

CCIL: Context-conditioned imitation learning for urban driving

Ke Guo*[†], Wei Jing[†], Junbo Chen[†], Jia Pan*

*The University of Hong Kong, Email: {kguo, jpan}@cs.hku.hk

[†] Alibaba Group, Email: 21wjing@gmail.com, chjb@zju.edu.cn

Abstract—Imitation learning holds great promise for addressing the complex task of autonomous urban driving, as experienced human drivers can navigate highly challenging scenarios with ease. While behavior cloning is a widely used imitation learning approach in autonomous driving due to its exemption from risky online interactions, it suffers from the covariate shift issue. To address this limitation, we propose a context-conditioned imitation learning approach that employs a policy to map the context state into the ego vehicle’s future trajectory, rather than relying on the traditional formulation of both ego and context states to predict the ego action. Additionally, to reduce the implicit ego information in the coordinate system, we design an ego-perturbed goal-oriented coordinate system. The origin of this coordinate system is the ego vehicle’s position plus a zero mean Gaussian perturbation, and the x-axis direction points towards its goal position. Our experiments on the real-world large-scale Lyft and nuPlan datasets show that our method significantly outperforms state-of-the-art approaches.

I. INTRODUCTION

Planning a safe, comfortable, and efficient trajectory for a self-driving vehicle (SDV) in complex urban environments is a challenging and critical task in autonomous driving [1]. Unlike highway driving [2], urban driving requires handling various road geometries, such as roundabouts and intersections, while interacting with traffic lights, pedestrians, and other vehicles. Conventional rule-based approaches [3] have achieved some success in industry but require extensive human engineering to deal with diverse real-world scenarios. Recent advances in deep learning techniques have motivated researchers [4, 5, 6] to employ neural networks to model complex driving policies. Imitation learning (IL) from human drivers’ demonstrations is a promising solution for learning these policies, as experienced drivers can handle even the most difficult situations, and their driving data can be easily collected at scale.

The simplest IL algorithm is the behavior cloning (BC) [7] method, which has wide applications in autonomous driving [8, 4, 9]. It learns a policy in a supervised fashion by minimizing the difference between the actions taken by the learner and those taken by an expert in the expert state distribution without potentially dangerous online interactions. Despite its simplicity, the BC method suffers from the *covariate shift* issue [10], *i.e.* the state induced by the learner’s policy cumulatively deviates from the expert’s distribution.

To overcome the *covariate shift* problem, existing methods such as DAGger [10] and DART [11] query supervisor corrections at the learner’s or perturbed expert’s states. Since human supervision is hard to collect, recent works like GAIL [12]

seek to provide feedback from a neural network-based discriminator to recover from out-of-distribution states induced by the learner’s policy. However, these data augmentation methods need either expert supervision or rolling out the learner’s policy in the real world, which is impractical in autonomous driving. Alternatively, some researchers have attempted to constrain the learned policy formulation to ensure its robustness to policy errors by incorporating control theoretic prior knowledge, as real-world systems usually have a robustness property. For example, Palan et al. [13], Havens and Hu [14] impose Kalman or linear matrix inequality constraints on learned linear policies to ensure closed-loop stability in a linear time-invariant system. Yin et al. [15] extends this methodology by formulating the policy as a simple feed-forward neural network. Furthermore, East [16] expands the method proposed in [14] to include polynomial policies and dynamical systems. However, the urban driving task is too complex to be handled by these naive linear or polynomial policy formulations.

To learn a stable and general urban driving policy by imitating offline human demonstrations, we propose a **context-conditioned imitation learning** (CCIL) method. Unlike BC, our approach utilizes a policy network that predicts the SDV’s future trajectory using only its context state, rather than taking both the ego and context state as inputs to generate the next action [17, 18]. In our method, the ego state represents the SDV’s historical trajectory, while the context state encompasses the states of all other observed objects and the goal positions of the SDV.

Our approach is primarily motivated by the fact that the ego state is highly susceptible to policy errors. Even a small distribution shift in policy inputs can result in greater action errors and eventually lead to out-of-distribution states. Besides, in the autonomous driving task, static elements in the context, such as lanes or crosswalks, remain unaffected by the SDV, while dynamic elements like human drivers attempt to recover from the SDV’s perturbation. This stability property in the traffic system can be leveraged to address the distribution shift issue by considering only the context as policy input. In addition, removing the ego state from policy inputs can also help overcome the inertia problem caused by causal confusion [19]. For example, when the ego vehicle comes to a stop, the training data often shows a high probability of it remaining static. This leads to the formation of a spurious correlation between low speed and no acceleration, making it challenging to restart under the imitative policy. However, our

method masks the ego history, thus preventing the formation of such spurious correlations.

However, simply removing the ego historical information from the policy inputs is not enough in practice. The policy inputs are usually transformed into the vehicle coordinate system, with the SDV’s rear axle midpoint as the origin and orientation as the x -axis direction [4, 20]. However, using such an ego-centric coordinate system can still implicitly leak ego information. This is because the observation distributions in the ego-centric coordinate can vary based on the choice of the coordinate origin, which planning policies can exploit to infer the ego information. An analogy with the non-inertial reference frame in physics can help in understanding this concept. In an ego-centric frame of reference, a moving subject (such as a car) that experiences acceleration can deduce some movement information about itself from local observations, thereby leaking ego motion information. To minimize such implicit leaked information, we develop an ego-perturbed goal-oriented coordinate system that is less influenced by ego motion. The origin of this coordinate system is the SDV’s current position plus a zero mean Gaussian perturbation, and the x -axis direction points towards its goal position.

The main contributions of our paper can be summarized as follows:

1. We present a novel context-conditioned imitation learning method to address the *covariate shift* issue in offline imitation learning. Our method learns a policy that predicts the SDV’s future trajectory using only its context as input, with a robustness assurance based on the assumption of context stability.

2. We apply our approach specifically to urban driving by removing the explicit ego state information in policy input and proposing the ego-perturbed goal-oriented coordinate system to minimize implicit ego information in the coordinate system.

3. We verify the effectiveness of our method on the real-world large-scale urban driving datasets, Lyft [21] and nuPlan [22], achieving state-of-the-art performance benchmarks. The videos and code for our method can be found at <https://sites.google.com/view/contextconditionedil>.

II. RELATED WORK

A. Imitation learning for autonomous driving

The objective of applying IL in autonomous driving is to teach an autonomous vehicle to drive by mimicking human drivers’ behavior. The most straightforward approach is BC, which minimizes the difference between the learner’s and expert’s actions in the expert states without requiring additional manually labeled data or online interaction. Early BC applications in autonomous driving, such as ALVINN [8] and PilotNet [4], used a large amount of human driving experience to learn an end-to-end policy that directly maps sensor inputs to vehicle control commands. Recently, ChauffeurNet [18] improved generalization and transparency by yielding intermediate planning using perception results. However, the BC approach often leads to a *covariate shift* between the training distribution and the deployment distribution. This means that even minor errors in the policy can cause the vehicle to deviate

from the expert’s state, leading to larger errors. To address this challenge, existing IL methods can be categorized into online methods, offline model-free methods, and offline model-based methods.

- 1) *Online IL*: Online methods aim to directly match the expert state-action distribution instead of matching the expert state-conditioned action distribution, as done in BC. For example, the method proposed in [23], based on DAgger [10], queries supervisor actions in the states visited by the learner and adds the new data into the dataset. This adjusts the expert state distribution to match the learner’s state distribution, eliminating the requirement of an interactive expert. To get rid of the requirement of an interactive expert, methods like [24] based on generative adversarial imitation learning (GAIL) [12] utilize a discriminator to measure the difference between the learner’s and the expert’s state-action distribution and then compute the reward in reinforcement learning. By increasing the policy’s accumulated reward, the policy’s state distribution will get closer to the expert distribution. However, deploying such reinforcement loops in safety-critical tasks like autonomous driving is challenging since it requires online interaction with the environment. In contrast, our proposed method matches the ego future trajectory’s distribution conditioned on context, which can be learned through supervised learning without interaction with the environment or access to expert supervision. This avoids the need for online interaction with the environment and facilitates deployment in safety-critical scenarios such as autonomous driving.

- 2) *Offline Model-free IL*: The most popular model-free IL methods are based on DART [11] which avoids the compounding error by providing synthetic examples of how to recover from the deviated state. In [9], temporally correlated noise is injected into the trajectory to simulate gradual drift away from the desired trajectory. Alternatively, ChauffeurNet [18] adds a uniform perturbation to the SDV’s current pose, *i.e.* its rear axle’s midpoint coordinate and orientation, and fits a new smooth trajectory that brings the SDV back to the original target location. However, these rule-based trajectory augmentation methods are challenging to cover the real motion distribution induced by the learner’s policy, and the policy is likely to develop a tendency for perturbed driving. In our approach, we also apply perturbation to the SDV’s current position, but the perturbation’s role is to blur the ego position information, rather than to provide data augmentation. Thus, our method does not require a trajectory smoothing process during training. Our method is also model-free, but we aim to endow the policy with robustness properties by constraining its formulation without relying on recovery examples.

- 3) *Offline Model-based IL*: To address the distribution shift problem, model-based IL methods minimize the difference between a trajectory rolling out in a differentiable learned or data-driven model and the expert trajectory. For example, PPUU [25] learns a data-driven dynamics model based on a variational autoencoder [26] and trains the policy network to output actions that generate a similar trajectory as the expert trajectory. As the dynamics model is differentiable, an action

can receive gradients from multiple time steps ahead, enabling the penalization of actions that will result in large divergences in the future, even if their instantaneous divergences are small. UrbanDriver [20], instead of learning a model, constructs a differentiable data-driven model using recorded perception data and High Definition maps, where new observations are calculated by a coordinate transformation based on the ego vehicle's pose and collected data. However, the model-based approaches' performance is limited by their models' accuracy.

B. Imitation learning with robustness

IL is different from supervised learning by deploying the policy under dynamics, whose robustness is considered to be the learned policy's ability to recover from policy errors. Control-theoretic methods have been explored to learn policies with stability guarantees by constraining policy and system dynamics. Taylor Series IL [27] shows that the induced trajectories of a learner and expert will be close if their derivative difference at expert states is small. However, computing high-order derivatives of the expert policy is difficult without sufficient data. Others learn a robust linear [13, 14] or simple feed-forward neural network [15] control policy for a linear dynamical system by posing constraints on the policy. Even though the authors in East [16] extend the robustness guarantee to the polynomial system and policy, obtaining guarantees on close-loop stability for the nonlinear autonomous driving system remains a challenge. Recently, the CMILe method [28] learns nonlinear policies with the same safety guarantees as the expert but requires online expert access, similar to DAGger [10]. In contrast, our method constrains our policy in the formulation of receiving context state and producing ego future trajectory, whose stability is guaranteed under a mild input-to-state stability assumption on the environment's dynamics.

III. THEORETICAL ANALYSIS

We first introduce the notations and definitions used in this paper. For any vector $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|_p$ denotes its L^p norm, while $\|\mathbf{x}\|$ represents L^2 norm. For any square matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\|\mathbf{A}\|$ represents its induced L^2 norm and $\rho(\mathbf{A})$ represents its spectral radius (the maximum of the absolute values of its eigenvalues). For every induced matrix norm, we have $\rho(\mathbf{A}) \leq \|\mathbf{A}\|$.

Definition 1 (Comparison functions). *A function $\gamma : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is class \mathcal{K} if it is continuous, strictly increasing and satisfies $\gamma(0) = 0$. A function $\beta(x, t) : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is class \mathcal{KL} if it is continuous, $\beta(\cdot, t)$ is class \mathcal{K} for each t and $\beta(x, \cdot)$ is decreasing for each x .*

Definition 2 (Input-to-state stability (ISS) [29]). *A discrete-time system $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t)$, $\mathbf{x}_0 = \boldsymbol{\xi}$ is input-to-state stable if there exists a class \mathcal{KL} function β and a class \mathcal{K} function γ such that, for each bounded input \mathbf{u} and initial condition $\boldsymbol{\xi}$ and $t \in \mathbb{Z}_+$, it holds that:*

$$\|\mathbf{x}_t(\boldsymbol{\xi}, \mathbf{u})\| \leq \beta(\|\boldsymbol{\xi}\|, t) + \gamma(\|\mathbf{u}\|_\infty), \quad (1)$$

where $\|\mathbf{u}\|_\infty = \sup_{t \in \mathbb{Z}_+} (\mathbf{u}_t)$ is the input's sup norm. If a system satisfies the ISS property, the distance between any two trajectories must eventually be bounded by its input signal and independent of initial conditions.

We consider performing imitation learning in a nonlinear discrete-time system as follows:

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t), \quad (2)$$

where $\mathbf{x}_t \in \mathbb{R}^n$ is the system state at time $t \in \mathbb{N}$ and $\mathbf{u}_t \in \mathbb{R}^u$ is the control input. For an autonomous driving system, its state can be separated into two parts $\mathbf{x}_t = (\mathbf{e}_t, \mathbf{c}_t)$: the ego state $\mathbf{e}_t \in \mathbb{R}^m$ controlled by the control input \mathbf{u} and the context state $\mathbf{c}_t \in \mathbb{R}^{n-m}$ influenced by the ego state. By viewing the policy error as a disturbance to the system, we apply linearization to the nonlinear system at the expert's state \mathbf{x}^* to obtain the following linear subsystem:

$$\mathbf{c}_{t+1} = \mathbf{A}\mathbf{c}_t + \mathbf{B}\mathbf{e}_t, \quad (3)$$

with solution:

$$\mathbf{c}_{t+1} = \mathbf{A}^{t+1}\mathbf{c}_0 + \sum_{j=0}^t \mathbf{A}^{t-j}\mathbf{B}\mathbf{e}_j. \quad (4)$$

Here with abuse of notation, we use the same symbol \mathbf{c}_t , \mathbf{e}_t for the context deviation $\mathbf{c}_t - \mathbf{c}_t^*$ and the ego state deviation $\mathbf{e}_t - \mathbf{e}_t^*$.

Because the real-world traffic system without the ego state deviation $\mathbf{c}_{t+1} = \mathbf{A}\mathbf{c}_t$ is stable, we have $\rho(\mathbf{A}) < 1$. For such a Schur matrix, there are constants $c > 0$ and $0 \leq \sigma < 1$ such that $\|\mathbf{A}^t\| \leq c\sigma^t$ [30]. Besides, the context in traffic systems generally includes static map elements that are not influenced by ego vehicles and intelligent human drivers who can quickly recover from the other vehicle's perturbation. It is natural to assume that the subsystem in Eq. (3) satisfying ISS property with $\beta(r, t) < \sigma^{t-1}r$ and $\gamma(r) \leq \epsilon r$, where $\epsilon \geq 0$. Combining Eq. (1) and Eq. (4), we have:

$$\begin{aligned} \beta(r, t) &= c\sigma^t r < \sigma^{t-1}r, \\ \gamma(r) &= \sum_{j=0}^{\infty} c\sigma^j \|\mathbf{B}\| r = \frac{c\|\mathbf{B}\|}{1-\sigma} r \leq \epsilon r, \end{aligned} \quad (5)$$

which implies that $c < \frac{1}{\sigma}$, and $\|\mathbf{B}\| \leq \frac{\epsilon(1-\sigma)}{c}$.

Next, we study how the policy in our formulation can be stable under this ISS system. We learn a context state feedback control $\mathbf{e}_t = \mathbf{u}_t = f_\theta(\mathbf{c}_t)$ that directly maps the context state to the ego state. For simplicity, we consider a linear policy $\mathbf{u}_t = \mathbf{K}\mathbf{c}_t$. Then, the linear subsystem in Eq. (3) can be simplified as $\mathbf{c}_{t+1} = (\mathbf{A} + \mathbf{B}\mathbf{K})\mathbf{c}_t$ whose stable condition is $\rho(\mathbf{A} + \mathbf{B}\mathbf{K}) < 1$. So, we have when $\|\mathbf{K}\| < \frac{c(1-\sigma)}{\epsilon(1-\sigma)}$,

$$\begin{aligned} \rho(\mathbf{A} + \mathbf{B}\mathbf{K}) &\leq \|\mathbf{A} + \mathbf{B}\mathbf{K}\| \leq \|\mathbf{A}\| + \|\mathbf{B}\|\|\mathbf{K}\| \\ &\leq c\sigma + \frac{\epsilon(1-\sigma)}{c}\|\mathbf{K}\| < 1. \end{aligned} \quad (6)$$

This result shows that the stability of closed-loop autonomous driving can be guaranteed as long as the context system is stable enough and the norm of our policy is sufficiently small.

In practice, we will add a L^2 norm regularization to the training loss of our neural network policy to internalize this stability prior.

Finally, we show how even with a fixed context $c_t = 0$, the naive BC policy formulation, which takes both the ego state and the context state and outputs the displacement from the current state, can fail. We consider a linear policy: $u_t = K_{bc}e_t$ and update the state $e_{t+1} = e_t + u_t$. Then, the dynamics is $e_{t+1} = (I + K_{bc})e_t$. Because $\|I\| = 1$, we cannot obtain the stable condition $\rho(I + K_{bc}) < 1$ by limiting the $\|K_{bc}\|$ as our policy formulation. Thus, the system’s stability cannot be guaranteed.

This simple analysis supports our decision to disregard the SDV’s historical trajectory during the learning process to avoid the *covariate shift*. Its advantage will also be validated in our experiments, showing that our method can significantly decrease the off-road rate.

IV. METHOD

In Fig. 1, we present an overview of our context-conditioned imitation learning method. Our policy network comprises two encoders, namely a spatial encoder and a temporal encoder based on the Transformer [31]. During training, we optimize the network parameters to minimize the L^1 distance between its predicted trajectories and the ground truth trajectories. During the evaluation, we employ a linear-quadratic regulator (LQR) [32] controller to yield smooth planning based on the predicted trajectory from the policy network.

A. Policy Inputs

Our policy network takes observed contexts from H previous time steps with a time step interval of I as inputs. Instead of directly learning an end-to-end policy from raw sensor data, such as camera images or lidar point clouds, we utilize intermediate input representations to improve generalization and interpretability, similar to the approach in [18].

The road networks consist of crosswalks and lanes, each represented as a polyline sequence of vectors with various features like VectorNet [33]. Crosswalk vector features include their initial and terminal points’ positions and sequence orders, while lane vectors also contain additional features such as their traffic light state and lane width.

Regarding the traffic participants, we divide them into the ego vehicle and other agents. As previously mentioned, we ignore the features of the ego vehicle that can be influenced by the learned policy and only consider its mission goal position without considering its historical trajectory. For other agents, we input their type, sizes, centroids, and orientations of bounding boxes at multiple past time steps.

Overall, our approach uses intermediate vector representations and disregards the SDV’s historical trajectory to improve model performance.

B. Coordinate System

After obtaining contexts without the explicit ego state for H time steps, we translate each context into its own current

coordinate system, allowing for feature reuse during closed-loop evaluation. To avoid leaking implicit ego information that could be detrimental to closed-loop performance, we develop an ego-perturbed goal-oriented coordinate system. The origin of this coordinate system is anchored at the ego vehicle’s position plus a zero-mean Gaussian perturbation, and the x -axis points to its goal position. The perturbation blurs the ego vehicle’s position information and ensures local observability. The goal position eliminates the influence of the SDV’s heading, which can be replaced with any point unaffected by the learned policy.

It’s worth noting that we apply the perturbation independently at all past H time steps to blur or conceal the implicit ego position information from the learned policy, rather than to augment training data. This is because perturbed data is nearly impossible to occur during training and evaluation. In contrast, ChauffeurNet [18] jitters the SDV’s current pose, fits a smooth trajectory to the perturbed point and the original start and end points, and generates realistic new training examples.

C. Policy Network

The policy network is designed to predict the SDV’s future trajectory using only its context. Since predicting the SDV’s future trajectory without its historical trajectory is more challenging, we need to make full use of spatial-temporal interactions between the SDV and its surroundings. To achieve this, we use a spatial encoder to embed observations at each time step into an observation feature. We then use a temporal encoder to embed the spatial observation features and generate predictions. The Transformer serves as the foundation for the network, as it has shown exceptional performance in autonomous driving [20, 34].

Spatial Encoder: To capture spatial relations between vectors of the same map polyline, we use a local Transformer encoder as the vector-vector interaction encoder. Next, we aggregate the features of vectors belonging to the same polyline by max-pooling to obtain polyline-level features. Additionally, we utilize two multi-layer perceptions (MLPs) to obtain agent and goal features. Furthermore, to model high-order interactions between the goal, agents, and map elements, we employ a global transformer encoder, with the goal embedding being output as observation features o_t .

Temporal Encoder: To capture temporal information and interaction, a Transformer encoder with a causal self-attention mask is harnessed to embed the H historical observation features. Finally, each hidden state of the Transformer is decoded by a full-connected linear layer to generate a prediction of the SDV’s poses for its relative future T time steps.

D. Training Process

We leverage offline supervised learning like BC to train the policy network by minimizing the L^1 distance between the predictions and ground truth trajectories. This is because the L^1 metric is more correlated with driving performance compared to the commonly used mean square error [35]. To help the network converge and generalize, we additionally

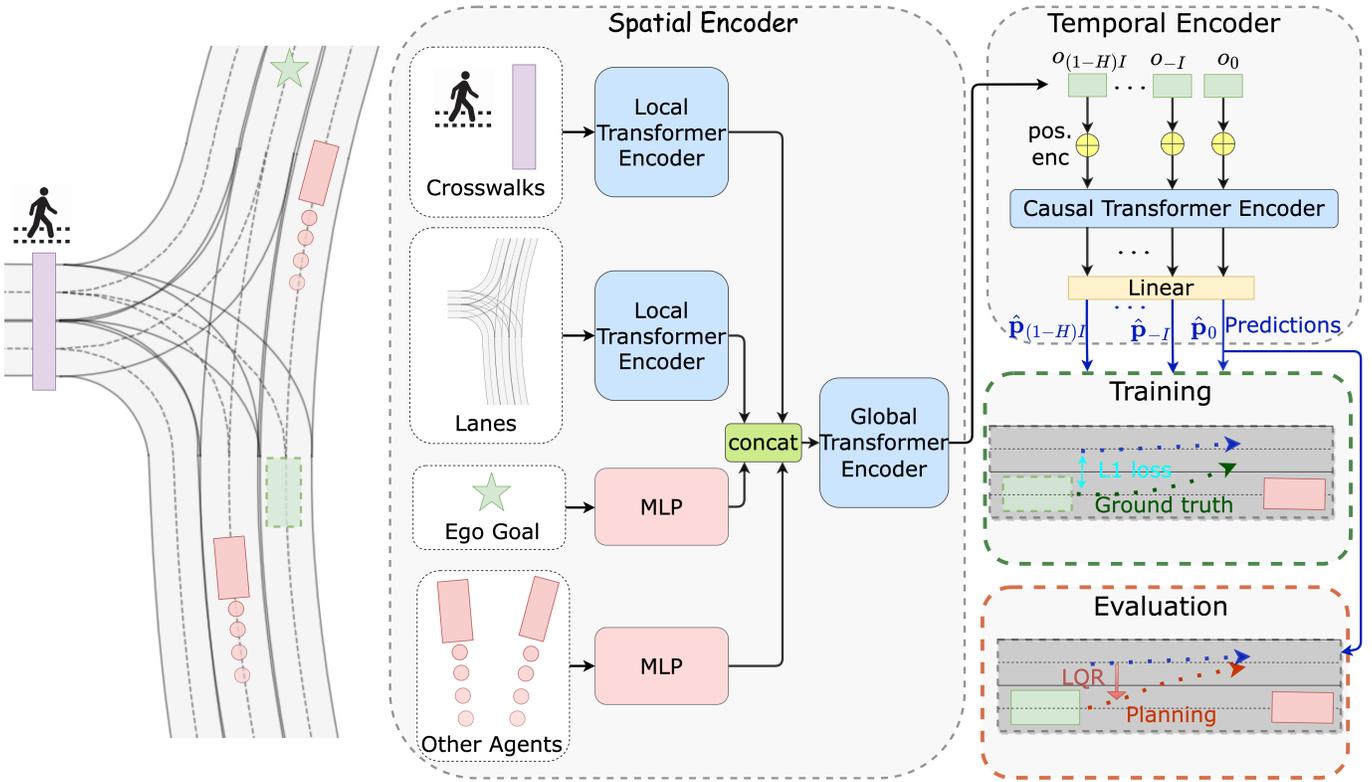


Fig. 1. Overview of our approach.

introduce an auxiliary task that minimizes the L^1 error at all previous time steps in the causal Transformer and apply a squared L^2 norm regularization to the network parameters θ , as inspired by Eq. (6). The final loss is:

$$L = \sum_{t=1}^T \left(\|\mathbf{p}_t - \hat{\mathbf{p}}_0^t\|_1 + \mu \sum_{h=1}^{H-1} \|\mathbf{p}_{t-hI} - \hat{\mathbf{p}}_{hI}^t\|_1 \right) + \frac{\lambda}{2} \|\theta\|^2, \quad (7)$$

where \mathbf{p}_t denotes the ground truth ego pose at time step t and $\hat{\mathbf{p}}_{hI}^t$ is the pose prediction at the past hI time step for its relative future t time step. μ is an auxiliary task hyperparameter, and λ is a regularization factor.

E. Evaluation Process

While the prediction without ego state input is more stable, it struggles to ensure the smoothness of the predicted trajectory from the current state. To address this, prior works [36, 25] have typically added a differentiable kinematic layer into the policy network to generate physically feasible planning. However, this differentiable kinematic layer would incorporate the SDV's current information, which is undesirable in learning a policy. Instead, we choose to obtain a smooth trajectory during the evaluation process by applying the LQR.

LQR is a computationally efficient method that minimizes the total commutative quadratic cost of a linear dynamic system. For simplicity, we consider a finite-horizon, discrete-

time linear system with dynamics described by:

$$\begin{bmatrix} \mathbf{p}_{t+1} \\ \dot{\mathbf{p}}_{t+1} \\ \ddot{\mathbf{p}}_{t+1} \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{D} & \mathbf{D}^2 \\ 0 & \mathbf{I} & \mathbf{D} \\ 0 & 0 & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{p}_t \\ \dot{\mathbf{p}}_t \\ \ddot{\mathbf{p}}_t \end{bmatrix} + \begin{bmatrix} \mathbf{D}^3 \\ \mathbf{D}^2 \\ \mathbf{D} \end{bmatrix} \mathbf{u}_t, \quad (8)$$

where \mathbf{D} is a diagonal matrix with the interval of each time step as diagonal entries, $\dot{\mathbf{p}} = (\omega_t, \mathbf{v}_t)$, $\ddot{\mathbf{p}}_t = (\alpha_t, \mathbf{a}_t)$, $\mathbf{u}_t = (\zeta_t, \mathbf{j}_t)$ represent the angular and positional velocity and acceleration, jerk, and control input, respectively. The system is subject to a quadratic cost function:

$$J = \sum_{t=1}^T \|\mathbf{p}_t - \hat{\mathbf{p}}_0^t\|^2 + \eta_\omega \omega_t^2 + \eta_\alpha \alpha_t^2 + \eta_a \|\mathbf{a}_t\|^2 + \eta_j \|\mathbf{j}_t\|^2, \quad (9)$$

where \mathbf{p}_t represents the planned pose, and the predicted pose $\hat{\mathbf{p}}_0^t$ from the policy network are regarded as target pose. We use weights η_ω , η_α , η_a and η_j to balance safety and smoothness of the planned trajectory. At the end of the optimization, a smooth trajectory with positions and headings can be generated for the SDV to follow.

V. EXPERIMENTS

A. Dataset

To evaluate the performance of our method, we conduct experiments on two large-scale real-world datasets:

Lyft Level 5 Prediction Dataset [21]: It contains approximately 1,000 h of urban driving demonstrations from Palo Alto, which have been separated into independent scenes of nearly 25 s at a frequency of 10 Hz. We train our network on

the provided 100 h subset (16,265 scenes) as UrbanDriver [20] and test it with all 16,220 validation scenes.

nuPlan Dataset [22]: This dataset consists of 1,312 h of human driving data collected in four cities (Boston, Pittsburgh, Las Vegas, and Singapore). Due to significant differences in traffic rules and patterns between different cities, we extract driving data in Las Vegas as Phan-Minh et al. [37]. Then, we segment the data into independent scenes of 25 s at a frequency of 10 Hz, as in the Lyft dataset, and filter out scenes without a mission goal. After filtering, we obtain 63,181 training, 4,774 validation, and 6,386 testing scenes.

More information on data preparation and another **toy Dataset** can be found in the appendix.

B. Closed-Loop Evaluation

To evaluate the closed-loop performance of our method, we use a log-replay simulator, as in prior work [20, 38]. At each step in the log-replay simulator, the SDV updates its pose according to the planned trajectory, while other agents are assumed to follow their recorded trajectories in the dataset. For both datasets, we evaluate our method for 25 s at a frequency of 10 Hz, using the following metrics:

Collision Rate: If the SDV collides with other agents at any time step during a scene, that scene is considered a collision scene. The collision rate is calculated by taking the ratio of the number of collision scenes to the total number of scenes.

Off-road Rate: In nuPlan, we use the official off-road metric, which considers the distance between a corner of the SDV’s bounding box and the drivable area. Specifically, if this distance exceeds 0.3 m, the scene is deemed off-road. However, in the Lyft dataset, there is no access to the drivable area. Thus, following UrbanDriver [20], we consider a scene off-road if the SDV deviates laterally from the human driver’s ground truth by more than 2 m in the scene.

Discomfort: To quantify the comfort and feasibility of the planned trajectory, we calculate the rate at which the absolute acceleration values exceed 3 m/s^2 across all time steps.

L2: We use the average L2 position errors between the roll-out trajectory and the human driver’s ground truth to quantify the similarity between our results and human driving.

C. Performance

On the Lyft dataset, we compare our methods against three state-of-the-art methods to demonstrate the advantage of the proposed framework:

Raster-perturb: An official baseline learned by BC for the Lyft dataset [21], based on ResNet-50 [39], which receives a Bird-Eye-View (BEV) representation of the scene surrounding the SDV and plans a trajectory with position and yaw displacements. To augment the data, a perturbation is applied to the current SDV position, and then a new kinematically feasible trajectory to reach the original endpoint is generated as ChauffeurNet.

BC-perturb: A BC method provided by UrbanDriver with the same trajectory perturbation and output as **Raster-perturb**

but its map and agent inputs are represented in vector formulations, which are processed by a VectorNet and Transformer. In addition, the SDV’s history is equipped with a dropout.

UrbanDriver: An offline policy gradient method to imitate the expert’s policy by exploiting a differentiable data-driven simulator using the same input and model structure as **BC-perturb**.

On the nuPlan dataset, we compare our methods against an **official** baselines learned by BC due to a lack of other prior works:

LaneGCN-perturb: A vector-based model that utilizes a series of MLPs to encode SDV and agent inputs and a LaneGCN [40] to encode vector-map elements, and then a fusion network to capture lane and agent intra- and inter-interactions through attention layers. To augment the training data, the SDV trajectories are perturbed and other agents are randomly dropped out.

In addition to these models provided by prior works, we also learn our model using a representative method in behavior cloning and offline reinforcement learning on both datasets:

TD3+BC [41]: It adds a BC term to the policy updating of the Twin Delayed Deep Deterministic Policy Gradient (TD3) [42] for implicit policy constraints. We construct the offline RL dataset by applying the trajectory perturbation augmentation and consider collision, and comfort for the reward design.

Vector-Chauffeur: Our network learned using the same data augmentation method as the ChauffeurNet including trajectory perturbation and ego past motion dropout. And we represent the data in the same coordinate system, which uses the SDV’s location as the origin and its heading perturbed by a uniform noise as the orientation.

Our proposed method outperformed all previous work significantly on both datasets, as shown in Tab. I. The nuPlan dataset had a higher collision rate than the Lyft dataset due to the scenarios being more complex with more road agents. The poor performance of the official baseline **LaneGCN-perturb** on the nuPlan dataset is due to using the official hyper-parameters without fine-tuning, as its training process is time-consuming.

D. Ablation Study

The following ablation experiments on the Lyft dataset are used to expose the significance of different components of our model, whose results are shown Tab. II:

Network Input: We first study the importance of removing explicit ego information from the network input. We consider two approaches to reintroducing explicit ego inputs. One is to directly input the SDV’s past positions and then process them in the same way that the goal position is processed. The other one is to additionally introduce a dropout of 50% at the ego input during training as ChauffeurNet. Note that our method is different from dropout methods because the dropout is only applied during training while we keep removing the ego state. We can observe that in both explicit ego information input

TABLE I
COMPARISON WITH BASELINES OF CLOSED-LOOP PERFORMANCE ON THE LYFT AND NUPLAN DATASETS.

Model	Num params	Collision(%)	Off-road(%)	Discomfort(%)	L2(m)
Raster-perturb*	23.6M	15.48	5.06	4.00	5.90
BC-perturb*	1.8M	9.38	6.77	39.10	4.77
UrbanDriver*	1.8M	13.28	7.27	39.41	5.74
TD3+BC	2.8M	22.53±1.76	15.21±0.97	4.86±0.47	6.34±0.41
Vector-Chauffeur	1.5M	10.12±0.23	3.40±0.32	5.42±0.44	5.03±0.43
CCIL (ours)	1.5M	3.32±0.15	0.62±0.13	4.33±0.22	1.23 ± 0.08
LaneGCN-perturb	2.0M	60.63±2.34	34.25±1.65	17.26±1.80	21.21±1.81
TD3+BC	2.8M	39.12±2.21	18.59±1.04	10.56±0.95	15.04±1.62
Vector-Chauffeur	1.5M	24.12±1.37	10.11±0.62	12.53±1.17	6.12±0.87
CCIL (ours)	1.5M	6.91±0.11	3.08±0.11	1.16±0.05	3.68±0.04

*There is no variance in Raster-perturb, BC-perturb, and UrbanDriver because we evaluate the deterministic pre-trained models in a deterministic simulator.

TABLE II
ABLATION EXPERIMENTS ON CLOSED-LOOP PERFORMANCE ON THE LYFT DATASET

Model	Perturb	Ego	Collision(%)	Off-road(%)	Discomfort(%)	L2(m)
w explicit ego	✓	✓	20.29±0.88	19.18±2.98	0.57±0.15	5.50±0.55
w ego dropout	✓	✓	14.05±1.53	5.02±0.88	0.63±0.20	4.15±0.47
w ego coordinate		✓	11.31±1.44	9.79±1.34	0.95±0.05	3.87±0.11
std=0			7.08±0.35	2.86±0.25	0.89±0.10	3.46±0.34
std=1	✓		3.39±0.17	1.00±0.16	1.99±0.15	2.10±0.05
std=2 (ours)	✓		3.32±0.15	0.62±0.13	4.33±0.22	1.23±0.08
std=3	✓		3.42±0.12	0.49±0.10	7.35±0.26	0.91±0.04
w/o causal Trans	✓		4.28±0.25	1.43±0.25	6.53±0.26	1.63±0.10
w/o LQR	✓		3.81±0.14	2.07±0.12	89.05±0.35	1.02±0.02
w/o regularization	✓		4.07±0.16	1.05±0.14	4.96±0.30	1.23±0.09
w/o auxiliary	✓		4.56±0.29	1.02±0.06	3.23±0.31	1.92±0.07

ways, there is a steep drop in the collision and off-road rate and L2 distance due to the *covariate shift* issue.

Coordinate System: To analyze the impact of the implicit ego information in the coordinate system, we first consider replacing our ego-perturbed goal-oriented coordinate system with the ego-centric coordinate system in ChauffeurNet using orientation uniformly around the heading. We observe that the ChauffeurNet coordinate system leads to inferior closed-loop performance due to the implicit ego information. Furthermore, in our coordinate system, we add a Gaussian noise with a zero mean to the SDV’s current position to obtain the origin. By increasing the standard deviation (std) of the Gaussian noise, we can reduce the implicit ego information. We observe that as the std increased, discomfort also increased, while the off-road rate and closed-loop L2 decreased. This suggests that although implicit ego information can improve instantaneous planning accuracy, it may worsen long-term closed-loop driving performance. We guess this is because during closed-loop driving, an auto-regressive model can generate temporal correlations in planning errors, leading to consistent bias in one direction that results in off-road events. However, temporally uncorrelated noise can cause slight oscillations around the expert trajectory but still enable successful driving. The phenomenon that offline instantaneous planning accuracy and actual long-term driving quality are weakly correlated was also noted by Codevilla et al. [35].

Model Architecture: To demonstrate the importance of the temporal information and the effectiveness of the causal Transformer in capturing the temporal interactions, we replace it with a MLP. Our findings suggest that the causal Transformer achieves better performance by integrating both spatial and temporal information. Moreover, we conducted an ablation study on the LQR module to examine its impact on comfort. We observed a significant decrease in comfort when the LQR mechanism was removed, as well as a small increase in collision and off-road rates. However, there was a slight improvement in the L2 metric. This highlights the importance of the LQR module in enhancing trajectory smoothness while acknowledging that the ego information introduced during evaluation can still impair closed-loop performance on the L2 metric.

Loss Term: To demonstrate the effectiveness of the additional loss terms, we remove the auxiliary and regularization loss term respectively. We find that both are beneficial to improving performance, although the improvement brought by norm regularization is limited, likely because our policy network has a modest parameter number and is well-regularized even without the loss term.

E. Qualitative Analysis

We visualize and compare the closed-loop trajectories of our method and the baselines in two Lyft scenarios. As shown in Fig. 2, our method generates a feasible trajectory

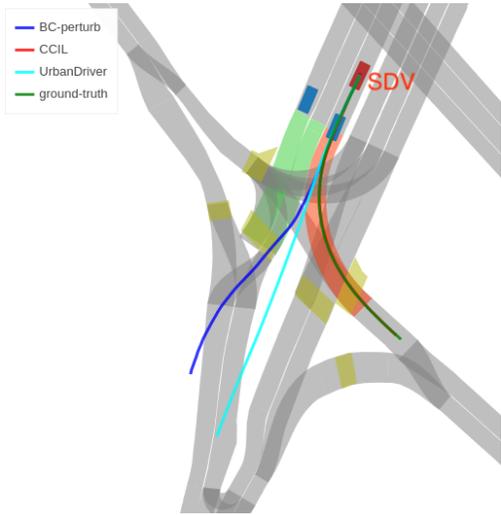


Fig. 2. Qualitative comparison between our method and the baselines for on-road driving.

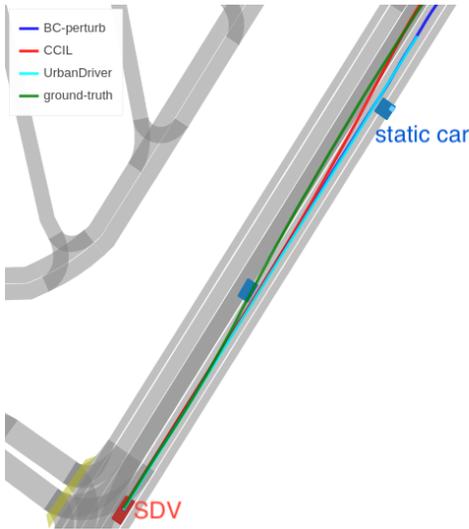


Fig. 3. Qualitative comparison between our method and the baselines for collision avoidance.

similar to that of a human driver, whereas the BC-perturb and UrbanDriver methods drive the SDV offroad. We can see in Fig. 3 that our method can avoid the collision with the static car at the shoulder by changing lanes while the baseline methods keep moving straight, leading to a crash. We hypothesize that off-road or collision events occur when the policy outputs an erroneous action at the turning point or lane-changing point, such as continuing to move straight forward. In such cases, the policy taken in out-of-distribution history may make a larger error, leading to going off-road or colliding with other objects. The trajectory perturbation in BC-perturb and the trajectory distribution matching in UrbanDriver do not help the SDV recover from the erroneous state, which may be caused by the strong correlation between past straight trajectories and moving straight action in the training data.

We also conducted a qualitative analysis of our model

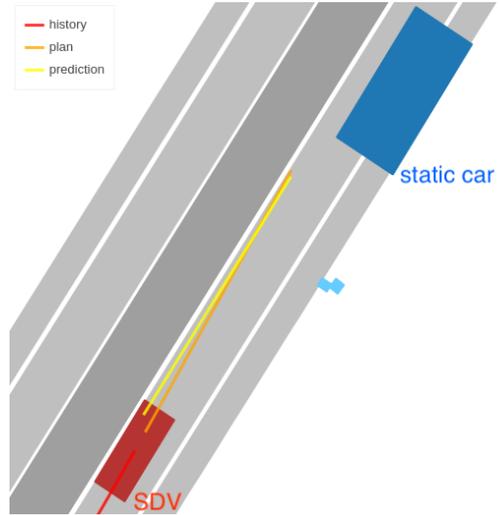


Fig. 4. Qualitative analysis of our model's prediction and planning performance. Please refer to Fig. 3 for the whole scenario.

prediction and planning in the collision avoidance scenario of Fig. 3. In Fig. 4, we demonstrate that our model can plan a smooth trajectory based on the policy network prediction. When faced with a static car ahead on the shoulder, our policy network infers that the SDV's future positions should be closer to the road center. Although the prediction is not dynamically feasible from the SDV's current state, it is close enough to it. This suggests that our policy can infer a reasonable ego state from the context.

VI. CONCLUSION

We have proposed a new offline imitation learning method to mitigate the distribution shift of behavior cloning. Our approach involves training a neural network to predict the future poses of the ego vehicle, without receiving explicit ego information. To remove implicit ego information and ensure local observability, we introduce a new ego-perturbed goal-oriented coordinate system for representation of the observation input. To tackle the challenging learning task, we design a Transformer-based network that leverages historical spatial-temporal context data effectively. Finally, we demonstrate the effectiveness of our approach using two real-world large-scale datasets, achieving state-of-the-art performance.

ACKNOWLEDGMENTS

This work was supported by Alibaba Group through Alibaba Innovative Research Program and Alibaba Research Intern Program.

REFERENCES

- [1] Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. A survey of autonomous driving: Common practices and emerging technologies. *IEEE access*, 8:58443–58469, 2020.

- [2] Mikael Henaff, Alfredo Canziani, and Yann LeCun. Model-predictive policy learning with uncertainty regularization for driving in dense traffic. In *ICLR*, 2019.
- [3] 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 arXiv:1807.08048*, 2018.
- [4] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [5] Yunpeng Pan, Ching-An Cheng, Kamil Saigol, Keuntaek Lee, Xinyan Yan, Evangelos A Theodorou, and Byron Boots. Imitation learning for agile autonomous driving. *The International Journal of Robotics Research*, 39(2-3): 286–302, 2020.
- [6] Lingping Gao, Ziqing Gu, Cong Qiu, Lanxin Lei, Shengbo Eben Li, Sifa Zheng, Wei Jing, and Junbo Chen. Cola-HRL: Continuous-Lattice Hierarchical Reinforcement Learning for Autonomous Driving. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 13143–13150, 2022.
- [7] Donald Michie, Michael Bain, and Jean Hayes-Michie. Cognitive models from subcognitive skills. *IEEE control engineering series*, 44:71–99, 1990.
- [8] Dean A Pomerleau. ALVINN: An autonomous land vehicle in a neural network. *NeurIPS*, 1, 1988.
- [9] Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *ICRA*, pages 4693–4700, 2018.
- [10] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *AISTATS*, pages 627–635, 2011.
- [11] Michael Laskey, Jonathan Lee, Roy Fox, Anca Dragan, and Ken Goldberg. DART: Noise injection for robust imitation learning. In *CoRL*, pages 143–156, 2017.
- [12] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *NeurIPS*, 29, 2016.
- [13] Malayandi Palan, Shane Barratt, Alex McCauley, Dorsa Sadigh, Vikas Sindhwani, and Stephen Boyd. Fitting a linear control policy to demonstrations with a Kalman constraint. In *Learning for Dynamics and Control*, pages 374–383, 2020.
- [14] Aaron Havens and Bin Hu. On imitation learning of linear control policies: Enforcing stability and robustness constraints via LMI conditions. In *ACC*, pages 882–887, 2021.
- [15] He Yin, Peter Seiler, Ming Jin, and Murat Arcak. Imitation learning with stability and safety guarantees. *IEEE Control Systems Letters*, 6:409–414, 2021.
- [16] Sebastian East. Imitation learning of stabilizing policies for nonlinear systems. *European Journal of Control*, page 100678, 2022.
- [17] Felipe Codevilla, Eder Santana, Antonio M López, and Adrien Gaidon. Exploring the limitations of behavior cloning for autonomous driving. In *CVPR*, pages 9329–9338, 2019.
- [18] Mayank Bansal, Alex Krizhevsky, and Abhijit Ogale. ChauffeurNet: Learning to drive by imitating the best and synthesizing the worst. In *RSS*, 2019.
- [19] Pim De Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. *NeurIPS*, 32, 2019.
- [20] Oliver Scheel, Luca Bergamini, Maciej Wolczyk, Błażej Osiński, and Peter Ondruska. Urban Driver: Learning to drive from real-world demonstrations using policy gradients. In *CoRL*, pages 718–728, 2022.
- [21] John Houston, Guido Zuidhof, Luca Bergamini, Yawei Ye, Long Chen, Ashesh Jain, Sammy Omari, Vladimir Iglovikov, and Peter Ondruska. One thousand and one hours: Self-driving motion prediction dataset. In *CoRL*, 2020.
- [22] Holger Caesar, Juraj Kabzan, Kok Seang Tan, Whye Kit Fong, Eric Wolff, Alex Lang, Luke Fletcher, Oscar Beijbom, and Sammy Omari. NuPlan: A closed-loop ML-based planning benchmark for autonomous vehicles. In *CVPR ADP3 workshop*, 2021.
- [23] Jiakai Zhang and Kyunghyun Cho. Query-Efficient Imitation Learning for End-to-End Simulated Driving. In *AAAI*, page 2891–2897, 2017.
- [24] Pin Wang, Dapeng Liu, Jiayu Chen, Hanhan Li, and Ching-Yao Chan. Decision making for autonomous driving via augmented adversarial inverse reinforcement learning. In *ICRA*, pages 1036–1042. IEEE, 2021.
- [25] Brandon Amos, Ivan Jimenez, Jacob Sacks, Byron Boots, and J Zico Kolter. Differentiable mpc for end-to-end planning and control. *NeurIPS*, 31, 2018.
- [26] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [27] Daniel Pfrommer, Thomas TCK Zhang, Stephen Tu, and Nikolai Matni. TaSIL: Taylor Series Imitation Learning. *arXiv preprint arXiv:2205.14812*, 2022.
- [28] Stephen Tu, Alexander Robey, Tingnan Zhang, and Nikolai Matni. On the sample complexity of stability constrained imitation learning. In *Learning for Dynamics and Control Conference*, pages 180–191, 2022.
- [29] Eduardo D Sontag. Input to state stability: Basic concepts and results. In *Nonlinear and optimal control theory*, pages 163–220. Springer, 2008.
- [30] Zhong-Ping Jiang and Yuan Wang. Input-to-state stability for discrete-time nonlinear systems. *Automatica*, 37(6): 857–869, 2001.
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017.
- [32] Karl Johan Åström and Richard M Murray. *Feedback systems: an introduction for scientists and engineers*.

Princeton University Press, 2021.

- [33] Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, and Cordelia Schmid. VectorNet: Encoding hd maps and agent dynamics from vectorized representation. In *CVPR*, pages 11525–11533, 2020.
- [34] Jiquan Ngiam, Vijay Vasudevan, Benjamin Caine, Zhengdong Zhang, Hao-Tien Lewis Chiang, Jeffrey Ling, Rebecca Roelofs, Alex Bewley, Chenxi Liu, Ashish Venugopal, et al. Scene transformer: A unified architecture for predicting future trajectories of multiple agents. In *International Conference on Learning Representations*, 2022.
- [35] Felipe Codevilla, Antonio M Lopez, Vladlen Koltun, and Alexey Dosovitskiy. On offline evaluation of vision-based driving models. In *ECCV*, pages 236–251, 2018.
- [36] Matt Vitelli, Yan Chang, Yawei Ye, Ana Ferreira, Maciej Wołczyk, Błażej Osiński, Moritz Niendorf, Hugo Grimmer, Qiangui Huang, Ashesh Jain, et al. SafetyNet: Safe planning for real-world self-driving vehicles using machine-learned policies. In *ICRA*, pages 897–904, 2022.
- [37] Tung Phan-Minh, Forbes Howington, Ting-Sheng Chu, Sang Uk Lee, Momchil S Tomov, Nanxiang Li, Caglayan Dicle, Samuel Fidler, Francisco Suarez-Ruiz, Robert Beaudoin, et al. Driving in Real Life with Inverse Reinforcement Learning. *arXiv preprint arXiv:2206.03004*, 2022.
- [38] Zhiyu Huang, Haochen Liu, Jingda Wu, and Chen Lv. Differentiable Integrated Motion Prediction and Planning with Learnable Cost Function for Autonomous Driving. *arXiv preprint arXiv:2207.10422*, 2022.
- [39] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [40] Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. In *ECCV*, pages 541–556, 2020.
- [41] Scott Fujimoto and Shixiang Shane Gu. A minimalist approach to offline reinforcement learning. *Advances in neural information processing systems*, 34:20132–20145, 2021.
- [42] Scott Fujimoto, Herke van Hoof, and David Meger. Addressing Function Approximation Error in Actor-Critic Methods. In *ICML*, volume 80, pages 1587–1596, 2018.

A. Map preprocessing

In Lyft, there are only two types of map elements: lane and crosswalk. In nuPlan, the map elements consist of a lane, lane connector, intersection, stop line, crosswalk, walkway, and car park. We categorize them into polyline elements including lanes and lane connectors, and polygon elements including stop lines, crosswalks, intersections, walkways, and car parks.

For each polyline element, it is approximated by a sequence of vectors with an interval of 3 m. For both datasets, we preprocess the polyline map into a graph to take advantage of the topological information. We connect each vector with its nearest left, right, and next vector if any exist. The nearest left or right vector is the nearest vector of its reachable left or right polyline. And the distance between vectors is represented by the Euclidean distance between their starting points. After building the vector graph, we can compute the travel distance between any two vectors using Dijkstra’s algorithm. For both datasets, the polyline vector has shared features, including coordinates of the starting and ending points, distances to the left and right vectors, distances to the mission goal, traffic light states (red or green), and the sequence order. However, lane vectors in Lyft have a lane width feature because the vectors approximate the computed middle lines of lanes, while the lane or lane connector vectors in nuPlan have a lane left and width feature because the vectors approximate an annotated baseline. In addition, the polyline vectors in nuPlan have an additional speed limit and type feature.

For each polygon element, we also approximate it with a sequence of vectors. In the Lyft dataset, the crosswalk vectors are directly constructed by connecting the original sequential annotation coordinates. But we approximate each polygon with a fixed set of 20 vectors because the annotation point number is too large for elements like intersections. For both datasets, each polygon vector has features including coordinates of starting and ending points and its order sequence, while vectors in nuPlan have additional type features.

The inputs to our neural network are composed of two types of map elements: polylines and polygons, and two types of agent elements: other agents and an ego goal. The missing inputs are padded with zeros and masked out when calculating the attention. The origin is the perturbed SDV position.

Polyline: 30 topologically nearest polylines with vectors whose starting points are within 35 m from the origin. The topological distance between the origin and a polyline is the minimum of the distances between the origin and its vectors.

Polygons: 20 polygons whose boundaries are within 35 m from the origin. If there are more than 20 polygons, they are selected according to the importance of their types: stop line, crosswalk, intersection, walkway, and car park.

Agents: the 30 nearest agents whose oriented boxes’ centroids are within 50 m from the origin. The agent features in Lyft include its centroid coordinates, yaws, shapes, types, and relative times in the past 2 and current steps. The nuPlan agents additionally have velocity features, as they are provided.

TABLE III
HYPER-PARAMETERS FOR BOTH DATASETS

Hyper-parameter	Value
Future steps T	15
All Transformers dropout	0.1
All Transformers head number	8
All Transformers hidden size	128
Local Transformer layer number	3
Global Transformer layer number	6
Causal Transformer layer number	3
Causal Transformer history length H	15
Causal Transformer interval I	2
Auxiliary weight μ	0.3
Regularization weight λ	0.0001
LQR angular velocity weight η_ω	0.1
LQR angular acceleration weight η_α	0.1
LQR positional acceleration weight η_a	0.1
LQR positional jerk weight η_j	0.01

Goal: the x, y coordinates of the SDV’s mission goal of a scene. In Lyft, the mission goal is not provided, so we regard the ending point of the lane where the SDV locates at the last time step of the scene as the mission goal. In nuPlan, we directly take the provided mission goal at the last time step of the scene as the scene mission goal.

B. Model

For both datasets, the same model architecture is used, whose hyper-parameters are listed in Tab. III.

C. Training

Our model is trained using the Adam optimizer with a learning rate of 0.0005, 10000 steps linear warm-up, $\beta = (0.9, 0.999)$, and batch size 128. We stop training after 30 epochs and select the model with the smallest validation collision rate for evaluation. We train all models except the pre-trained model in the Lyft dataset independently 3 times and then report the mean and std of their performances.

D. Evaluation

For Lyft baselines, we directly evaluate the pre-trained model provided by the Lyft dataset and UrbanDriver. The Raster-perturb is the **model trained on train.zarr for 2 epochs** at <https://github.com/woven-planet/15kit/blob/master/examples/planning/train.ipynb>. The BC-perturb and UrbanDriver are the **Open Loop with history** and the **Urban Driver** at https://github.com/woven-planet/15kit/blob/master/examples/urban_driver/train.ipynb.

We evaluate all Lyft models from the second time step to compute current velocity using position information as input to the LQR with zero assumed initial acceleration.

For nuPlan baselines, we train the official models by ourselves using provided hyper-parameters and evaluate from the first time step because the velocity and acceleration information is provided.

E. Runtime

We conduct runtime experiments using a single Nvidia GeForce GTX 1080 GPU and an Intel i7-8700@3.2GHz CPU on the Lyft dataset. We measure the runtime of each method its mean and std over all time steps in an evaluation. The runtime results shown in Tab. IV consider all components in each model including data-preprocessing, model inference, and control. We observe that our architecture can achieve higher data processing efficiency and medium model inference efficiency compared with other methods. Our method takes a longer total execution time due to the extra LQR control module which does not exist in prior works because they only focus on optimizing positional accuracy but not comfort. However, our approach can still be executed in real-time on this hardware.

F. Toy Experiment

We design a toy experiment to vividly show our method’s ability to reduce compounding error. In the experiment, we use synthetic data from a very ideal and simplified scenario where a SDV moves under a ring road network with a fixed speed of 1 m/s at 1 Hz, as shown in Fig. 5. During training, the radius of the ring is a variable with a range from 10 m to 100 m and the circular lane is represented as a series of fixed lane points with the same interval of nearly 1m.

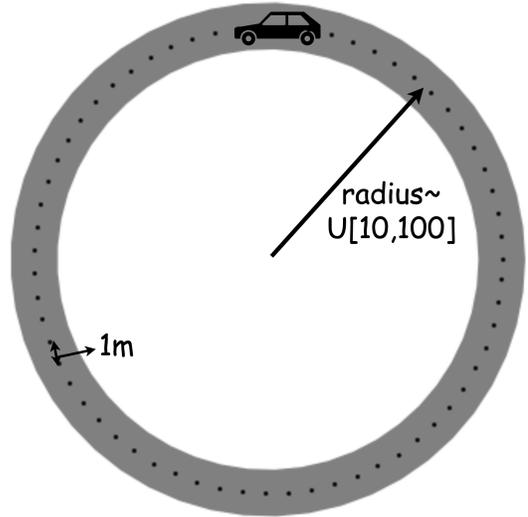


Fig. 5. Ring road

We compare our CCIL methods with several baseline methods introduced above, including BC, BC-perturb, and UrbanDriver. The inputs of these baseline methods are composed of two parts: ego state (its position and orientation in the past 10 time steps) and context (the nearest 10 lane points) in the ego coordinate system, while our method only takes the context in the past 10 time steps as inputs in the ego-perturbed center-oriented coordinate system. The ego-perturbed origin-oriented coordinate system means using the SDV position plus a zero-mean one-std Gaussian perturbation as the origin and

TABLE IV
AVERAGED RUNTIME PER FRAME OF INDIVIDUAL COMPONENTS FOR EACH METHOD

Model	Data process (ms)	Model inference (ms)	Control (ms)	Total (ms)
Raster-perturb	6.03±0.61	4.62±0.16	-	10.65±0.64
BC-perturb	6.69±0.72	6.78±5.26	-	13.47±5.65
UrbanDriver	6.33±1.41	12.78±8.80	-	19.11±9.23
CCIL	4.92±0.63	6.74±0.34	11.09±0.27	22.75±1.07

orientation to the center of the circular road as the x -axis direction. The trajectory perturbation is applied to augment the data in the BC-perturb method. For outputs, the BC and CCIL methods generate the relative position and yaw at the next time step and the BC-perturb method produces the next 10 time steps. In the UrbanDriver method, we unroll the policy for 32 time steps.

In each method, we employ a two-layer MLP with a hidden size of 128 as a policy network. We train the neural networks using the Adam optimizer with a learning rate of 0.0001 and random initial weights 100 times. We stop training after 10000 steps and then unroll the policy from one random starting point on the circle of radius 50 m for 100 time steps. The 100 closed-loop trajectories for each method are depicted in Fig. 6. We can observe that some trajectories in the baseline methods deviate from the road due to the covariate shift issue, but the trajectories in our method keep following the route.

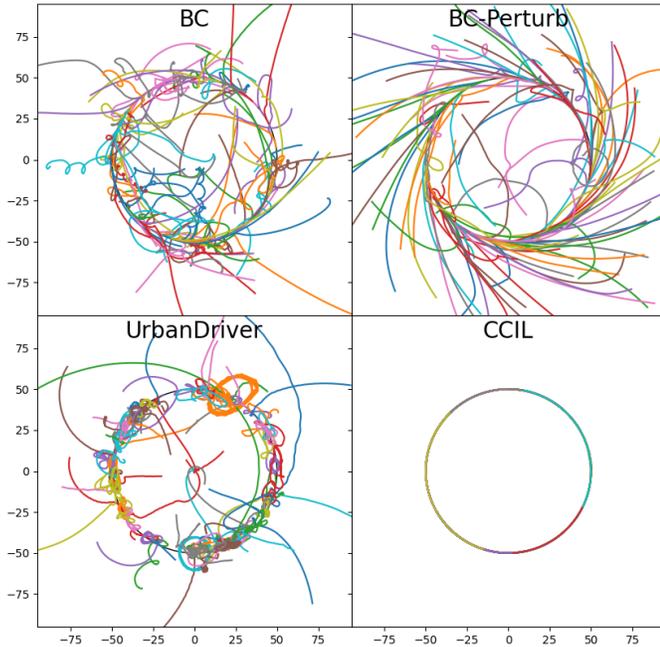


Fig. 6. Closed-loop trajectories of each model trained on the toy dataset.