

Supplementary Material for Parting with Misconceptions about Learning-based Vehicle Motion Planning

Anonymous Author(s)

Affiliation

Address

email

Abstract: In this **supplementary document**, we first provide additional details about the nuPlan dataset and evaluation schemes. Next, we provide information on the on our standardized benchmark for training and evaluation. Following this, we describe the implementation and training schedules for our baselines. Additionally, we provide detailed descriptions of PDM’s components. Finally, we describe our preliminary version of PDM-Hybrid, which won the 2023 nuPlan Challenge. The **supplementary video** compares the 3 evaluation modes in nuPlan and visualizes the operation of rule-based, learned and hybrid planners on different types of scenarios.

1 nuPlan

1.1 Dataset

The nuPlan dataset comprises 1300 hours of driving data from Las Vegas, Boston, Pittsburgh, and Singapore. The data logs are separated into three sets for training, validation, and testing. The data includes human driving logs, HD map information, and auto-labelled data from an offline perception system. Based on the logs, short driving scenarios of 15s duration can be generated to simulate and measure the capabilities of various planners in open- and closed-loop. The scenarios are categorized into 70 types, of which 14 were used for the 2023 nuPlan challenge. The planner receives (1) a goal position, (2) a sequence of roadblocks to follow, and (3) observations over the past 2s. A planner is required to output an 8s trajectory.

1.2 Simulation

A planner in nuPlan can be simulated in three modes: open-loop, closed-loop non-reactive, closed-loop reactive.

Open-Loop. During open-loop simulation, the ego vehicle is completely replayed from the human recordings. Thus, the planner’s output does not influence the vehicle’s movement.

Closed-Loop. To simulate the planner in closed-loop, nuPlan uses a two-stage simulation pipeline. First, an LQR controller [1] converts the trajectory into actions. Second, the actions are applied to propagate a kinematic bicycle model [2, 3], approximating the vehicle’s motion in coordinate space. The controller and motion model parameters are fixed and cannot be tuned for the challenge. The environment is either non-reactive (replaying the recordings) or reactive by simulating the surrounding vehicles with a lightweight planner.

31 2 Training & Evaluation

32 In this section, we describe our training and evaluation benchmark. We will publicly release the
33 configurations and settings required to reproduce our benchmark.

34 **Training.** Our proposed training set contains a maximum of 4k samples per scenario, using all
35 70 scenario types from nuPlan’s "training" database files. Overall, this dataset consists of 177,435
36 training samples.

37 **Evaluation.** We evaluate the planners with a maximum of 100 scenarios of the 14 challenge types
38 from nuPlan’s validation split, resulting in 1,118 scenarios. To ensure that the scenarios do not
39 overlap, we use a threshold of 15 seconds between the initial frames of subsequent scenarios. At the
40 current stage, we strongly recommend evaluating planners on the validation split and preserving the
41 test split’s integrity. We found our benchmark well aligned with the 2023 nuPlan leaderboard.

42 **Hardware.** For simulation, training, and runtime analysis, we use an AMD Ryzen 9 7950X
43 processor with 64GB memory and a single NVIDIA RTX 3090.

44 3 Baselines

45 In this section, we provide detailed information about the training and implementation of the baseline
46 planners presented in our work.

47 3.1 Learned Planner

48 All learning-based planners are trained on our standardized training set, as described in Section 2.
49 The models output an 8s trajectory at 2Hz.

50 **PlanCNN.** The PlanCNN model is provided in the framework [4]. We use a pre-trained ResNet-50
51 as CNN encoder. A single linear layer decodes the trajectory. We train PlanCNN with the Adam
52 optimizer, a batch size of 64, and a learning rate of $1e^{-4}$ for 100 epochs. We divide the learning rate
53 by 10 after 50 and 75 epochs. The CNN is trained with an L1-Loss.

54 **Urban Driver.** The Urban Driver model was introduced in [5] and uses PointNet layers to
55 encode the ego vehicles motion history, surrounding agents, and the map, which is represented by
56 polylines. Subsequently, an attention layer aggregates the encoded features before decoding the
57 output trajectory. We use the implementation provided by the nuPlan framework [4] and train the
58 model with an L1-Loss using a batch size of 64 and a learning rate of $1e^{-4}$ for 60 epochs.

59 **GC-PGP.** The GC-PGP model was introduced in [6], and extends the state-of-the-art prediction
60 model, called PGP [7], for goal-directed ego-forecasting. We refer to Section 5 for a detailed de-
61 scription of the model. Note that a modified variant of GC-PGP was used in our preliminary planner.
62 We train the model as proposed in [6] for 70 epochs with a batch size of 32 and a learning rate of
63 $1e^{-4}$ that is decayed after 40, 50, and 55 epochs by a factor of 0.5.

64 3.2 IDM

65 The IDM planner utilizes the centerline of lanes as lateral path, and applies the policy from [8] for
66 longitudinal control.

67 **Lateral:.** The planner retrieves a sequence of lanes, by applying a Breadth-First-Search (BFS) from
68 the current lane to any lane at the end of the route. If no route is found, the longest lane sequence is
69 returned.

70 **Longitudinal.** IDM iteratively applies a policy to calculate the longitudinal velocity \dot{x}_α and accel-
71 eration \dot{v}_α for the ego vehicle α . By integrating the velocities over time, the planner retrieves the
72 longitudinal ego position x_α , which is interpolated along the centerline to calculate the trajectory
73 samples. Since IDM is a parameterized car-following model, each unrolling step requires to extract

Parameter	Value	Description
v_0	v_{lane}	Desired velocity. Either the current speed-limit, or $v_{\text{lane}} = 10$ m/s if speed-limit not available.
s_0	1.0 m	Desired net distance to the leading agent $\alpha - 1$.
T	1.5 s	Desired time headway to leading agent $\alpha - 1$.
a	1.0 m/s ²	Maximum acceleration of ego vehicle α
b	3.0 m/s ²	Maximum deceleration (positive) of ego vehicle α
δ	4.0	Acceleration exponent.

Table 1: IDM Parameters.

the states of the leading agent $\alpha - 1$, resulting in the net distance s_α and approaching rate Δv_α :

$$s_\alpha := x_{\alpha-1} - x_\alpha - l_{\alpha-1}, \quad (1)$$

$$\Delta v_\alpha := v_\alpha - v_{\alpha-1}, \quad (2)$$

where $l_{\alpha-1}$ is the length of the leading vehicle. Finally, the IDM output can be expressed by the following equations:

$$\dot{x}_\alpha = \frac{dx_\alpha}{dt} = v_\alpha \quad (3)$$

$$\dot{v}_\alpha = \frac{dv_\alpha}{dt} = a \left(1 - \left(\frac{v_\alpha}{v_0} \right)^\delta - \left(\frac{s^*(v_\alpha, \Delta v_\alpha)}{s_\alpha} \right)^2 \right) \quad (4)$$

$$\text{with } s^*(v_\alpha, \Delta v_\alpha) = s_0 + v_\alpha T + \frac{v_\alpha \Delta v_\alpha}{2\sqrt{ab}} \quad (5)$$

where parameters in red are manually selected and summarized in Table 1. Intersections on the route with red traffic lights are considered to be stationary obstacles, to obey traffic rules.

4 PDM

4.1 PDM-Closed

Path-Planning. Instead of BFS, PDM-Closed utilizes Dijkstra, with the lane length as edge-weights, to search a sequence of lanes for centerline extraction. We found Dijkstra slightly more suitable due to the avoidance of detours while having no substantial effect on run-time.

Observation & Forecasting. PDM-Closed generates a forecast of dynamic agents for the planning horizon of 8s with a temporal resolution of 2Hz, and stores the bounding boxes together with static obstacles in occupancy maps. We only consider the nearest 50 vehicles, 10 pedestrians, 10 bicycles, and 50 static objects to the ego agent. Thereby, the planner avoids exploding computation costs when being nearby a large number of entities (e.g. a crowd of pedestrians). Like IDM, we add intersections on the route with a red traffic light as stationary objects.

Proposals. We generate proposals by pairing 3 centerline offsets, and 5 IDM policies at varying target speeds, resulting in 15 proposals. The parameters are summarized in Table 2. Note that we use higher acceleration parameters to favor progress. We iteratively unroll the proposals for 4s at 10Hz, where we update the leading agent states at 5Hz. The shorter proposal horizon and lower frequency of updating the leading agent help in reducing computation cost.

Simulation. We simulate all trajectory proposals for 4s at 10Hz using a faster re-implementation of nuPlan’s LQR controller and kinematic bicycle model. Thereby, the proposals are converted into the actual expected trajectory in closed-loop.

Scoring. Our scoring function closely resembles the nuPlan closed-loop metrics [9]. However, we leverage a computationally efficient re-implementation of the metrics to meet the strict runtime requirements of the competition. The scoring considers at-fault collisions, drivable area infractions, and

Parameter	Value	Description
o	$\{-1, 0, 1\}$ m	Centerline offsets
v_0	$\{\frac{i}{5} v_{\text{lane}}\}_{i=1, \dots, 5}$	Desired velocity. Either the current speed-limit, or $v_{\text{lane}} = 15$ m/s if speed-limit not available.
s_0	1.0 m	Desired net distance to the leading agent $\alpha - 1$.
T	1.5 s	Desired time headway to leading agent $\alpha - 1$.
a	1.5 ms^{-2}	Maximum acceleration of ego vehicle α
b	3.0 ms^{-2}	Maximum deceleration (positive) of ego vehicle α
δ	10.0	Acceleration exponent.

Table 2: PDM-Closed Parameters.

driving direction compliance as multiplicative metrics. Furthermore, the scoring evaluates progress, time-to-collision, and comfortability as weighted metrics. We normalize the progress metric with the highest progress of a proposal which is free of multiplicative infractions. We use the same weights as nuPlan, but ignore speed-limit compliance and the binary no-progress metric, since the IDM proposals are naturally bound to comply with the current speed limit and the no-progress metric cannot be evaluated without privileged knowledge of the human expert’s behavior.

Trajectory Selection. Finally, PDM-Closed outputs the highest-scoring proposal, which is extended to the entire planning horizon of 8s with the corresponding IDM policy. If the best trajectory is expected to collide within 2s, the output is overwritten with an emergency brake maneuver.

4.2 Architectures & Training

We propose two MLP models to investigate open-loop scoring: ϕ_{open} and ϕ_{offset} . Both models share architectural attributes and the same training schedule.

Architecture. The ϕ_{open} receives a 120m centerline with 1m resolution (c), and the ego history states (h) consisting of past waypoints and longitudinal, lateral, and angular velocity and acceleration, for the past 2s at 5Hz. Additionally, ϕ_{offset} receives the trajectory (w_{closed}) of PDM-Closed as input, with an 8s duration downsampled to 5Hz. The input features (c , h , and optionally w_{closed}) are first projected to a 512-dimensional vector with linear layers and then concatenated and forwarded into the MLP. Hereinafter, we apply the same architecture proposed in [10], consisting of two 512-dimensional linear layers and dropout. We use a ReLU activation function for all linear layers. The ϕ_{open} model outputs waypoints w_{open} relative to the ego position, whereas ϕ_{offset} predicts offsets to w_{closed} .

Training. We train ϕ_{open} and ϕ_{offset} on our standardized training split, as described in Section 2. The models with the Adam optimizer, a batch size of 64, and a learning rate of $1e^{-4}$ for 100 epochs. We divide the learning rate by 10 after 50 and 75 epochs. We use an L1-Loss function for both models. Note that we add w_{closed} to the output of ϕ_{offset} during training.

5 2023 nuPlan Challenge

Our preliminary version, which won the 2023 nuPlan competition, was composed of PDM-Closed and used a modified version of GC-PGP for long-term trajectory correction. The code for our preliminary version and for GC-PGP will be publicly released.

Goal-conditioned Ego-Forecasting via Graph-based Policy. GC-PGP extends the state-of-the-art prediction model, called PGP [7], for goal-directed ego-forecasting.

The model receives an ego-centered lane-graph representation, together with observed states of surrounding agents and the ego vehicle. The nodes in the lane-graph compromise polylines of similar length, with directed edges for lanes in proximity or the direction of traffic flow. The lane-nodes and dynamics of the surrounding agents and the ego vehicle are encoded with separate Gated Recurrent

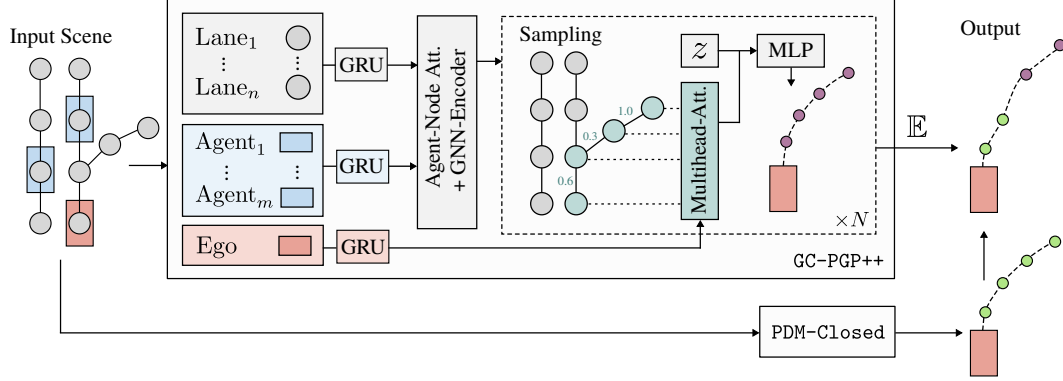


Figure 1: **PDM-Hybrid (Graph)**. GC-PGP++ encodes the input lane-graph, and generates trajectory proposals based on lane-traversals. The long-term trajectory of PDM-Closed is corrected by the expectation of GC-PGP++.

Method	OLS \uparrow	CLS-R \uparrow	CLS-NR \uparrow	Runtime (ms) \downarrow
GC-PGP [6]	82	54	57	100
GC-PGP++	84	49	50	84
PDM-Hybrid (Graph)	84	92	93	172
PDM-Hybrid (Centerline)	84	92	93	96

Table 3: **Val14 Benchmark**. The preliminary PDM-Hybrid (Graph) planner integrates the improved GC-PGP++ for long-term correction, leading to the same performance as PDM-Hybrid.

Units (GRUs). The model aggregates the information by applying Agent-to-Node Attention and Graph Neural Network (GNN) layers, yielding a per-node feature representation.

The node-feature are used to estimate transition probabilities for outgoing edges. Subsequently, traversals across the lane-graph are sampled. During inference, GC-PGP masks out off-route edges to ensure goal-compliant traversals. Then, a latent-variable model decodes trajectories based on the traversals and the ego-motion encoding. The output trajectories are obtained after a k -means clustering. The original version of GC-PGP selects the cluster centers with the highest rank as output trajectory.

Modifications. We observe that the goal-conditioned trajectories do not describe disjoint behaviors. Hence we omit the k -means clustering and instead calculate the expectation of all decoded trajectories by averaging, as shown in Fig. 1. Our modification paired with the hard route constraints of GC-PGP lead to a small OLS performance increase while being computationally more efficient. We refer to our modified version as GC-PGP++.

References

- [1] Y. Tassa, N. Mansard, and E. Todorov. Control-limited differential dynamic programming. In *Proc. IEEE International Conf. on Robotics and Automation (ICRA)*, 2014.
- [2] R. Rajamani. *Vehicle dynamics and control*. Springer Science & Business Media, 2011.
- [3] P. Polack, F. Althché, B. d’Andréa Novel, and A. de La Fortelle. The kinematic bicycle model: A consistent model for planning feasible trajectories for autonomous vehicles? In *Proc. IEEE Intelligent Vehicles Symposium (IV)*, 2017.
- [4] H. Caesar, J. Kabzan, K. S. Tan, W. K. Fong, E. M. Wolff, A. H. Lang, L. Fletcher, O. Beijbom, and S. Omari. nuplan: A closed-loop ml-based planning benchmark for autonomous vehicles. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2021.
- [5] O. Scheel, L. Bergamini, M. Wolczyk, B. Osipiński, and P. Ondruska. Urban driver: Learning to drive from real-world demonstrations using policy gradients. In *Proc. Conf. on Robot Learning (CoRL)*, 2021.
- [6] M. Hallgarten, M. Stoll, and A. Zell. From Prediction to Planning With Goal Conditioned Lane Graph Traversals. *arXiv.org*, 2302.07753, 2023.
- [7] N. Deo, E. M. Wolff, and O. Beijbom. Multimodal Trajectory Prediction Conditioned on Lane-Graph Traversals. In *Proc. Conf. on Robot Learning (CoRL)*, 2021.
- [8] M. Treiber, A. Hennecke, and D. Helbing. Congested traffic states in empirical observations and microscopic simulations. *Physical review E*, 2000.
- [9] Motional. nuplan metrics, 2023. URL https://nuplan-devkit.readthedocs.io/en/latest/metrics_description.html.
- [10] J.-T. Zhai, Z. Feng, J. Du, Y. Mao, J.-J. Liu, Z. Tan, Y. Zhang, X. Ye, and J. Wang. Re-thinking the open-loop evaluation of end-to-end autonomous driving in nuscenese. *arXiv.org*, 2305.10430, 2023.