

# ParkDiffusion++: Ego Intention Conditioned Joint Multi-Agent Trajectory Prediction for Automated Parking using Diffusion Models

Jiarong Wei<sup>1,2</sup>, Anna Rehr<sup>2</sup>, Christian Feist<sup>2</sup>, and Abhinav Valada<sup>1</sup>

**Abstract**—Automated parking is a challenging operational domain for advanced driver assistance systems, requiring robust scene understanding and interaction reasoning. The key challenge is twofold: (i) predict multiple plausible ego intentions according to context and (ii) for each intention, predict the joint responses of surrounding agents, enabling effective what-if decision-making. However, existing methods often fall short, typically treating these interdependent problems in isolation. We propose ParkDiffusion++, which jointly learns a multi-modal ego intention predictor and an ego conditioned multi-agent joint trajectory predictor for automated parking. Our approach makes several key contributions. First, we introduce an ego intention tokenizer that predicts a small set of discrete endpoint intentions from agent histories and vectorized map polylines. Second, we perform ego intention conditioned joint prediction, yielding socially consistent predictions of the surrounding agents for each possible ego intention. Third, we employ a lightweight safety-guided denoiser with different constraints to refine joint scenes during training, thus improving accuracy and safety. Fourth, we propose counterfactual knowledge distillation, where an EMA teacher refined by a frozen safety-guided denoiser provides pseudo-targets that capture how agents react to alternative ego intentions. Extensive evaluations demonstrate that ParkDiffusion++ achieves state-of-the-art performance on the Dragon Lake Parking (DLP) dataset and the Intersections Drone (inD) dataset. Importantly, qualitative what-if visualizations show that other agents react appropriately to different ego intentions.

## I. INTRODUCTION

Automated parking systems are essential for managing shared spaces between human drivers and automated vehicles. A typical automated parking pipeline consists of three stages: searching, decision making, and maneuvering. The core challenge of decision making lies in coupling ego vehicle intention prediction with multi-agent joint what-if trajectory prediction conditioned on ego intention. Many existing trajectory prediction approaches operate in an open loop [1]–[3], where future trajectories are predicted from past agent motions. However, they rarely model how surrounding agents would respond conditioned on counterfactual ego intentions. This shortcoming is critical in parking lots, where weak lane semantics and mixed traffic induce diverse interactions including following a leading vehicle, yielding to oncoming traffic, and avoiding pedestrians. In addition, existing parking prediction works [4]–[6] primarily yield marginal predictions without explicit scene-level compatibility or safety constraints.

The task of combining ego intention prediction with ego intention conditioned joint trajectory prediction is challenging in two aspects. First, real-world datasets provide ground truth

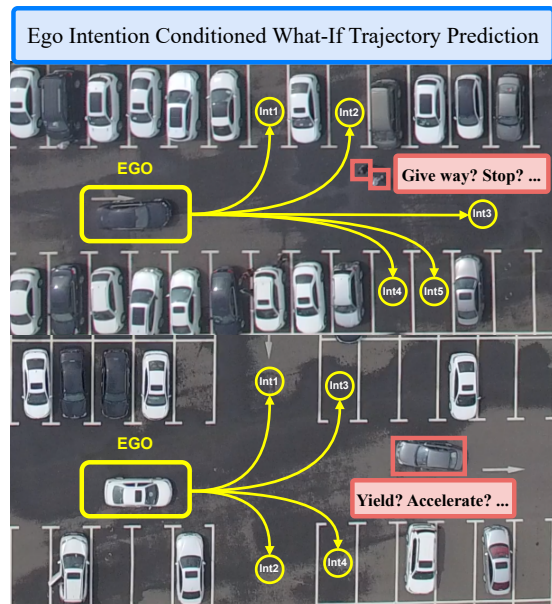


Fig. 1. Parking lot environments, characterized by their unstructured nature relative to on-road driving, require complex reasoning for autonomous navigation. The ego vehicle must perform what-if prediction, a process of evaluating multiple feasible ego future intentions while conditionally predicting the reactive behaviors of other agents to ensure safe and efficient decision making.

only for the realized ego intention. Counterfactual intentions lack supervision, so training relies on reliable pseudo-targets to cover unobserved what-if queries. Second, the introduction of intention selection and conditional decoding must generate scene-level trajectories that achieve high accuracy, while also ensuring safety and consistency.

In this work, we propose ParkDiffusion++, which performs ego intention conditioned multi-agent joint trajectory prediction for automated parking, addressing both challenges. ParkDiffusion++ comprises two stages. Stage 1 trains an ego intention tokenizer that predicts a discrete set of ego intentions from local context features. Stage 2 trains an ego intention conditioned joint predictor using ground truth and counterfactual intentions from Stage 1. Specifically, we build joint prediction on per-agent marginals combined via a searching and selecting process. Then, an EMA teacher decoder refined by a safety-guided denoiser provides pseudo targets for unobserved intentions. We perform extensive evaluations on the Dragon Lake Parking and Intersections Drone datasets, where results show that ParkDiffusion++ sets the state-of-the-art for ego intention conditioned joint trajectory prediction in parking scenarios. We present detailed ablation studies and

<sup>1</sup> Department of Computer Science, University of Freiburg, Germany.

<sup>2</sup> CARIAD SE, Germany.

This work was partially funded by CARIAD SE

qualitative analysis of what-if predictions, demonstrating how other agents react to different ego intentions.

To summarize, our main contributions are as follows:

- 1) The novel ParkDiffusion++ method for ego intention conditioned joint trajectory prediction.
- 2) A joint trajectory selector built on marginal models to ensure joint trajectory compatibility.
- 3) A safety-guided denoiser to improve prediction safety.
- 4) A counterfactual knowledge distillation module conditioned on different ego intentions using the safety-guided denoiser teacher.
- 5) A new baseline for ego intention conditioned joint trajectory prediction in automated parking and state-of-the-art performance compared to existing methods.

## II. RELATED WORK

In this section, we review the recent works that are related to and motivate our key contributions.

*Joint Trajectory Prediction:* Joint trajectory prediction targets a scene-consistent distribution over agent futures, rather than independent marginal forecasts. M2I [7] achieves joint trajectory prediction by identifying influencer-reactor agent pairs and combining marginal and conditional forecasts. JFP [8] builds on MultiPath++ [9], modeling pairwise interactions with a graphical formulation to produce scene-consistent trajectories. BiFF [10] fuses high-level future intentions with low-level interactive behaviors to yield joint, scene-consistent forecasts. OptTrajDiff [11] performs joint trajectory prediction using an optimized Gaussian diffusion process with clean-manifold guidance. Trajectory Mamba [12] replaces quadratic self-attention with a Selective-SSM encoder-decoder and joint polyline encoding to improve efficiency. Despite these advances, most joint predictors treat ego intent as latent and therefore cannot perform what-if trajectory prediction.

*Ego Conditioned Trajectory Prediction:* Ego conditioned prediction models how other agents respond when an explicit ego intention or trajectory is provided. PiP [13] conditions multi-agent forecasting on candidate ego plans produced by a planner, coupling prediction with planning. CBP [14] predicts other agents' future trajectories conditioned on a queried ego trajectory and quantifies interactivity. WIMP [15] injects counterfactual geometric goals or lanes to produce what-if forecasts of multi-actor reactions. SceneTransformer [16] unifies marginal, joint, and conditional modes and supports AV goal-conditioned prediction where other agents adapt to the chosen AV goal. ScePT [17] is a scene-consistent, policy-based joint predictor that supports conditional prediction by clamping one agent's (e.g., the ego's) future and rolling out coherent trajectories for the rest. MotionLM [18] casts multi-agent forecasting as language modeling over discrete motion tokens and supports temporally causal conditional rollouts for what-if queries. DTPP [19] employs a query-centric Transformer to perform efficient ego conditioned prediction for each branch in an ego trajectory tree. Most recent methods assume an externally specified ego trajectory or goal at test time or rely on planner-generated branches. To the best of our knowledge, none jointly learns a multi-modal

ego intention predictor and yields ego conditioned joint scene predictions.

*Counterfactual Knowledge Distillation:* Peysakhovich *et al.* [20] propose robust multi-agent counterfactual prediction from logged data, estimating outcomes under altered incentives or rules. Hart and Knoll [21] evaluate autonomous-driving policies by constructing counterfactual worlds to assess agent reactions without extra labels. Monti *et al.* [22] distill a long-observation teacher into a short-horizon student to enable accurate, low-latency trajectory forecasts. Bender *et al.* [23] distill causal behavior by generating counterfactual variants to correct Clever-Hans predictors. Jung *et al.* [24] analyze counterfactual fairness and align classifiers to fairness objectives via counterfactual knowledge distillation. To the best of our knowledge, counterfactual knowledge distillation has not been applied to ego intention conditioned "what-if" trajectory prediction.

*Guided Diffusion Models for Trajectory Prediction:* Guided diffusion steers the reverse denoising process with auxiliary signals to generate intention or safety-aware motion trajectories. MotionDiffuser [25] applies constrained sampling with differentiable costs to produce rule-compliant, physically plausible multi-agent futures. Diffusion-ES [26] employs gradient-free evolutionary search with black-box rewards to guide diffusion-based trajectory generation. Zheng *et al.* [27] integrate classifier guidance into trajectory scores for a transformer-based diffusion planner, enabling safe, multi-modal closed-loop driving. Plainer *et al.* [28] introduce an energy-based diffusion model with a Fokker-Planck derived consistency regularizer in physical simulation, conceptually related to consistency-aware guidance. Inspired by these approaches, we incorporate safety-guidance components into our denoising process to more effectively enforce safety constraints and enhance model performance.

*Trajectory Prediction for Automated Parking:* ParkPredict [29] pioneered parking prediction, analyzing how model complexity and feature sets affect intent estimation and long-term forecasts in CARLA. ParkPredict+ [4] addresses multimodal intent and trajectory prediction by combining CNNs and Transformers on a real-world parking dataset. ParkDiffusion [5] accurately predicts the multimodal trajectories of vehicles and pedestrians using diffusion models. However, prior parking work predicts trajectories marginally, without enforcing scene compatibility or interactivity. Furthermore, these methods do not model how surrounding agents adapt to changes in ego intention, nor do they explicitly incorporate safety constraints.

## III. METHOD

In this section, we present our method for ego intention conditioned joint trajectory prediction, illustrated in Fig. 2. We begin by formally defining the problem, then detail the components of our approach: an ego intention tokenizer, a conditional joint predictor featuring a reactive modeling and scoring mechanism, a pretrained safety-guided denoiser, and a Counterfactual Knowledge Distillation (CKD) scheme. The CKD queries alternative ego tokens and uses an Exponential Moving

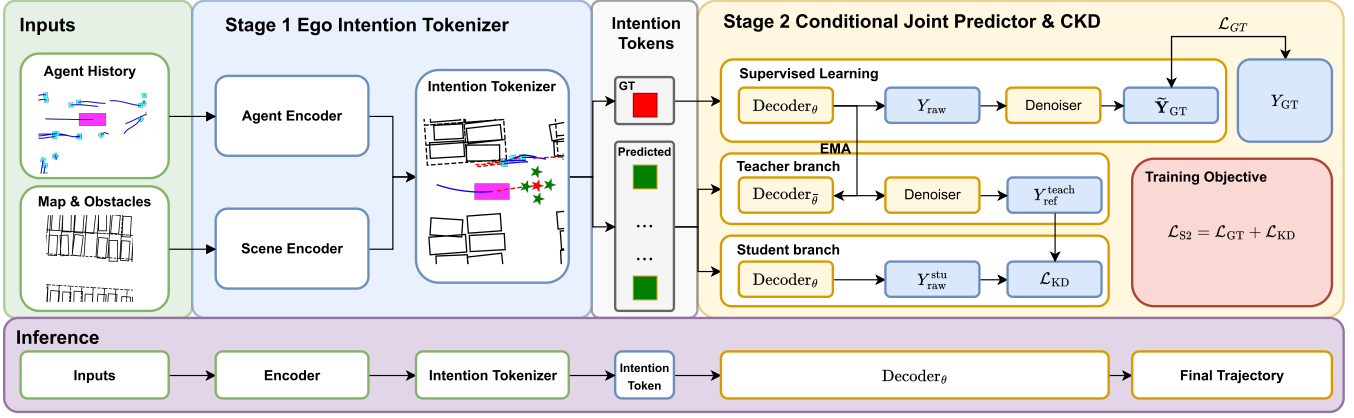


Fig. 2. **Overview of ParkDiffusion++.** Given agent histories and a vectorized map with parking slots and obstacles, Stage 1 predicts a bank of ego endpoint tokens, with the ground truth token in red and others in green. During Stage 2 training, we denote  $\text{Decoder}_\theta$  as the student conditional predictor, and  $\text{Decoder}_{\bar{\theta}}$  as EMA-updated teacher of  $\text{Decoder}_\theta$ . In the supervised learning branch, the ground truth token drives  $\text{Decoder}_\theta$  to yield a raw joint scene. We compute  $\mathcal{L}_{GT}$  on the raw output and use the frozen Denoiser to form a stop-gradient consistency target (not shown for simplicity). The Counterfactual Knowledge Distillation (CKD) module consists of a teacher and a student branch. The teacher branch runs  $\text{Decoder}_{\bar{\theta}}$  and refines its output with the same frozen Denoiser as supervised learning to form a teacher target, and the student branch runs  $\text{Decoder}_\theta$  and learns to match that teacher target featuring an additional safety penalty. The primary training objective of Stage 2 comprises the supervised learning loss and the counterfactual knowledge distillation loss ( $\mathcal{L}_{KD}$ ). At inference (bottom row), a token from Stage 1 is selected and fed to  $\text{Decoder}_\theta$  to yield the final joint trajectory. Internal steps such as per-agent marginals, beam-based assembly of joint scenes, and scene selection are abstracted within this block for clarity.

Average (EMA) teacher decoder refined by a frozen denoiser to produce safe pseudo-targets. Finally, we describe the two-stage training procedure that integrates these components.

### A. Problem Formulation

Given a scene observed in an ego-centric frame, the inputs consist of the past trajectories of  $N$  agents,  $\mathcal{X} = \{\mathbf{x}_{i,t}\}_{i=0}^{N-1} \}_{t=1}^{T_p}$ , where  $\mathbf{x}_{i,t} \in \mathbb{R}^D$  is the state of agent  $i$  (ego is  $i = 0$ ) containing its position, heading, velocity, and acceleration at time  $t$ . The inputs are augmented with scene context,  $\mathcal{C}_{\text{scene}}$ , composed of agent attributes (e.g., category  $c_i$ , size  $s_i$ ) and a vectorized map representation describing parking slots and static obstacles. Our core objective is two-fold. First, we aim to infer a probability distribution over a discrete set of  $K_{\text{intent}}$  high-level future intentions for the ego vehicle,  $\mathcal{Z} = \{z_k\}_{k=1}^{K_{\text{intent}}}$ , based on the observed past. Second, and centrally, we aim to model a conditional distribution over the joint future trajectories of all agents,  $\mathbf{Y} = \{\mathbf{y}_{i,t}\}_{i=0}^{N-1} \}_{t=1}^{T_f}$ , given a specific ego intention  $z_k$ :

$$p_\theta(\mathbf{Y} \mid \mathcal{X}, \mathcal{C}_{\text{scene}}, z_k). \quad (1)$$

### B. Ego Intention Prediction

The first stage of our method is an intention tokenizer that proposes a discrete set of plausible ego endpoints. By representing future possibilities as a compact set of intention points rather than complete trajectories, we create a tractable and interpretable interface for downstream planning modules. This can then perform efficient what-if reasoning over a finite set of outcomes.

1) *Scene Encoding:* The scene encoding architecture builds upon ParkDiffusion [5] to learn a comprehensive scene representation. Each agent’s history  $\mathbf{x}_{i,1:T_p}$  is processed by a temporal 1-D convolution followed by a GRU. The resulting sequence representation is then passed to a Transformer encoder over agents to capture local interactions, yielding  $\mathbf{f}_i \in \mathbb{R}^d$ . These features are then modulated by agent type via

a learned affine transform to yield type-aware features  $\tilde{\mathbf{f}}_i$ . The vectorized map is processed in two streams to distinguish between soft constraints (drivable boundaries  $\mathcal{M}_{\text{soft}}$ ) and hard constraints (physical obstacle segments  $\mathcal{M}_{\text{hard}}$ ). The ego feature  $\tilde{\mathbf{f}}_0$  queries each map stream via cross-attention. A global map summary  $\mathbf{m}$  is obtained by max-pooling over these embedded map features. Finally, the ego-agent’s feature, a max-pooled summary of features from all other agents  $\mathbf{f}_{\text{social}}$ , and the global map summary are fused to obtain a holistic scene context vector  $\mathbf{c}$ .

2) *Discrete Ego Intention Tokenization:* We predict a discrete set of ego endpoints using a bank of learned mode embeddings  $\{\mathbf{e}_k \in \mathbb{R}^{d_e}\}_{k=1}^{K_{\text{intent}}}$ , which allows the model to discover a compact basis of common maneuvers directly from data. For each mode  $k$ , we compute  $\mathbf{h}_k = f_{\text{mode}}(\mathbf{c}; \mathbf{e}_k)$ ; linear heads regress the endpoint token and a logit  $s_k$ , and a softmax over  $\{s_k\}_{k=1}^{K_{\text{intent}}}$  yields the categorical distribution.

### C. Conditional Joint Prediction

The second stage of our model is the conditional predictor, which generates full joint trajectories for all agents, contingent on a single ego intention token from the first stage.

1) *Ego Conditioned Context:* The selected ego intention token (denoted as  $\hat{\mathbf{g}}_k$ ) is first processed by an MLP to yield a rich embedding  $\boldsymbol{\tau}_k \in \mathbb{R}^{d_\tau}$ . For each agent  $i$  in the scene, we first form the context  $\mathbf{h}_i = [\tilde{\mathbf{f}}_i; \mathbf{f}_{\text{social}}; \mathbf{m}] \in \mathbb{R}^{d_h}$  as described in Sec. III-B.1. From  $\boldsymbol{\tau}_k$ , we apply feature-wise linear modulation (FiLM) to  $\mathbf{h}_i$ , yielding the intention conditioned feature vector  $\mathbf{u}_{i,k}$ .

2) *Reactive Social Modeling:* Our token-conditioned exposure gate learns to identify which agents are most relevant to the ego’s intention, preventing over-sensitivity to irrelevant agents and thus preserving joint prediction accuracy. Before computing the exposure, we define the terms in our formulation:  $\sigma(\cdot)$  is the sigmoid function,  $[\cdot]_+ = \max(0, \cdot)$  is the ReLU function,  $d_{\text{line}}(\cdot, \ell_k)$  is the minimum Euclidean distance

from a point to the line segment  $\ell_k$ , and  $\alpha, \beta, R_{\text{path}}, R_{\text{end}}$  are learnable scalar parameters. Let  $\ell_k$  be the line segment from the ego’s last position to the intention token  $\hat{\mathbf{g}}_k$ . We compute an exposure scalar for agent  $i$ ,  $e_i \in [0, 1]$ , based on its last observed position  $\mathbf{p}_i^{\text{last}}$ :

$$e_i = \sigma(\alpha[R_{\text{path}} - d_{\text{line}}(\mathbf{p}_i^{\text{last}}, \ell_k)]_+ + \beta[R_{\text{end}} - \|\mathbf{p}_i^{\text{last}} - \hat{\mathbf{g}}_k\|_2]_+). \quad (2)$$

A small gating network then uses this scalar to modulate the agent’s feature vector  $\mathbf{u}_{i,k}$ . This gate effectively learns a continuous, soft attention weight for each agent, conditioned on the ego’s specific intention. Unlike hard attention, which would select a binary subset of agents, this soft mechanism allows the model to fluidly modulate the degree of influence for all agents based on their geometric relevance to the ego’s intended maneuver.

3) *Joint Scene Prediction via Selection*: While we can decode  $M$  plausible marginal futures for each of the  $N$  agents, combining them into coherent joint scenes presents a combinatorial challenge. A brute-force approach would require evaluating  $M^N$  possible scenes, which is computationally intractable for even modest numbers of agents and modes. To address this, we adopt a tractable three-stage pipeline. First, for each agent, we generate a set of  $M$  marginal trajectory proposals along with compact per-mode features. Second, we use a fast beam search over the top- $R$  marginals per agent to construct a small, diverse set of  $K_{\text{scene}}$  high-quality joint scene candidates. We use  $K_{\text{scene}}$  to denote the number of joint candidates. Third, a learned scene selector scores and ranks these candidates. To evaluate a candidate scene, the selector scores each candidate using the global context and the sum of selected per-agent features. A softmax function is then applied to these scores for probability learning.

#### D. Safety-Guided Denoiser

We develop a safety-guided denoiser, a powerful tool for refining trajectories by combining a learned data-driven prior with explicit rule-based guidance. The network  $\epsilon_\psi$  is framed within the paradigm of score-based generative diffusion modeling and is trained on a simple objective to predict the Gaussian noise  $\xi$  that was added to a clean trajectory  $\mathbf{Y}_0$  at a given noise level  $\sigma$ . This process implicitly teaches the model the score function of the data distribution  $\nabla_{\mathbf{Y}} \log p(\mathbf{Y})$ , allowing it to understand what constitutes a “plausible” trajectory.

1) *Pretrained Denoiser*: We follow the structure of the Leapfrog Diffusion Model [30], similar to ParkDiffusion, to pretrain the denoiser  $\epsilon_\psi$  independently and freeze it during Stage 2. When we use the denoiser to refine a trajectory during training, its output serves as a non-differentiable target, therefore no gradients flow into  $\psi$  in Stage 2. When used for refinement, the denoiser is conditioned on the same intention token as the decoder so that  $C_{\text{tube}}$  and  $C_{\text{end}}$  use the ego-to-token line segment  $\ell_k$  and the token location  $\mathbf{g}_k$ .

2) *Deterministic Refinement with Geometric Guidance*: We denote this deterministic refinement operator by  $\text{Denoiser}_\psi(\cdot)$ , which implements the project-then-guide steps using the

pretrained score network  $\epsilon_\psi$ . Instead of a full, stochastic generation process, we use a few steps of a deterministic refinement algorithm called “project-then-guide.” This process combines the learned data prior from the denoiser with an explicit guidance step from a differentiable potential function,  $C(\mathbf{Y})$ . The refinement at each step  $s$  consists of two parts. First, in the projection step, we use the pretrained denoiser network to predict the noise in the current trajectory estimate  $\mathbf{Y}^{(s)}$ . Subtracting this noise yields an intermediate, cleaner trajectory,  $\mathbf{Y}^{(s+1/2)}$ , effectively projecting it onto the learned data manifold. Second, in the guidance step, we apply an explicit correction to this plausible trajectory using the gradient of our geometric potential function. This function acts as a differentiable critic, and its gradient provides a corrective force that pushes the trajectory away from constraint violations. The full update is:

$$\begin{aligned} \mathbf{Y}^{(s+1/2)} &= \mathbf{Y}^{(s)} - \epsilon_\psi(\mathbf{Y}^{(s)}, \sigma_s; \mathcal{C}_{\text{scene}}, \mathbf{g}_k), \\ \mathbf{Y}^{(s+1)} &= \mathbf{Y}^{(s+1/2)} - \eta_s \nabla_{\mathbf{Y}} C(\mathbf{Y}^{(s+1/2)}). \end{aligned} \quad (3)$$

3) *Geometric Potential Functions*: The potential function  $C(\mathbf{Y})$  sums five key terms that enforce safety, intention consistency, and physical plausibility:

- Agent-Agent Overlap ( $C_{\text{ov}}$ ): A hinge loss on the distance between agents’ safety radii ( $r_i, r_j$ ), denoted as

$$C_{\text{ov}} = \sum_{t, i \neq j} [(r_i + r_j - \|\mathbf{y}_{i,t} - \mathbf{y}_{j,t}\|_2)_+]^2. \quad (4)$$

- Obstacle Clearance ( $C_{\text{obs}}$ ): A hinge loss penalizing proximity to static obstacle segments, defined as

$$C_{\text{obs}} = \sum_{i,t} [(m_{\text{obs}} - d_{i,t}^{\min})_+]^2, \quad (5)$$

where  $m_{\text{obs}}$  and  $d_{i,t}^{\min}$  denote obstacle clearance margin, and Euclidean distance from the agent  $i$ ’s rectangular footprint at time  $t$ , respectively.

- Ego Path-Tube ( $C_{\text{tube}}$ ): Encourages ego trajectory  $\mathbf{y}_{0,t}$  to stay within a radius  $R_{\text{tube}}$  of its intended path  $\ell_k$ .
- Ego Endpoint Anchoring ( $C_{\text{end}}$ ): Ensures ego trajectory terminates near the conditioning intention token  $\hat{\mathbf{g}}_k$ .
- Motion Smoothness ( $C_{\text{sm}}$ ): Enforces kinematic plausibility by penalizing velocity and acceleration magnitudes.

#### E. Counterfactual Knowledge Distillation

A key limitation of imitation learning is that datasets only provide supervision for a single reality that occurred. To enable our model to reason about what would have happened under different ego intentions, we propose a Counterfactual Knowledge Distillation (CKD) framework.

Let  $\text{Decoder}_\theta$  denote the student decoder. We maintain an exponential moving average (EMA) teacher  $\text{Decoder}_{\bar{\theta}}$  of the student  $\text{Decoder}_\theta$ , updated after each step by

$$\bar{\theta} \leftarrow \tau \bar{\theta} + (1 - \tau) \theta, \quad \tau \in [0.99, 0.999]. \quad (6)$$

The denoiser and tokenizer remain frozen in Stage 2.

We sample a counterfactual token  $\hat{\mathbf{g}}_{\tilde{k}}$  from the Stage 1 bank ( $\tilde{k} \neq k^*$ ), then form a teacher target by first predicting with the EMA teacher-decoder and then refining with the frozen denoiser:

$$\mathbf{Y}_{\text{raw}}^{\text{teach}} = \text{Decoder}_{\bar{\theta}}(\mathcal{X}, \mathcal{C}_{\text{scene}}, \hat{\mathbf{g}}_k), \quad (7)$$

$$\mathbf{Y}_{\text{ref}}^{\text{teach}} = \text{Denoiser}_{\psi}(\mathbf{Y}_{\text{raw}}^{\text{teach}}, \mathcal{C}_{\text{scene}}, \hat{\mathbf{g}}_k). \quad (8)$$

We detach this refined target  $\text{sg}[\mathbf{Y}_{\text{ref}}^{\text{teach}}]$  to prevent gradients flowing into the teacher. The *student* prediction is produced by the current decoder without refinement:

$$\mathbf{Y}_{\text{raw}}^{\text{stu}} = \text{Decoder}_{\theta}(\mathcal{X}, \mathcal{C}_{\text{scene}}, \hat{\mathbf{g}}_k). \quad (9)$$

The distillation loss encourages the student to approximate the teacher’s refined scene while remaining collision-free:

$$\mathcal{L}_{\text{KD}} = \lambda_{\text{kd}} \left\| \mathbf{Y}_{\text{raw}}^{\text{stu}} - \text{sg}[\mathbf{Y}_{\text{ref}}^{\text{teach}}] \right\|_2^2 + \lambda_{\text{safe}} \text{CoL}_{\delta}(\mathbf{Y}_{\text{raw}}^{\text{stu}}), \quad (10)$$

where  $\text{CoL}_{\delta}(\cdot)$  denotes a differentiable collision-overlap penalty with clearance  $\delta$ . Only  $\theta$  is updated by  $\mathcal{L}_{\text{KD}}$ ,  $\bar{\theta}$  is updated by EMA (6) and  $\psi$  is frozen.

### F. Training Scheme

Our model is trained in two distinct stages to decouple intention discovery from intention conditioned dynamics.

1) *Stage 1: Ego Intention Tokenizer*: The objective is to accurately model the distribution of realized ego intentions. We use a “winner-takes-all” loss to optimize SmoothL1 for endpoint regression, cross-entropy for mode selection, and a diversity regularizer, weighted by  $\lambda_{\text{xy}}$ ,  $\lambda_{\text{cls}}$ ,  $\lambda_{\text{div}}$ .

2) *Stage 2: Conditioned Joint Prediction*: The primary objective  $\mathcal{L}_{\text{SUP}}$  uses the ground truth ego intention for direct supervision on both the per-agent marginal predictions and the final selected joint scene. We train with two branches: a supervised (GT) branch and, with probability  $p_{\text{cf}}$ , an additional counterfactual KD branch. Given the ground truth intention token denoted as  $\hat{\mathbf{g}}_{k^*}$ , the student decoder predicts a raw joint scene  $\mathbf{Y}_{\text{raw}} = \text{Decoder}_{\theta}(\mathcal{X}, \mathcal{C}_{\text{scene}}, \hat{\mathbf{g}}_{k^*})$ , which we refine with the frozen denoiser to obtain  $\tilde{\mathbf{Y}}_{\text{GT}} = \text{Denoiser}_{\psi}(\mathbf{Y}_{\text{raw}}, \mathcal{C}_{\text{scene}}, \hat{\mathbf{g}}_{k^*})$ . We supervise the raw prediction against ground truth and include a stop-gradient consistency term that encourages cooperation with the frozen refiner:

$$\begin{aligned} \mathcal{L}_{\text{GT}} = & \lambda_{\text{raw}} \mathcal{L}_{\text{SUP}}(\mathbf{Y}_{\text{raw}}, \mathbf{Y}^{\text{gt}}) \\ & + \lambda_{\text{cons}} \left\| \mathbf{Y}_{\text{raw}} - \text{sg}[\tilde{\mathbf{Y}}_{\text{GT}}] \right\|_2^2 \\ & + \lambda_{\text{safe}} \text{CoL}_{\delta}(\mathbf{Y}_{\text{raw}}). \end{aligned} \quad (11)$$

Gradients from  $\mathcal{L}_{\text{GT}}$  update  $\theta$  only;  $\psi$  remains frozen.

For each batch, we always run the supervised branch and, with probability  $p_{\text{cf}}$ , additionally run the counterfactual distillation branch. The Stage 2 objective is

$$\mathcal{L}_{\text{S2}} = \mathcal{L}_{\text{GT}} + \mathbb{I}\{\text{CF}\} \cdot \mathcal{L}_{\text{KD}}, \quad (12)$$

where  $\text{CF} \sim \text{Bernoulli}(p_{\text{cf}})$  indicates whether the counterfactual branch is used. Only the student parameters  $\theta$  receive gradients, and the denoiser  $\psi$  and tokenizer are frozen. The teacher-decoder parameters  $\bar{\theta}$  are updated by the EMA rule in (6). This hybrid scheme ensures the model accurately predicts realized outcomes while learning to generate safe and plausible responses to a wide range of counterfactual plans.

## IV. EXPERIMENTS

This section first presents the datasets, metrics, and baselines that we use. We then detail the implementation, followed by quantitative and qualitative results and an ablation study on each component in our model.

### A. Experimental Setup

1) *Datasets and Baselines*: We evaluate our approach on the Dragon Lake Parking (DLP) dataset [4], which features dense, unstructured parking lots, and the Intersections Drone (inD) dataset [31], which comprises urban intersections with challenging vehicle–pedestrian interactions, complementing the parking scenarios in DLP. We follow the dataset split of ParkDiffusion [5] for both datasets. For DLP, we use 51750 samples for training and 5750 samples for validation. For inD, we choose “Bendplatz” and “Frankenburg” scenarios, with 34014 samples for training and 4859 samples for validation.

To the best of our knowledge, there are no existing baselines that jointly learn ego intention prediction and conditional joint trajectory prediction. We therefore compare to strong ego-conditioned predictors that accept externally supplied ego inputs: WIMP [15], SceneTransformer [16], ScePT [17], MotionLM [18], and DTPP [19]. Since our implementation builds on ParkDiffusion [5], which yields per-agent multimodal marginal predictions, we include a naive random joint baseline that assembles  $K_{\text{scene}}$  joint scenes by uniformly sampling one marginal per agent. We denote this adapted baseline ParkDiffusion<sup>†</sup> and report its performance. For comparability in parking layouts, we use a shared vectorized-map frontend taken from ParkDiffusion for all methods and keep each baseline’s core architecture intact. This avoids layout-specific feature biases while preserving each method’s modeling capacity.

2) *Evaluation Metrics*: We predict  $K_{\text{scene}}$  joint scenes per data sample. Oracle (min–) metrics take the best over the  $K_{\text{scene}}$  candidates. The final (f–) metrics evaluate the top-1 scene selected by the model. ADE/FDE are average/final displacement errors in meters. MR counts a miss if none of the  $K_{\text{scene}}$  candidates have all valid agents with FDE  $\leq 2.0$  m (scene-level miss rate). OR denotes the overlap (collision) rate on the evaluated scene, where a scene is flagged if any vehicle-vehicle, vehicle-pedestrian, or vehicle-obstacle contact occurs. This extent of OR beyond agent-agent contacts to include agent-obstacle contacts is a special design feature from us for parking environments, where collisions between vehicles and obstacles frequently occur. mAP is the area under the precision–recall curve over the  $K_{\text{scene}}$  joint candidates per sample, labeling a candidate positive if all valid agents have FDE  $\leq 2.0$  m.

3) *Implementation Details*: All positions are in an ego-centric frame, pre-aligned by preprocessing. We observe  $T_p=10$  past steps and predict  $T_f=10$  future steps at 0.4 s intervals for a horizon of 4.0 s. Stage 1 trains the tokenizer for 20 epochs and proposes  $K_{\text{intent}}=6$  tokens. Stage 2 freezes the tokenizer and trains the conditional predictor for 20 epochs, decoding  $M=6$  marginals per agent and assembling  $K_{\text{scene}}=6$  joint candidates. Counterfactual Knowledge Distillation (CKD) is applied with probability  $p_{\text{cf}}=0.5$  per batch. The teacher path uses the EMA decoder under a sampled non-GT token and refines that prediction with  $t_{\text{steps}}=5$  denoiser steps. These parameters remain the same in benchmarking experiments, and in ablation studies, we state otherwise. We use AdamW for training with a learning rate of  $1e-4$  on a cosine

TABLE I  
BENCHMARKING RESULTS ON THE DRAGON LAKE PARKING (DLP) DATASET.

Model	minADE[m]	minFDE[m]	minMR	minOR	mAP	f-ADE[m]	f-FDE[m]	f-MR	f-OR	f-mAP
ParkDiffusion <sup>†</sup> [5]	0.52	1.06	0.48	0.21	0.68	0.90	1.78	0.63	0.28	0.57
WIMP [15]	0.62	1.20	0.46	0.18	0.70	0.84	1.60	0.55	0.24	0.62
SceneTransformer [16]	0.36	0.79	0.28	0.07	0.92	0.55	1.08	0.41	0.12	0.84
ScPT [17]	0.38	0.86	0.30	0.08	0.90	0.61	1.28	0.37	0.11	0.88
MotionLM [18]	<u>0.32</u>	0.80	0.29	0.08	<u>0.95</u>	<b>0.34</b>	<u>0.88</u>	0.35	0.10	<u>0.94</u>
DTPP [19]	<u>0.32</u>	0.82	<u>0.27</u>	<u>0.06</u>	0.94	0.50	1.21	<u>0.33</u>	<u>0.09</u>	<u>0.92</u>
ParkDiffusion++ (Ours)	<b>0.29</b>	<b>0.56</b>	<b>0.23</b>	<b>0.03</b>	<b>0.97</b>	<u>0.36</u>	<b>0.66</b>	<b>0.29</b>	<b>0.05</b>	<b>0.95</b>

We report results for multiple baselines on the Dragon Lake Parking (DLP) dataset [4] validation split. Evaluation metrics include oracle (min-) and final (f-) ADE, FDE, MR, OR, mAP. Lower is better for ADE/FDE/MR/OR, and higher is better for mAP. Bold and underline denote best and second-best results, respectively.

TABLE II  
BENCHMARKING RESULTS ON THE INTERSECTIONS DRONE (INd) DATASET.

Model	minADE[m]	minFDE[m]	minMR	minOR	mAP	f-ADE[m]	f-FDE[m]	f-MR	f-OR	f-mAP
ParkDiffusion <sup>†</sup> [5]	0.86	1.60	0.58	0.31	0.66	1.12	2.05	0.68	0.38	0.54
WIMP [15]	0.78	1.48	0.52	0.28	0.72	0.99	1.82	0.61	0.33	0.60
SceneTransformer [16]	0.58	1.10	0.43	0.19	0.84	0.86	1.65	0.51	0.22	0.69
ScPT [17]	<b>0.40</b>	0.88	0.36	0.15	0.91	0.67	1.35	0.47	0.18	0.74
MotionLM [18]	0.48	1.05	0.39	0.15	0.89	<u>0.63</u>	1.31	<u>0.43</u>	0.18	<u>0.81</u>
DTPP [19]	0.49	0.96	0.38	<u>0.13</u>	0.87	<b>0.61</b>	<u>1.22</u>	0.45	<u>0.16</u>	0.76
ParkDiffusion++ (Ours)	<u>0.42</u>	<b>0.78</b>	<b>0.32</b>	<b>0.07</b>	<b>0.93</b>	0.66	<b>1.18</b>	<b>0.38</b>	<b>0.09</b>	<b>0.83</b>

We also report results on the inD (drone-based intersections) dataset [31] validation split. Evaluation metrics include oracle (min-) and final (f-) ADE, FDE, MR, OR, mAP. Lower is better for ADE/FDE/MR/OR, and higher is better for mAP. Bold and underline denote best and second-best results, respectively.

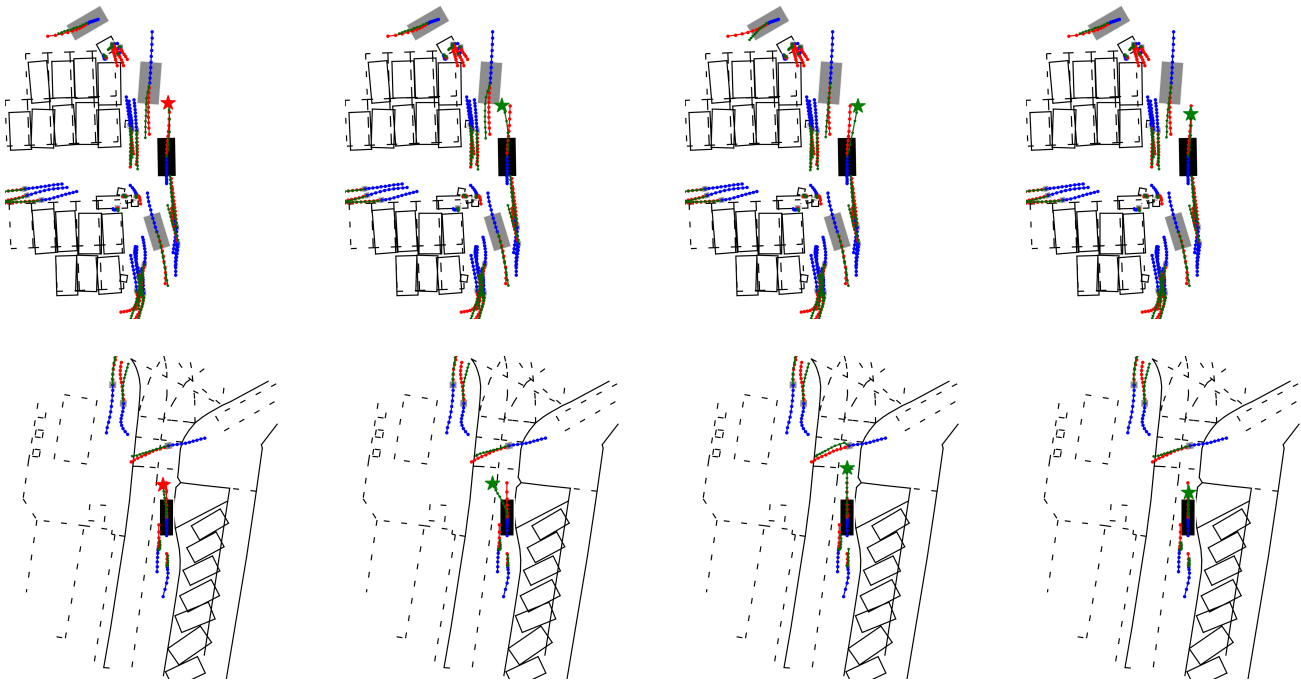


Fig. 3. Visualizations of ParkDiffusion++ predictions on the DLP (top row) and inD (bottom row) datasets. For both rows, the leftmost figure shows how our model reacts to the most likely intention (red star) from Stage 1. The other three figures in each row show the joint prediction result conditioned on other representative predicted intentions (green star). The ego vehicle is depicted in solid black at the center, while other agents are shown in gray. Trajectories are color-coded as follows: past in blue, ground-truth future in red, and predicted in green.

schedule. We train the model on two NVIDIA A100 GPUs.

### B. Quantitative Analysis

On the DLP validation split, ParkDiffusion++ achieves the best oracle metrics among the compared methods. For

final metrics, ParkDiffusion++ yields the best f-FDE, f-MR, f-OR, and f-mAP, with f-ADE of 0.36m following the best of 0.34m by MotionLM. On the inD validation split, ParkDiffusion++ achieves the best oracle metrics in minFDE, minMR, minOR, mAP, and the second-best minADE of

TABLE III  
MAIN COMPONENTS ABLATION

Method	minADE	minFDE	MR	OR	mAP
ParkDiffusion <sup>†</sup> [5]	0.52	1.06	0.48	0.21	0.68
+ JS	0.37	0.75	0.38	0.12	0.85
+ JS + SGD	0.32	0.60	0.26	0.05	0.92
+ JS + SGD + CKD	<b>0.29</b>	<b>0.56</b>	<b>0.23</b>	<b>0.03</b>	<b>0.97</b>

Ablation on the main components. Starting from the vanilla ParkDiffusion<sup>†</sup>, we add in sequence Joint Selector (JS), Safety-Guided Denoiser (SGD), and CKD to quantify each contribution. Evaluation metrics are the oracle ones including minADE, minFDE, MR, OR, and mAP.

TABLE IV  
EGO INTENTION TOKENIZER ABLATION

Setting	minADE	minFDE	MR	OR	mAP
<i>Capacity <math>K_{intent}</math></i>					
1 (GT token)	0.31	0.60	0.27	0.04	0.96
6	<b>0.29</b>	<b>0.56</b>	<b>0.23</b>	<b>0.03</b>	<b>0.97</b>
12	0.30	0.58	<b>0.23</b>	<b>0.03</b>	<b>0.97</b>
18	0.34	0.62	0.25	0.04	0.94
<i>Token source at <math>K_{intent}=6</math></i>					
Stage 1 ranking	<b>0.29</b>	<b>0.56</b>	<b>0.23</b>	<b>0.03</b>	<b>0.97</b>
Stage 1 random	0.33	0.63	0.34	0.07	0.92
GT + noise	0.31	0.60	0.27	0.04	0.96

Ablation on the Ego Intention Tokenizer. Intention capacity describes how many intention tokens we generate for the downstream CKD module. Token source includes three token selection strategies, namely ranking Stage 1 outputs and select the highest possible ones (Stage 1 ranking), random sampling from the Stage 1 outputs (Stage 1 random), and sampling around ground truth token with random noise (GT + noise).

0.42m. For final metrics, we achieve the best in f-FDE, f-MR, f-OR, f-mAP. DTPP model achieves the best f-ADE of 0.61m. We observe consistent trends across datasets. In both datasets, we achieve the best FDE, MR, OR, mAP for both oracle and final. However, minADE and f-ADE are less satisfying because the knowledge learned from counterfactual knowledge distillation and the safety-guided denoiser emphasize safety instead of efficiency. In real-world driving, human drivers sometimes adopt a more aggressive strategy for efficiency.

### C. Qualitative Analysis

Fig. 3 visualizes the prediction results of ParkDiffusion++ on both datasets. On the DLP dataset, our model accurately predicts the ego intention, and the resulting conditional joint trajectory predictions closely align with the ground truth. When conditioned on counterfactual ego intentions, other agents exhibit plausible reactive behaviors to avoid potential collisions. For instance, in the second panel of the top row, a vehicle positioned to the ego’s front-left adjusts its trajectory to yield, thereby preventing a potential conflict. Similarly, on the inD dataset, the model effectively captures the scene’s dynamic patterns when conditioned on the most likely intention. In the third panel, a pedestrian is predicted to yield in response to the ego vehicle’s potential acceleration.

We find that our model presents a conservative tendency in the predictions across both datasets. For example, on the DLP dataset, the vehicle yielding to the ego agent we mentioned actually yields further left to create more space

TABLE V  
COUNTERFACTUAL KNOWLEDGE DISTILLATION ABLATION

Setting	minADE	minFDE	MR	OR	mAP
<i>Counterfactual sampling probability <math>p_{cf}</math></i>					
0.00	0.32	0.60	0.26	0.05	0.92
0.25	0.31	0.59	0.28	0.04	0.96
0.50	<b>0.29</b>	<b>0.56</b>	<b>0.23</b>	<b>0.03</b>	<b>0.97</b>
0.75	0.35	0.64	0.27	0.04	0.94
<i>Teacher configuration</i>					
None	0.32	0.60	0.26	0.05	0.92
EMA only	0.32	0.59	0.26	0.05	0.93
EMA + Denoiser(U)	0.30	0.58	0.25	0.05	0.93
EMA + SGD	<b>0.29</b>	<b>0.56</b>	<b>0.23</b>	<b>0.03</b>	<b>0.97</b>

Ablation on the Counterfactual Knowledge Distillation. We show the ablation from two perspectives: counterfactual sampling probability describing the weight of counterfactual learning over supervised learning, and teacher configurations including None (no CKD at all), EMA only, EMA + Denoiser(U) (EMA + Denoiser without safety guided potential functions), and EMA + SGD (EMA + Safety-Guided Denoiser).

from the ego vehicle, even though this brings it closer to more vulnerable pedestrian agents. This behavior may be attributed to a lack of prioritized collision penalties in the model’s training objective. A similar conservative trend is observed on the inD dataset, where a pedestrian behind the ego vehicle is predicted to move slower than its ground truth, and two other pedestrians on the upper left adjust their paths to avoid a mutual collision. Collectively, these qualitative examples demonstrate the effectiveness of the conditional joint prediction model with safety constraints, matching the strong quantitative results where our method achieved the lowest oracle and final OR over all baselines.

### D. Ablation Studies

We perform ablation studies on the DLP validation dataset to provide insights into the impact of the design of individual components.

1) *Components Analysis*: We ablate on different components and present results in Tab. III. Random composition of independent marginals (ParkDiffusion<sup>†</sup>) typically yields poor MR/OR and weak mAP. Adding the Joint Selector (JS) converts it to a truly learned joint prediction, improving all the metrics by a large margin. The Safety-Guided Denoiser (SGD) further improves the overall performance, and most notably, the safety-critical OR metrics. The Counterfactual Knowledge Distillation (CKD) transfers teacher-refined behavior to the student, continues to improve the general performance across all dimensions.

2) *Ego Intention Tokenizer*: We ablate the ego intention tokenizer capacity and token source in Tab. IV. We observe that the performance saturates around  $K_{intent}=6$ . Increasing the number of counterfactual tokens to 12 brings negative effects to accuracy, and when it is increased to 18, performance across all dimensions drops by a large margin, indicating that the counterfactual knowledge starts to confuse the main joint predictor. For the token selection strategy in CKD training, the Stage 1 ranking strategy outperforms the other two, reflecting that this strategy brings extra environment awareness knowledge to Stage 2 to help the learning process.

3) *Counterfactual Knowledge Distillation (CKD)*: We ablate on the counterfactual knowledge distillation from two perspectives, counterfactual sampling probability and teacher configuration. We present the results in Tab. V. We find that balancing the supervised learning with the CKD modules achieves the best performance. When CKD module has larger weights, the performance drops quickly. Regarding different teacher configurations, adding EMA module alone helps stabilize the training but leads to no explicit gain in the metrics. Adding a denoiser without safety guidance based on EMA helps to improve all metrics slightly. With the Safety-Guided Denoiser (SGD) added in the last step, the MR, OR, mAP scores improve by a large margin.

## V. CONCLUSION

We presented ParkDiffusion++, a two-stage approach for ego intention conditioned joint trajectory prediction in automated parking. Stage 1 predicts a compact set of ego intention tokens, and Stage 2 performs intention conditioned joint prediction augmented by counterfactual knowledge distillation and safety-guided modules. Experiments show that ParkDiffusion++ achieves the state-of-the-art quantitative results, and present qualitative what-if visualizations. Limitations include reliance on handcrafted geometric potentials in the denoiser and the remaining gap between offline counterfactual supervision and closed-loop control. Future work could explore how to learn guidance potentials from data and how to couple the conditional predictor with smarter decision making, planning, and control.

## REFERENCES

- [1] A. Distelzweig, A. Look, E. Kosman, F. Janjoš, J. Wagner, and A. Valada, "Stochasticity in motion: An information-theoretic approach to trajectory prediction," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2025, pp. 6211–6218.
- [2] S. Marvi, C. Rist, J. Schmidt, J. Jordan, and A. Valada, "Evidential uncertainty estimation for multi-modal trajectory prediction," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2025, pp. 785–792.
- [3] A. Distelzweig, E. Kosman, A. Look, F. Janjoš, D. K. Manivannan, and A. Valada, "Motion forecasting via model-based risk minimization," in *IEEE Int. Conf. on Robotics and Automation*, 2025, pp. 7011–7018.
- [4] X. Shen, M. Lacayo, N. Guggilla, and F. Borrelli, "Parkpredict+: Multimodal intent and motion prediction for vehicles in parking lots with cnn and transformer," in *IEEE Int. Conf. on Intelligent Transportation Systems*, 2022, pp. 3999–4004.
- [5] J. Wei, N. Vödisch, A. Rehr, C. Feist, and A. Valada, "Parkdiffusion: Heterogeneous multi-agent multi-modal trajectory prediction for automated parking using diffusion models," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2025.
- [6] N. Radwan, W. Burgard, and A. Valada, "Multimodal interaction-aware motion prediction for autonomous street crossing," *The International Journal of Robotics Research*, vol. 39, no. 13, pp. 1567–1598, 2020.
- [7] Q. Sun, X. Huang, J. Gu, B. C. Williams, and H. Zhao, "M2i: From factored marginal trajectory prediction to interactive prediction," in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2022, pp. 6543–6552.
- [8] W. Luo, C. Park, *et al.*, "Jfp: Joint future prediction with interactive multi-agent modeling for autonomous driving," in *Int. Conf. on Learning Representations*, 2023, pp. 1457–1467.
- [9] B. Varadarajan, A. Hefny, *et al.*, "Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction," in *IEEE Int. Conf. on Robotics and Automation*, 2022, pp. 7814–7821.
- [10] Y. Zhu, D. Luan, and S. Shen, "Biff: Bi-level future fusion with polyline-based coordinate for interactive trajectory prediction," in *IEEE/CVF Int. Conf. on Computer Vision*, 2023, pp. 8260–8271.
- [11] Y. Wang, C. Tang, L. Sun, S. Rossi, Y. Xie, C. Peng, T. Hannagan, S. Sabatini, N. Poerio, M. Tomizuka, *et al.*, "Optimizing diffusion models for joint trajectory prediction and controllable generation," in *Eur. Conf. on Computer Vision*, 2024, pp. 324–341.
- [12] Y. Huang, Y. Cheng, and K. Wang, "Trajectory mamba: Efficient attention-mamba forecasting model based on selective ssm," in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2025, pp. 12 058–12 067.
- [13] H. Song, W. Ding, Y. Chen, S. Shen, M. Y. Wang, and Q. Chen, "Pip: Planning-informed trajectory prediction for autonomous driving," in *Eur. Conf. on Computer Vision*, 2020, pp. 598–614.
- [14] E. Tolstaya, R. Mahjourian, C. Downey, B. Vadarajan, B. Sapp, and D. Anguelov, "Identifying driver interactions via conditional behavior prediction," in *IEEE Int. Conf. on Robotics and Automation*, 2021, pp. 3473–3479.
- [15] S. Khandelwal, W. Qi, J. Singh, A. Hartnett, and D. Ramanan, "What-if motion prediction for autonomous driving," *arXiv preprint arXiv:2008.10587*, 2020.
- [16] J. Ngiam, B. Caine, *et al.*, "Scene transformer: A unified architecture for predicting multiple agent trajectories," in *Int. Conf. on Learning Representations*, 2022.
- [17] Y. Chen, B. Ivanovic, and M. Pavone, "Scept: Scene-consistent, policy-based trajectory predictions for planning," in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2022, pp. 17 103–17 112.
- [18] A. Seff, B. Cera, D. Chen, M. Ng, A. Zhou, N. Nayakanti, K. S. Refaat, R. Al-Rfou, and B. Sapp, "Motionlm: Multi-agent motion forecasting as language modeling," in *IEEE/CVF Int. Conf. on Computer Vision*, 2023, pp. 8579–8590.
- [19] Z. Huang, P. Karkus, B. Ivanovic, Y. Chen, M. Pavone, and C. Lv, "Dtpp: Differentiable joint conditional prediction and cost evaluation for tree policy planning in autonomous driving," in *IEEE Int. Conf. on Robotics and Automation*, 2024, pp. 6806–6812.
- [20] A. Peysakhovich, C. Kroer, and A. Lerer, "Robust multi-agent counterfactual prediction," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [21] P. Hart and A. Knoll, "Counterfactual policy evaluation for decision-making in autonomous driving," *arXiv preprint arXiv:2003.11919*, 2020.
- [22] A. Monti, A. Porrello, S. Calderara, P. Coscia, L. Ballan, and R. Cucchiara, "How many observations are enough? knowledge distillation for trajectory forecasting," in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2022, pp. 6553–6562.
- [23] S. Bender, C. J. Anders, P. Chormai, H. A. Marxfeld, J. Herrmann, and G. Montavon, "Towards fixing clever-hans predictors with counterfactual knowledge distillation," in *IEEE/CVF Int. Conf. on Computer Vision*, 2023, pp. 2607–2615.
- [24] S. Jung, S. Yu, S. Chun, and T. Moon, "Do counterfactually fair image classifiers satisfy group fairness?—a theoretical and empirical study," *Advances in Neural Information Processing Systems*, vol. 37, pp. 56 041–56 053, 2024.
- [25] C. Jiang, A. Corrman, C. Park, B. Sapp, Y. Zhou, D. Anguelov, *et al.*, "Motiondiffuser: Controllable multi-agent motion prediction using diffusion," in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2023, pp. 9644–9653.
- [26] B. Yang, H. Su, N. Gkanatsios, T.-W. Ke, A. Jain, J. Schneider, and K. Fragkiadaki, "Diffusion-es: Gradient-free planning with diffusion for autonomous driving and zero-shot instruction following," *arXiv preprint arXiv:2402.06559*, 2024.
- [27] Y. Zheng, R. Liang, K. Zheng, J. Zheng, L. Mao, J. Li, W. Gu, R. Ai, S. E. Li, X. Zhan, *et al.*, "Diffusion-based planning for autonomous driving with flexible guidance," *arXiv preprint arXiv:2501.15564*, 2025.
- [28] M. Plainer, H. Wu, L. Klein, S. Günemann, and F. Noé, "Consistent sampling and simulation: Molecular dynamics with energy-based diffusion models," *arXiv preprint arXiv:2506.17139*, 2025.
- [29] X. Shen, I. Batkovic, *et al.*, "Parkpredict: Motion and intent prediction of vehicles in parking lots," in *IEEE Intelligent Vehicles Symposium*, 2020, pp. 1170–1175.
- [30] W. Mao, C. Xu, Q. Zhu, S. Chen, and Y. Wang, "Leapfrog diffusion model for stochastic trajectory prediction," in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2023, pp. 5517–5526.
- [31] J. Bock, R. Krajewski, T. Moers, S. Runde, L. Vater, and L. Eckstein, "The ind dataset: A drone dataset of naturalistic road user trajectories at german intersections," in *IEEE Intelligent Vehicles Symposium*, 2020, pp. 1929–1934.