in close proximity to each other.

To tackle these limitations, *occupancy-based* approaches [5], [13], [23], [33] have proposed a more interpretable object-free paradigm. Occupancy describes the probability a point in space and time is occupied by any traffic participant. This enables planners to decide on consistent and effective actions with respect to this interpretable representation, serving as an explanation of their decision. This paradigm is also more robust to shifts in the ego state distribution due to the supervision of the intermediate representations, making it better suited for closed-loop deployment. However, representing the occupancy in dense spatio-temporal grids over a large region of interest is resource-intensive since a high resolution is needed to attain accurate plans.

We propose QUAD, an interpretable, effective and efficient neural motion planner. QUAD diverges from prior works that first perceive, then predict, and finally plan. Instead, our unified autonomy first generates candidate trajectories respecting kinematic constraints and traffic rules, and then queries an implicit occupancy model only at spatio-temporal points needed for planning, which is used to rank the safety of the candidates. Fig. 1 shows an example of the candidate trajectories, which we can see only occupy a small portion of the spatio-temporal volume prior works predict. Moreover, we note that many candidate trajectories heavily overlap, which motivates us to quantize the spatio-temporal query points to reduce redundant computation.

Through extensive evaluation we show that QUAD is able to achieve better closed loop performance in a state-of-the-

*Denotes equal contribution

art highway driving simulator while attaining better runtime than competitive baselines.

II. RELATED WORK

This section reviews prior approaches to end-to-end autonomy from sensor data.

Object-based autonomy: Most previous approaches have employed object-based reasoning using a three-stage pipeline: (1) object detection [19], [42] and tracking [35], [39], (2) trajectory prediction based on past tracks [6], [26], [36], [46], and (3) motion planning to decide which action the self-driving vehicle should take [10], [30], [31], [34]. This paradigm faces three key challenges [33], [37]: (i) the uncertainty from detection and tracking is not propagated to downstream prediction, (ii) the predicted future distributions tend to be overly simple in practice for computational tractability in scenes with many actors, and (iii) the planner is blind to any objects that fall below the confidence threshold used to determine whether an object exists or not in detection. Several works [3], [4], [16], [20], [22] have addressed (i) by optimizing jointly through multiple stages. The recent work of PlanT [29] tackles (ii) by learning a Transformer that plans a trajectory for the ego vehicle from object and route tokens coming from the detector and the map, skipping the second stage of trajectory prediction. However, (iii) is fundamentally difficult to address in this autonomy paradigm as the core assumption is that there is a finite set of objects with a clear boundary between positive and negative instances.

Sensor-to-plan autonomy: Autonomy approaches in this family learn to map sensor data directly to plans, without any other intermediate interpretable representations. The pioneering work of ALVINN [28] proposed the use of a single neural network that directly outputs driving control commands. Following the triumph of deep learning, direct control-based methods have made strides through the implementation of more expressive networks, advanced sensors, and scalable learning methods [2], [7], [12], [14], [25], [45]. To leverage prior knowledge about vehicle kinematics, recent methods such as NMP [43] and Lift-Splat-Shoot [27] leverage analytical trajectory samplers to propose candidate plans, reducing the learning problem to predicting the cost of the trajectories. Despite not requiring any manual annotation for training, using only expert demonstrations as supervision tends to result in policies with stability and robustness issues, as they become highly vulnerable to distributional shift [5], [8], [32]. Note that some methods in this family [14], [43] perform auxiliary tasks for additional supervision, but the auxiliary outputs are not used during inference by the planner and therefore no consistency between those outputs and the plans is guaranteed. Thus, the planner’s decisions cannot be explained by those auxiliary outputs.

Occupancy-based autonomy: These methods predict spatio-temporal occupancy from sensor data, and plan a trajectory that is safe with respect to the predicted occupancy *without considering individual actors*. The occupancy is then

used to assess the risk of a trajectory by measuring the probability of collision [5], [13], [33]. This approach has proven very effective and robust to shifts in state distribution during deployment. Different architectures have been proposed to predict occupancy. P3 [33], MP3 [5], FIERY[15] and OcFlow [23] all propose to predict similar variations of 3D spatio-temporal occupancy grids using convolutional neural networks (CNN). However, this relies on the CNN receptive field being large enough to “transport” the occupancy from where the evidence is in the sensor data, to where the object would be at the end of the forecasting horizon. Moreover, predicting a dense 3D grid is computationally demanding and unnecessary. To tackle these two shortcomings, ImplicitO [1] proposes an implicit occupancy model that can be queried at continuous spatio-temporal points by leveraging deformable attention [47], yet it does not propose how to make use of this model for driving. In this work, we propose a unified autonomy framework which utilizes a state-of-the-art implicit occupancy architecture to obtain a strong understanding of the scene along with a query point quantization strategy to attain superior plans in practical runtimes without degrading driving quality.

III. QUERY-BASED AUTONOMOUS DRIVING (QUAD)

The goal of our motion planner (QUAD) is to find the best trajectory plan to execute given a short history of LiDAR observations, L , and a high-definition (HD) map, M . As shown in Fig. 2, with every new sensor observation, QUAD plans a trajectory by optimizing the following objective over a finite set of trajectory candidates \mathcal{T} generated based on the current ego kinematic state and the map:

$$\tau^* = \arg \min_{\tau \in \mathcal{T}} J(\tau, \mathbf{Z}, M), \quad (1)$$

where $\mathbf{Z} \in \mathbb{R}^{H \times W \times C}$ is a bird’s eye view (BEV) latent representation of the environment extracted from a voxelized multi-sweep LiDAR and a raster map by a learned scene encoder.

We define our objective function J as a linear combination of N explainable costs f_i , weighted by learnable coefficients w_i

$$J(\tau, \mathbf{Z}, M) = f(\tau, \psi(\mathcal{Q}_\tau, \mathbf{Z}), M) \quad (2)$$

$$= \sum_i^N w_i f_i(\tau, \psi(\mathcal{Q}_\tau^{(i)}, \mathbf{Z}), M). \quad (3)$$

The learned occupancy model ψ provides flexibility, expressivity and interpretability in our planner. It models occupancy *implicitly* and can be queried for occupancy at any set of continuous spatio-temporal locations \mathcal{Q} . More precisely, $\psi(\mathbf{q}, \mathbf{Z})$ outputs the probability that the continuous spatio-temporal location $\mathbf{q} = (x, y, t)$ lies within an object surface in BEV. Given this implicit occupancy representation, our planner can understand complex environments by having the different costs f_i reason about occupancy at different spatio-temporal locations $\mathcal{Q}_\tau^{(i)}$.

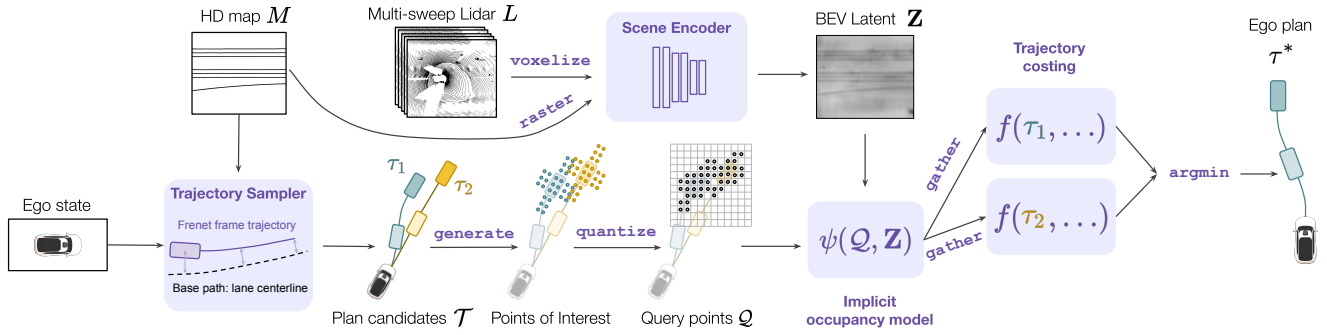


Fig. 2: **QUAD inference.** Given the ego state and the map, the trajectory sampler generates candidate plans. These plans are converted into query points that cover the relevant areas around the ego vehicle future positions. Leveraging multi-sweep LiDAR and HD map, a scene encoder builds a BEV latent representation which we then use to query an implicit occupancy model. Finally, we gather the occupancy relevant to each trajectory, cost them, and select the one with the lowest cost.

A. Trajectory Sampling and Query Points

Trajectory Sampler: We consider a set of candidate trajectories $\tau \in \mathcal{T}$ starting from the current ego state and going 5 seconds into the future. The ego state consists of the SDV current location in BEV, speed, and steering. A trajectory τ contains a sequence of kinematic bicycle model [18] states for each time step in the planning horizon. It is crucial that the set of sampled trajectories, while small enough to enable real-time computation, encompasses a range of maneuvers, including lane following, lane changes, nudges to avoid encroaching objects, hard brakes. To accomplish this efficiently, we opt for a sampling approach that takes into account the lane structure. More precisely, we use the lane centerlines from the HD map as base paths and sample longitudinal and lateral profiles in Frenet frame [34], [40]. As a result, the sampled trajectories align with appropriate lane-based driving, while incorporating lateral variations.

Points of Interest: The goal of these points of interest is to cover the relevant areas around the ego vehicle throughout the candidate plans. For every trajectory τ and time step t (sampled every 0.5s), we consider points within the ego bounding box as well as points forwards, backwards, and to the sides of the ego. For simplicity, we sample a uniform grid of points within the ego box at a certain resolution, and simply shift this grid forward/backward by the length of the ego vehicle and right/left by the width to obtain all the points of interest. This process is shown in Fig. 2 for the final time step of two trajectories.

Point Quantization: Since the points of interest are sampled along and around the trajectories \mathcal{T} , which are generated to ensure coverage of the available actions, the distance between multiple pairs of query points $\|\mathbf{q}_j - \mathbf{q}_k\|_2$ from different trajectories can be very small, as depicted in Fig. 1. We leverage this observation to improve the efficiency of our planner by quantizing the query points with a certain spatial resolution and only querying ψ with the unique set of points after quantization. We tune the quantization resolution to maximize efficiency without sacrificing driving performance.

Empirically, we find this to reduce the number of queries by two orders of magnitude, from millions to tens of thousands.

B. Implicit Occupancy Model

Our implicit occupancy model consists of a scene encoder that provides a BEV latent representation of the environment \mathbf{Z} , and an implicit occupancy decoder ψ that attends to the latent scene representation to predict occupancy probability at query points. By maintaining an intermediate occupancy representation, we can make planning decisions which are interpretable and consistent with the occupancy predictions.

Scene Encoder: We use as input a sequence of LiDAR point clouds, containing the 5 latest LiDAR sweeps. Each sweep contains a set of points with coordinates (p_x, p_y, p_h) , where the (p_x, p_y) is the point location in the SDV coordinate frame while p_h is the height over the ground. We then voxelize [41] the LiDAR in BEV to obtain a 3D tensor where the different sweeps are concatenated along the height dimension. Since the behavior of other traffic participants is highly influenced by the road topology, we make use of the prior knowledge stored in the HD map to provide important cues about the regions they might occupy and how they could move. More precisely, we raster the polylines representing the lane centerlines in the HD map as a BEV binary map with the same spatial resolution as the LiDAR. Following [1], [4], our scene encoder uses two convolutional stems for processing the voxelized LiDAR and map raster respectively. The resulting feature maps are concatenated along the channel dimension and passed through a lightweight Feature Pyramid Network [21] to get a fused BEV feature map \mathbf{Z} containing information from both modalities at half resolution of the inputs. Intuitively, the latent scene embeddings \mathbf{Z} contain local geometry, motion and semantic descriptors from the area within the receptive field of our encoder.

Implicit Occupancy Decoder: Leveraging the latent scene embedding \mathbf{Z} , our implicit occupancy decoder predicts the occupancy probabilities at a set of query points $\mathcal{Q} = \{\mathbf{q}_j\}_{j \in [1, |\mathcal{Q}|]}$. In more detail, each query point $\mathbf{q} =$

$(x, y, t) \in \mathbb{R}^3$ denotes a spatio-temporal point in BEV at a future time t . We exploit ImplicitO [1], a recently proposed state-of-the-art architecture for occupancy prediction. Given a query point, it bilinearly interpolates a latent vector at the query point BEV location (x, y) , and uses it to predict locations to attend using deformable attention [47]. With the attended latent vector, an MLP decoder predicts occupancy for a particular query point. This simple architecture has a great advantage over prior methods that predict occupancy grid maps using CNNs: it can attend anywhere in the BEV latent instead of being limited by the CNN’s receptive field to a small region. This is crucial since vehicles can travel very fast, so to accurately predict the occupancy into the future (e.g., at $t = 5s$), the model needs to find the original LiDAR evidence at $t = 0s$, which may be 150-200 meters behind. The specific query points \mathcal{Q} used during inference depend on the trajectory sampler and costs explained next.

C. Trajectory Costing

In order to plan an effective trajectory, we must consider various factors of driving such as collision likelihood, traffic violations, goal location, and comfort. To meet this desiderata in a way that the decisions are explainable, we consider a set of interpretable costs. We split costs f_i into *agent-agnostic* costs and *agent-aware* costs. At a high level, agent-agnostic costs describe the comfort, rule compliance, and progress of a candidate trajectory. Agent-aware costs evaluate the safety of the trajectories with respect to other agents using the outputs of our implicit occupancy model ψ at the query point locations \mathcal{Q} . In the next paragraph we describe our agent-aware costs at a high-level.

A *collision* cost considers the maximum probability of collision for each time step t of each trajectory candidate τ . Specifically, we gather the occupancy at the query points within the ego bounding box B_r^t , and take the maximum probability. For each trajectory, we aggregate the maximum probabilities over time steps with a cumulative sum to further penalize trajectories that collide earlier on. A *longitudinal buffer* cost penalizes trajectories with agents too close in front or behind the ego vehicle by gathering the occupancy at those locations. We apply a linear decay to the cost based on the distance with respect to the ego. Similarly, *lateral buffer* cost penalizes trajectories that remain in close lateral proximity to other agents in the scene.

D. Learning

We optimize our motion planner in two stages. We first train the implicit occupancy model to learn to perceive and forecast. In a second stage, we freeze the occupancy model and train the cost aggregation weights $\{w_i\}$ to imitate an expert driver. This two-stage training maintains the interpretability of the occupancy intermediate representation and allows the cost aggregation weights to train with stable occupancy predictions.

Occupancy: An advantage of having intermediate representations is that one can use much richer and denser supervision to perceive the world and understand its dynamics rather

than just imitating the expert trajectory. For each continuous query point $\mathbf{q} \in \mathcal{Q}$, occupancy is supervised with binary cross entropy loss. Following [1], we train with a batch of continuous query points \mathcal{Q} , uniformly sampled across the spatio-temporal volume.

Costing: We train the cost aggregation such that the behavior of our planner imitates that of an expert. Because selecting the trajectory with the minimum cost from a discrete set is not a differentiable process, we use the max-margin loss to penalize trajectories that are either unsafe or have a low cost but differ significantly from the expert driving trajectory following prior works [33], [34]. Intuitively, this loss incentivizes the expert trajectory τ_e to have a smaller cost J than all other trajectories. More precisely, our objective is

$$\mathcal{L}_w = \max_{\tau} \left[\Delta J_r(\mathbf{x}, \tau, \tau_e) + l_{\text{im}} + \sum_t [\Delta J_c^t(\mathbf{x}, \tau, \tau_e) + l_c^t]_+ \right], \quad (4)$$

where $\Delta J(\mathbf{x}, \tau, \tau_e) = J(\mathbf{x}, \tau_e) - J(\mathbf{x}, \tau)$ is the difference between the cost of the expert trajectory τ_e and the candidate trajectory τ ; J_c^t is the collision cost at a particular time step into the future, and J_r are the rest of the costs, aggregated; $[\cdot]_+$ represents the ReLU function; and l_{im} and l_c^t are the imitation and safety margins, respectively. The imitation margin is simply the distance between the trajectory waypoints in τ_e and τ , and the safety margin is whether the candidate trajectory τ collides with any ground-truth object.

As a way to mitigate distribution shift from open-loop learning to closed-loop deployment, we exploit dataset aggregation [32] by combining the initial set consisting of expert demonstrations on expert states with another dataset generated with states visited by our learned autonomy model.

IV. EXPERIMENTS

In this section, we first compare QUAD to state-of-the-art autonomy models, measuring their ability to drive in closed-loop in both safety-focused and canonical highway driving. Then, we analyze trade-off between runtime and driving quality of different methods, and ablate the importance of QUAD’s query point quantization. Finally, we show qualitative results.

A. Data

We utilize our high-fidelity end-to-end simulator to generate datasets for open-loop evaluation and training as well as a closed-loop benchmark. Our simulator can generate both LiDAR sensor data [24] as well as intelligent actor behaviors [17], [38]. We generate several datasets and benchmarks with distinct purposes.

Safety-focused set: We generate a dataset composed of ~ 900 safety-focused scenarios, split into ~ 700 for training and ~ 200 for evaluation. The scenarios in this set contain scripted interactions such as cut ins, blocked lanes, agents merging onto the highway, ego merging onto the highway through an off-ramp, and exiting through an off-ramp, aggressively slow/fast moving actors. The ego is provided with

	Mission			Safety and Compliance				Progress, Consistency and Comfort				
	GSR↑	ECR↓	PCR↓	MinTTC↑				TVR↓	Progr.↑	L2E↓	P2P↓	Jerk↓
				<i>p</i> 10	< 1s	< 2s	< 5s					
EXPERT	91.6%	0.0%	2.7%	4.64	0.0%	1.1%	17.9%	6.8%	454.2	0.0	15.1	0.30
PLANT [29]	23.7%	10.0%	2.9%	3.42	8.4%	8.9%	19.5%	63.7%	331.7	95.0	42.8	1.03
CIL [7]	0.0%	16.3%	5.6%	1.55	8.4%	15.3%	36.8%	100.0%	64.6	298.2	19.7	1.38
NMP [43]	51.6%	14.2%	21.4%	0.50	10.5%	16.3%	62.1%	18.4%	337.6	133.3	44.2	1.21
P3 [33]	76.3%	6.8%	8.3%	3.44	4.7%	5.3%	29.5%	11.6%	436.1	37.5	14.7	0.17
OccFlow [23]	60.5%	30.5%	28.0%	0.40	24.2%	28.4%	44.2%	36.8%	385.3	55.2	12.8	0.14
QUAD (Ours)	84.7%	2.1%	1.0%	4.67	1.6%	1.6%	14.2%	8.4%	430.6	36.6	15.6	0.29

TABLE I: [Safety-focused set] Closed-loop simulation results

a variety of mission routes including keep lane, lane change, or lane merge. These scripted interactions are parameterized by highly controllable parameters such as time-to-collision, time-to-arrival to a merge point, initial speeds. Training and validation utilize non-overlapping parameter values to ensure generalization is tested.

Canonical driving set: Contains a total of ~ 1700 scenarios, ~ 1400 for training and ~ 300 for evaluation, each lasting around 20 seconds. This set is generated from a set of randomly distributed discrete and continuous parameters which control the map (using a mix of simulated and real maps to get exposure to diverse curvatures, speed limits, number of lanes, and distinct topologies), ego starting location, initial agent conditions such as speed, location, and maximum acceleration limits to ensure a diverse distribution of scenarios, and a variety of vehicle types. In this set, the ego vehicle does not have a particular goal or mission (in contrast to safety-focused set), making the task easier. Similar to the safety-focused set, training and evaluation are non-overlapping by ensuring different random parameter values.

Open-loop vs. closed-loop: Both safety-focused and canonical scenario sets can be used for open-loop as well as closed-loop execution. In open-loop, an expert planner is in charge of driving and the policy under train/test is only proposing plans that do not get executed and hence do not affect the next state of the world. The expert is a privileged planner that has access to the internal state of the simulation, receiving ground truth current states for other actors in the world as well as their most recent future plan. In contrast, during closed-loop simulation, the planner policy under test is driving the ego vehicle. In other words, the states the ego vehicle visits are dependent on the policy under test. We rely on closed-loop simulation as our ultimate benchmark, since it best captures the driving performance a planner would have in the real-world, where the planner actions affect the environment.

B. Experimental setup

Baselines: There are three main families of baselines we consider. (1) *Object-based autonomy* composed of separate object detection and planning modules. We compare against PlanT [29], a recent method that leverages an object detector and then a Transformer-based planner that reasons about the

detected objects together with the route. (2) *Sensor-to-plan autonomy* approaches including direct regression of the plan through Conditional Imitation Learning (CIL) [7] as well as NMP [43], which learns a cost-prediction network to rank samples. (3) *Occupancy-based autonomy* that leverages grid-based occupancy and a sample-based planner [34]. We consider OccFlow [23] and P3 [33] as two possible architectures to predict temporal occupancy grids.

For P3 and OccFlow we train in two stages. In the first stage, we train the occupancy module on nominal scenarios with dense traffic. In the second stage, we train the planner on safety focused scenarios. For the remaining baselines PlanT, CIL, and NMP since they are trained in a single stage we train the entire module on the combined set of canonical and safety training scenarios.

Metrics: We consider a comprehensive set of metrics to provide a holistic view of autonomy performance.

To evaluate safety we propose the following metrics. *Execution Collision Rate (ECR)* measures the percentage of scenarios with a collision. *Plan Collision Rate (PCR)* calculates the collisions of our planned trajectory with respect to the simulated actors’ future trajectories. *Minimum Time-To-Collision (MinTTC)* computes the minimum time buffer (in seconds) –across a scenario– before a collision occurs, taking into account the ego plan and the actor plan. To better illustrate the worst-case we compute percentile 10 (*p*10), and to get an idea of different levels of risk exposure we report cumulative TTC buckets ($< 1s$, $< 2s$, $< 5s$), in decreasing order of severity. To measure compliance, we use *Traffic Violation Rate (TVR)* to show the percentage of scenarios where the ego vehicle violates a lane boundary or speed limit, goes off-road, or collides.

On top of safety and compliance, it is also important driving makes sufficient progress towards the goal and is smooth. We compute *progress within map (Progr.)* as the meters traveled without going off-road. *L2 distance to Expert (L2E)* measures distance from the expert demonstrations. In closed-loop, this corresponds to the distance between the executed trajectories. *Plan-to-plan consistency (P2P)* computes the distance over the finite horizon planned trajectories for consecutive plans. *Jerk* measures the level of discomfort of the planned trajectories. Note secondary metrics are only meaningful when primary metrics are similar, since avoiding collisions naturally causes higher jerk and plan-to-plan dis-

	ECR↓	PCR↓	MinTTC↑ <i>p</i> 10	TVR↓	Progr.↑	L2E↓	Jerk↓
EXPERT	0.0%	0.0%	10.00	2.2%	317.6	0.0	0.40
PLANt	29.6%	18.2%	0.40	86.6%	214.6	66.3	1.08
CIL	49.0%	28.8%	0.40	92.0%	50.7	255.1	1.35
NMP	7.0%	22.8%	1.67	8.3%	214.9	80.6	1.20
P3	2.9%	6.0%	3.88	4.8%	317.3	36.2	0.28
OCCFLOW	20.4%	20.9%	0.40	22.9%	306.0	39.6	0.23
QUAD (Ours)	0.0%	0.5%	4.64	2.2%	299.1	33.7	0.43

TABLE II: [Canonical set] Closed-loop simulation results




Fig. 3: Driving quality vs. runtime comparison

tance thus safety and compliance must be prioritized at all times. Finally, for self-driving vehicles it is critical to reach a target location. To measure this, we propose *Goal Success Rate (GSR)* as the ratio of scenarios where the ego vehicle reaches the goal without colliding or violating traffic rules.

C. Results

Comparison against state-of-the-art: Table I shows our main benchmark results, where we evaluate the baselines and our method in closed-loop, on the safety-focused highly-interactive scenarios. We find our method is much safer than the baselines, as showcased by lower ECR, lower PCR as well as higher TTC. Moreover, it complies better with the rules of the road as showcased by our lower TVR. On top of being safer and more compliant, QUAD achieves the highest GSR indicating the maneuvers we pick at a behavioral level are effective. On the secondary metrics our model ranks high in progress, best in expert imitation, and has fairly consistent plans over time. While we do observe higher jerk than some baselines, our jerk is closer to that exhibited by the expert.

Table II showcases our closed-loop results in canonical driving scenarios. Similar to previous results, we observe a substantial safety improvement, being able to attain very close safety and compliance performance to that of the expert. QUAD has the closest imitation to expert and ranks highly on progress too, striking a good balance between safety and progress towards the goal. When taking into consideration the priority safety has over progress, our method clearly outperforms the baselines. Note that CIL and PlanT do not use lane based trajectory samples and this results in much worse TVR metric due to lane boundary violations.

Inference runtime: Fig. 3 shows how methods compare when considering the balance between driving quality and runtime, an important factor for real-world deployment. We profile in a crowded scenario with 74 agents and 11 lanes to stress-test the different autonomy systems. We highlight that QUAD is faster than other occupancy-based autonomy




Fig. 4: Qualitative results. We visualize the LiDAR point cloud, map, predicted occupancy (by querying ψ at a regular grid, solely for illustration purposes), and cost associated with the trajectory samples. From top to bottom: a lane change, a re-incorporation near an off-ramp, and a merge.

models (P3, OCCFLOW), while attaining the best driving quality both in terms of mission achievement and safety.

Query point quantization ablation: Our proposed query point quantization is necessary to reap the benefits from an implicit occupancy decoder without inquiring a prohibitive runtime. The naive alternative of using the continuous query points directly after the generation step incurs very high inference times of ~ 700 ms. This is due to the high number of points, caused by the high overlap between the candidate plans especially at the beginning of the temporal horizon. Quantization becomes necessary in practice. Utilizing a resolution of 0.5m (the one used in the rest of the experiments), autonomy runtime reduces to ~ 250 ms while the execution collision rate only increases from $\sim 1\%$ to $\sim 2\%$, thus providing significant improvements over the baselines.

Qualitative results: Fig. 4 showcases example scenarios where the goals instructed to the ego consist of semantically diverse and interesting maneuvers. From the cost distribution, we can observe that our planner can understand the mission route and progress towards the goal, it can speed up to pass slow moving vehicles when instructed to lane-change, and can plan smooth merge trajectories at on-ramps.

V. CONCLUSION

In this paper, we have proposed an interpretable motion planner leveraging spatio-temporal occupancy queries to effectively understand the current and future free-space from sensor data. We showcased our proposed autonomy can drive more safely and progress further than contemporary object-based, sensor-to-plan and occupancy-based autonomy models. Our framework also achieves faster runtime than its closest competitors.

REFERENCES

- [1] Ben Agro, Quin Sykora, Sergio Casas, and Raquel Urtasun. Implicit occupancy flow fields for perception and prediction in self-driving. *CVPR*, 2023. 2, 3, 4
- [2] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseen Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint*, 2016. 2
- [3] Sergio Casas, Cole Gulino, Simon Suo, Katie Luo, Renjie Liao, and Raquel Urtasun. Implicit latent variable model for scene-consistent motion forecasting. In *ECCV*, 2020. 2
- [4] Sergio Casas, Wenjie Luo, and Raquel Urtasun. Intentnet: Learning to predict intention from raw sensor data. In *CoRL*, 2018. 2, 3
- [5] Sergio Casas, Abbas Sadat, and Raquel Urtasun. Mp3: A unified model to map, perceive, predict and plan. In *CVPR*, 2021. 1, 2
- [6] Yuning Chai, Benjamin Sapp, Mayank Bansal, and Dragomir Anguelov. Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. *CoRL*, 2019. 2
- [7] Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *ICRA*, 2018. 1, 2, 5
- [8] Felipe Codevilla, Eder Santana, Antonio M López, and Adrien Gaidon. Exploring the limitations of behavior cloning for autonomous driving. In *ICCV*, 2019. 2
- [9] Alexander Cui, Sergio Casas, Abbas Sadat, Renjie Liao, and Raquel Urtasun. Lookout: Diverse multi-future prediction and planning for self-driving. In *ICCV*, 2021. 1
- [10] Haoyang Fan, Zhongpu Xia, Changchun Liu, Yaqin Chen, and Qi Kong. An auto-tuning framework for autonomous vehicles. *arXiv preprint*, 2018. 2
- [11] Haoyang Fan, Fan Zhu, Changchun Liu, Liangliang Zhang, Li Zhuang, Dong Li, Weicheng Zhu, Jiangtao Hu, Hongye Li, and Qi Kong. Baidu apollo em motion planner. *arXiv preprint*, 2018. 1
- [12] Jeffrey Hawke, Richard Shen, Corina Gurau, Siddharth Sharma, Daniele Reda, Nikolay Nikolov, Przemyslaw Mazur, Sean Micklethwaite, Nicolas Griffiths, Amar Shah, et al. Urban driving with conditional imitation learning. *arXiv preprint*, 2019. 2
- [13] Stefan Hoermann, Martin Bach, and Klaus Dietmayer. Dynamic occupancy grid prediction for urban autonomous driving: A deep learning approach with fully automatic labeling. In *ICRA*, 2018. 1, 2
- [14] Anthony Hu, Gianluca Corrado, Nicolas Griffiths, Zak Murez, Corina Gurau, Hudson Yeo, Alex Kendall, Roberto Cipolla, and Jamie Shotton. Model-based imitation learning for urban driving. *arXiv preprint arXiv:2210.07729*, 2022. 1, 2
- [15] 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. 2
- [16] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, Lewei Lu, Xiaosong Jia, Qiang Liu, Jifeng Dai, Yu Qiao, and Hongyang Li. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 2
- [17] Arne Kesting, Martin Treiber, and Dirk Helbing. General lane-changing model mobil for car-following models. *Transportation Research Record*, 1999(1), 2007. 4
- [18] Jason Kong, Mark Pfeiffer, Georg Schilb, and Francesco Borrelli. Kinematic and dynamic vehicle models for autonomous driving control design. In *2015 IEEE intelligent vehicles symposium (IV)*, pages 1094–1099. IEEE, 2015. 3
- [19] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *CVPR*, 2019. 2
- [20] Ming Liang, Bin Yang, Wenyuan Zeng, Yun Chen, Rui Hu, Sergio Casas, and Raquel Urtasun. Pnpnet: End-to-end perception and prediction with tracking in the loop. In *CVPR (CVPR)*, June 2020. 2
- [21] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 3
- [22] Wenjie Luo, Bin Yang, and Raquel Urtasun. Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. In *CVPR*, 2018. 2
- [23] Reza Mahjourian, Jinkyu Kim, Yuning Chai, Mingxing Tan, Ben Sapp, and Dragomir Anguelov. Occupancy flow fields for motion forecasting in autonomous driving. *IEEE Robotics and Automation Letters*, 2022. 1, 2, 5
- [24] Sivabalan Manivasagam, Shenlong Wang, Kelvin Wong, Wenyuan Zeng, Mikita Sazanovich, Shuhan Tan, Bin Yang, Wei-Chiu Ma, and Raquel Urtasun. Lidarsim: Realistic lidar simulation by leveraging the real world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11167–11176, 2020. 4
- [25] Matthias Müller, Alexey Dosovitskiy, Bernard Ghanem, and Vladen Koltun. Driving policy transfer via modularity and abstraction. *arXiv preprint*, 2018. 2
- [26] Tung Phan-Minh, Elena Corina Grigore, Freddy A Boulton, Oscar Beijbom, and Eric M Wolff. Governet: Multimodal behavior prediction using trajectory sets. *arXiv preprint*, 2019. 2
- [27] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *ECCV*, 2020. 1, 2
- [28] Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. In *NIPS*, 1989. 2
- [29] Katrin Renz, Kashyap Chitta, Otniel-Bogdan Mercea, A Koepke, Zeynep Akata, and Andreas Geiger. Plant: Explainable planning transformers via object-level representations. *arXiv preprint arXiv:2210.14222*, 2022. 1, 2, 5
- [30] Nicholas Rhinehart, Jeff He, Charles Packer, Matthew A Wright, Rowan McAllister, Joseph E Gonzalez, and Sergey Levine. Contingencies from observations: Tractable contingency planning with learned behavior models. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13663–13669. IEEE, 2021. 2
- [31] Nicholas Rhinehart, Rowan McAllister, and Sergey Levine. Deep imitative models for flexible inference, planning, and control. *arXiv preprint*, 2018. 2
- [32] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011. 1, 2, 4
- [33] Abbas Sadat, Sergio Casas, Mengye Ren, Xinyu Wu, Pranaab Dhawan, and Raquel Urtasun. Perceive, predict, and plan: Safe motion planning through interpretable semantic representations. In *ECCV*, 2020. 1, 2, 4, 5
- [34] Abbas Sadat, Mengye Ren, Andrei Pokrovsky, Yen-Chen Lin, Ersin Yumer, and Raquel Urtasun. Jointly learnable behavior and trajectory planning for self-driving vehicles. In *IRIOS*, 2018. 2, 3, 4, 5
- [35] Sarthak Sharma, Junaid Ahmed Ansari, J Krishna Murthy, and K Madhava Krishna. Beyond pixels: Leveraging geometry and shape cues for online multi-object tracking. In *ICRA*, 2018. 2
- [36] Charlie Tang and Russ R Salakhutdinov. Multiple futures prediction. In *Advances in Neural Information Processing Systems*, 2019. 2
- [37] Ameni Trabelsi, Ross J Beveridge, and Nathaniel Blanchard. Drowned out by the noise: Evidence for tracking-free motion prediction. *arXiv preprint*, 2021. 2
- [38] Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62(2), aug 2000. 4
- [39] Xinhua Weng, Jianren Wang, David Held, and Kris Kitani. 3d multi-object tracking: A baseline and new evaluation metrics. In *2020 IEEE/RSSJ IROS*, 2020. 2
- [40] Moritz Werling, Julius Ziegler, Sören Kammel, and Sebastian Thrun. Optimal trajectory generation for dynamic street scenarios in a frenet frame. In *ICRA*, 2010. 3
- [41] Bin Yang, Ming Liang, and Raquel Urtasun. Hdnet: Exploiting hd maps for 3d object detection. In *CoRL*, 2018. 3
- [42] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3d object detection from point clouds. In *CVPR*, 2018. 2
- [43] Wenyuan Zeng, Wenjie Luo, Simon Suo, Abbas Sadat, Bin Yang, Sergio Casas, and Raquel Urtasun. End-to-end interpretable neural motion planner. In *CVPR*, 2019. 1, 2, 5
- [44] Wenyuan Zeng, Shenlong Wang, Renjie Liao, Yun Chen, Bin Yang, and Raquel Urtasun. Dsdnet: Deep structured self-driving network. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 156–172. Springer, 2020. 1
- [45] Chris Zhang, Runsheng Guo, Wenyuan Zeng, Yuwen Xiong, Binbin Dai, Rui Hu, Mengye Ren, and Raquel Urtasun. Rethinking closed-loop training for autonomous driving. In *ECCV 2022*. 2
- [46] Tianyang Zhao, Yifei Xu, Mathew Monfort, Wongun Choi, Chris Baker, Yibiao Zhao, Yizhou Wang, and Ying Nian Wu. Multi-agent tensor for contextual trajectory prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*,

2019. [2](#)
- [47] 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. [2](#), [4](#)