
Robust Trajectory Prediction against Adversarial Attacks

Yulong Cao^{*1,2}, Danfei Xu^{2,3}, Xinshuo Weng², Z. Morley Mao¹, Anima Anandkumar^{2,4}, Chaowei Xiao^{2,5}, and Marco Pavone^{2,6}

¹University of Michigan

²NVIDIA

³Georgia Institute of Technology

³California Institute of Technology

⁴Arizona State University

⁵Stanford University

Abstract

Trajectory prediction using deep neural networks (DNNs) is an essential component of autonomous driving (AD) systems. However, these methods are vulnerable to adversarial attacks, leading to serious consequences such as collisions. In this work, we identify two key ingredients to defend trajectory prediction models against adversarial attacks including (1) designing effective adversarial training methods and (2) adding domain-specific data augmentation to mitigate the performance degradation on clean data. We demonstrate that our method is able to improve the performance by 46% on adversarial data and at the cost of only 3% performance degradation on clean data, compared to the model trained with clean data. Additionally, compared to existing robust methods, our method can improve performance by 21% on adversarial examples and 9% on clean data. Our robust model is evaluated with a planner to study its downstream impacts. We demonstrate that our model can significantly reduce the severe accident rates (e.g., collisions and off-road driving)¹.

1 Introduction

Trajectory prediction is a critical component of modern autonomous driving (AD) systems. It allows an AD system to anticipate the future behaviors of other nearby road participants and plan its actions accordingly. Recent trajectory prediction models built on Deep Neural Networks (DNN) have demonstrated state-of-the-art performance on large-scale benchmarks [1–7], showing a promising path towards learning-based trajectory prediction for AD systems. As trajectory prediction plays an important role in AD systems, accurate predictions are required for making safe driving decisions. It is crucial to understand how unknown scenarios will affect trajectory predictions and then bolster the robustness of such trajectory predictions in return.

To achieve this goal, adversarial attacks [8–10] are often used as a proxy to measure the worst-case performance of the model when facing unseen scenarios. Similarly, we use a standard adversarial attack setup [11] for trajectory predictions. As illustrated in Fig. 1, an adversarial agent (red vehicle) aims to cause a traffic accident. It drives along a carefully designed trajectory (i.e., adv history) to influence the trajectory prediction model of the Autonomous Vehicle (green vehicle). Such an

^{*}Work done during an internship at NVIDIA

¹Our project website is at <https://robustav.github.io/RobustTraj>




Figure 1: Overview of *RobustTraj* preventing Autonomous Vehicle (AV) from collisions when its trajectory prediction model is under adversarial attacks. When the trajectory prediction model is under attack, the AV predicts the wrong future trajectory of the other agent turning right (yellow vehicle). This results in AV speeding up instead of slowing down, and eventually colliding into the other vehicle.

adversary can critically compromise the predicted trajectories of all other agents by altering its route in inconspicuous ways. By fooling the trajectory prediction models, it can further affect downstream planning of the AV systems and cause serious consequences. Using the adversarial attack as the proxy, this work aims to develop effective techniques to bolster the robustness of trajectory prediction models against adversarial attacks and improve the AD’s safety under uncertain scenarios.

At the same time, adversarial robustness for machine learning is a widely-studied area, but most works focus on classification tasks [12–23]. Among the proposed techniques, adversarial training [8] remains the most effective and widely used method to defend classifiers against adversarial attacks. The general strategy of adversarial training is to solve a min-max game by generating adversarial examples for a model at each training step and then optimizing the model to make correct predictions for these samples. However, directly applying adversarial training to trajectory prediction presents a number of critical technical challenges.

First, most trajectory prediction methods employ probabilistic generative models to cope with the uncertainty in motion forecasting [2–7]. As we will show in this paper, the stochastic components of these models (e.g., posterior sampling in VAEs) can obfuscate the gradients that guide the adversarial generation, making naïve adversarial training methods ineffective. Second, adversarial training on trajectory prediction task aims to model joint data distribution of future trajectories and adversarial past trajectories. However, the co-evolution of the adversarial sample distribution and the prediction model during the training process makes the joint distribution hard to model and often destabilizes the adversarial training. Finally, prior work [23] shows that adversarial training often leads to degraded performance on clean (unperturbed) data, while retaining good performance in benign cases is crucial due to the critical role of trajectory prediction for AVs. Hence, an effective adversarial training method must carefully balance the benign and the adversarial performance of a model.

Our approach. We propose an adversarial training framework for trajectory predictions named *RobustTraj*, by addressing the aforementioned challenges. First, to address the issue of an obfuscated gradient in adversarial generation due to stochastic components, we devise a *deterministic attack* that creates a deterministic gradient path within a probabilistic model to generate adversarial samples. Second, to address the challenge of an unstable training process due to shift in adversarial distributions, we introduce a hybrid objective that interleaves the adversarial training and learning from clean data to anchor the model output on stable clean data distribution. Finally, to achieve balanced performances on both adversarial and clean data, we introduce a domain-specific data augmentation technique for trajectory prediction via a dynamic model. This data augmentation technique generates diverse, realistic, and dynamically-feasible samples for training and achieves a better performance trade-off on clean and adversarial data.

We empirically show that *RobustTraj* can effectively defend two different types of probabilistic trajectory prediction models [4, 7] against adversarial attacks, while incurring minimal performance degradation on clean data. For instance, *RobustTraj* can increase the adversarial performance of AgentFormer [4], a state-of-the-art trajectory prediction model, by 46% at the cost of 3% performance

drop on clean data. To further show impacts of our method on the AD stack, we plug our robust trajectory prediction model into a planner and demonstrate that our model reduces serious accidents rates (e.g., collisions and off-road driving) under attacks by 100%, compared to the standard non-robust model trained using only clean data.

2 Related Work

Adversarial attacks and defenses on trajectory prediction. A recent work began to study the adversarial robustness of trajectory prediction models [11]. Zhang et al. [11] demonstrated that perturbing agents’ observed trajectory can adversarially impact the prediction accuracy of a DNN-based trajectory forecasting model. To mitigate the issue, Zhang et al. [11] proposed several defense methods such as data augmentation and trajectory smoothing. However, these methods are less effective when facing adaptive attacks [24]. In our work, we propose to use adversarial training which provides the general adversarial robustness that can resist adaptive attacks.

Adversarial scenario generation. A few recent studies work on generating adversarial traffic scenarios such that the autonomous driving systems fail to make safe driving decisions [25, 26]. However, generating realistic traffic scenarios is challenging and the generated adversarial scenarios can be unrealistic and violate traffic rules by directly optimizing the latent vectors of the traffic model [26]. In this work, we consider defending against realistic adversarial scenarios grounded on the scenarios from a dataset. Wang et al. [25] perturb the raw input data to mislead the full stack AV system. However, in this work, our primary goal is to study and improve the robustness of trajectory prediction models. To obtain salient and unambiguous insights, we minimize the conflating factors in our analysis without considering the perception model.

Adversarial training. A variety of adversarial training methods have been proposed to defend DNN-based models against adversarial attacks [8, 19, 12–23]. The most common strategy is to design a min-max game with the inner maximization process and outer minimization process. The inner maximization process generates adversarial examples that maximize an adversarial objective (e.g., make wrong prediction). The outer minimization process then updates the model parameters to minimize the error on the adversarial examples. Several recent works also propose to mix clean data and adversarial examples for improving robustness [27, 28] and performance on clean data [18]. Although there exists a large body of literature in studying adversarial robustness for machine learning, most focus on the problem of discriminative model (e.g., object recognition), leaving other problem domains (e.g., conditional generative models) largely unexplored. In this work, we develop a novel adversarial training method for trajectory prediction models, where most state-of-the-art trajectory prediction models are generative and probabilistic, by addressing a number of critical technical challenges.

3 Preliminaries and Formulation

Trajectory prediction. The goal is to predict future trajectory distribution $\mathcal{P}_\theta(Y|X)$ of N agents conditioned on their H history time states $\mathbf{X} = (\mathbf{X}^{-H+1}, \dots, \mathbf{X}^0)$, and other environment context such as maps² to predict T time-step future trajectories $\mathbf{Y} = (\mathbf{Y}^1, \dots, \mathbf{Y}^T)$. For observed time steps $t \leq 0$, we denote the agent states as $\mathbf{X}^t = (x_1^t, \dots, x_i^t, \dots, x_N^t)$, where x_i^t is the state of agent i at the time step t . Similarly, $\mathbf{Y}^t = (y_1^t, \dots, y_N^t)$ denotes the states of N agents at a future time step t ($t > 0$). We denote the ground truth and the predicted future trajectories as \mathbf{Y} and $\hat{\mathbf{Y}}$, respectively. We denote the history information encoded by a function f as the decision context $\mathbf{C} = f(\mathbf{X})$.

Probabilistic trajectory prediction models. In this work, we focus on defending generative, probabilistic trajectory prediction models, as they have demonstrated superior performance in modeling uncertainty in predicting future motions [2–7]. We consider the two most popular types of generative models: conditional variational encoders (CVAEs) and conditional GANs (cGANs), both can be viewed as latent variable models. We define latent variables $\mathbf{Z} = \{z_1, \dots, z_i, \dots, z_N\}$ where z_i represents the latent variable of the agent i . CVAE formulates the generative problem as: $p_\theta(\mathbf{Y}|\mathbf{X}) = \int p_\theta(\mathbf{Y}|\mathbf{X}, \mathbf{Z}) \cdot p_\theta(\mathbf{Z}|\mathbf{X})d\mathbf{Z}$, where $p_\theta(\mathbf{Z}|\mathbf{X})$ is a conditional Gaussian prior

²For simplicity, we ignore contextual information.

$(\mathcal{N}(p_\theta^\mu(\mathbf{Z}|\mathbf{X}), p_\theta^\sigma(\mathbf{Z}|\mathbf{X})))$ with mean $p_\theta^\mu(\mathbf{Z}|\mathbf{X})$ and standard deviation $p_\theta^\sigma(\mathbf{Z}|\mathbf{X})$; $p_\theta(\mathbf{Y}|\mathbf{X}, \mathbf{Z})$ is a conditional likelihood model. The model is usually trained through optimizing a negative evidence lower objective [4]:

$$\begin{aligned} \mathcal{L}_{\text{total}} &= \mathcal{L}_{\text{elbo}} + \mathcal{L}_{\text{diversity}} \\ &= -\mathbb{E}_{q_\phi(\mathbf{Z}|\mathbf{Y}, \mathbf{X})}[\log p_\theta(\mathbf{Y}|\mathbf{Z}, \mathbf{X})] + \text{KL}(q_\phi(\mathbf{Z}|\mathbf{Y}, \mathbf{X}) \parallel p_\theta(\mathbf{Z}|\mathbf{X})) + \min_k \|\hat{\mathbf{Y}}^{(k)} - \mathbf{Y}\|^2, \end{aligned} \quad (1)$$

where $q_\phi(\mathbf{Z}|\mathbf{Y}, \mathbf{X})$ is an approximate posterior parameterized by ϕ , $p_\theta(\mathbf{Z}|\mathbf{X})$ is a conditional Gaussian prior parameterized by θ , and $p_\theta(\mathbf{Y}|\mathbf{Z}, \mathbf{X})$ is a conditional likelihood modeling future trajectory \mathbf{Y} via the latent codes \mathbf{Z} and past trajectory \mathbf{X} . Additionally, $\mathcal{L}_{\text{diversity}} = \min_k \|\hat{\mathbf{Y}}^{(k)} - \mathbf{Y}\|^2$ is a diversity loss, which encourages the network to produce diverse samples. Given each past trajectory X , the model generates K sets of latent codes $\{\mathbf{Z}^{(1)}, \dots, \mathbf{Z}^{(k)}, \dots, \mathbf{Z}^{(K)}\}$ from the conditional Gaussian prior $\mathcal{N}(p_\theta^\mu(\mathbf{Z}|\mathbf{X}), p_\theta^\sigma(\mathbf{Z}|\mathbf{X}))$, where $\mathbf{Z}^{(k)} = \{z_1^k, \dots, z_n^k\}$, resulting in K future trajectories $\hat{\mathbf{Y}}^{(k)}$.

Similarly, in a conditional Generative Adversarial Net (cGAN)-based model (e.g., Social-GAN [1]), it uses a loss function as follows:

$$\begin{aligned} \mathcal{L}_{\text{total}} &= \mathcal{L}_{\text{gan}} + \mathcal{L}_{\text{diversity}} \\ &= \mathbb{E}_{\mathbf{Y} \sim p_{\text{data}}}[\log D_\theta(\mathbf{Y}|\mathbf{X})] + \mathbb{E}_{\mathbf{Z} \sim p_{\mathbf{Z}}}[\log(1 - D_\theta(G_\phi(\mathbf{Y}|\mathbf{X}, \mathbf{Z})))] + \min_k \|\hat{\mathbf{Y}}^{(k)} - \mathbf{Y}\|^2, \end{aligned} \quad (2)$$

where G represents the generator and D represents the discriminator. $\hat{\mathbf{Y}}^{(k)} = G(\mathbf{Y}|\mathbf{X}, \mathbf{Z}^{(k)})$ is one of the predicted trajectories in which $\mathbf{Z}^{(k)}$ is randomly sampled from $\mathcal{N}(0, 1)$. During the training, \mathcal{L}_{gan} is maximized to train D and $\mathcal{L}_{\text{total}}$ is minimized to train G .

Threat model. We follow the setup in prior work [11] and adopt an idealized threat model, where the adversary alters its observed history \mathbf{X} by adding a perturbation δ bounded by the adversarial set $\mathbb{S}_p^\epsilon = \{\delta \mid \|\delta\|_p \leq \epsilon\}$, where ϵ is the perturbation bound. The perturbation aims to mislead the prediction $\hat{\mathbf{Y}}$. A naïve adversarial attack is to find the perturbation through an adversarial objective $\delta = \arg \max_{\delta \in \mathbb{S}} \{\min_{k \in \{1, \dots, K\}} \|\hat{\mathbf{Y}}^{(k)}(\mathbf{X} + \delta) - \mathbf{Y}\|^2\}$, where $\hat{\mathbf{Y}}^{(k)}(\mathbf{X} + \delta)$ is the predicted trajectory conditioned on the random variable $\mathbf{Z}^{(k)}$ and adversarial history trajectory $\mathbf{X} + \delta$.

Naïve adversarial training. Adversarial training formulates a min-max game with an inner maximization process that optimizes the perturbation δ to generate adversarial examples for misleading the model at each training iteration, and an outer minimization process that optimizes the model parameters to make correct predictions for these examples. We follow the standard adversarial training formulation [8]:

$$\min_{\theta, \phi} \max_{\delta \in \mathbb{S}} \mathcal{L}_{\text{total}}(\mathbf{X} + \delta, \mathbf{Y}). \quad (3)$$

4 RobustTraj: Robust Trajectory Prediction

As stated earlier, applying adversarial training for trajectory prediction presents three critical challenges: (1) gradient obfuscation due to model stochasticity, (2) unstable learning due to changing adversarial distribution, and (3) performance loss in the benign situation. In this section, we describe each challenge in more detail and present the corresponding solutions in our *RobustTraj* method.

Improve adversarial generation with *Deterministic Attack*. Since trajectory prediction is inherently uncertain and there is no single correct answer, probabilistic generative models are usually used to cope with the stochastic nature of the trajectory prediction task. Such stochasticity will obfuscate the gradients that are used to generate effective adversarial examples in the inner maximization process of adversarial training. The naïve attack mentioned in section 3 is a straightforward way to achieve this goal. However, this optimization involves a stochastic sampling process $\mathbf{Z}^{(k)} \sim \mathcal{N}(p_\theta^\mu(\mathbf{Z}|\mathbf{X}), p_\theta^\sigma(\mathbf{Z}|\mathbf{X}))$. Such a stochastic process will obfuscate the gradients for finding the optimal adversarial perturbation δ , making the outer minimization (robust training) less effective. In order to sidestep such stochasticity, we propose the *deterministic attack* that creates a deterministic




Figure 2: Visualizations of the CVAE models trained with clean (a) data, *Salt and pepper* noise (b), and adversarial perturbations (c); Quantitative results of the correlation between the label of the generated images and conditioned images at different noise levels (d).

gradient path within the model to generate the adversarial perturbation. $\hat{\mathbf{Z}}$. Specifically, we use a *deterministic latent code* by replacing the sampling process $\mathbf{Z}^{(k)} \sim \mathcal{N}(p_{\theta}^{\mu}(\mathbf{Z}|\mathbf{X}), p_{\theta}^{\sigma}(\mathbf{Z}|\mathbf{X}))$, with the maximum-likelihood sample (here, i.e $\hat{\mathbf{Z}} = p_{\theta}^{\mu}(\mathbf{Z}|\mathbf{X})$). The objective for generating the adversarial perturbation is thus:

$$\delta = \arg \max_{\delta \in \mathbb{S}} \mathcal{L}_{\text{adv}}(\mathbf{X} + \delta, \mathbf{Y}) = \arg \max_{\delta \in \mathbb{S}} \| p_{\theta}(\mathbf{Y}|\hat{\mathbf{Z}}, \mathbf{X} + \delta) - \mathbf{Y} \|^2, \text{ where } \hat{\mathbf{Z}} = p_{\theta}^{\mu}(\mathbf{Z}|\mathbf{X} + \delta). \quad (4)$$

We empirically show that gradients from this deterministic gradient path can effectively guide the generation of adversarial examples. We name our attack as *Deterministic Attack*.

Stabilize adversarial training with bounded noise and hybrid objective. During the adversarial training process, the distribution of the perturbed input $\mathbf{X} + \delta$ coevolves with the training process as δ is calculated via an inner maximization process at each training iteration. Although δ is bounded by the adversarial set $\mathbb{S}_{\epsilon}^{\delta}$, the resulting latent condition variable $\mathbf{C} = f(\mathbf{X} + \delta)$ can be arbitrarily noisy since the Lipschitz constant of neural network layers (f) is not bounded during training (See *Lemma 1.* in Appendix A). Since $\mathbf{C} = f(\mathbf{X} + \delta)$ is noisy, it is a less informative signal compared to the deterministic signal X . Thus, modeling $p_{\theta}(\mathbf{Y}|\mathbf{X} + \delta, \mathbf{Z})$ becomes substantially harder. In an extreme case that $\mathbf{C} = f(\mathbf{X} + \delta)$ is super noisy and contains no information, the training process can degenerate to model $p_{\theta}(\mathbf{Y}|\mathbf{Z})$, resulting in the undesirable worse performance on the clean data.

To further validate the above hypothesis that it is hard to model $p_{\theta}(\mathbf{Y}|\mathbf{X} + \delta, \mathbf{Z})$ with a changing data distribution of $\mathbf{X} + \delta$, we conduct an additional experiment. For simplicity, we use MNIST [29] as the dataset. As shown in Fig. 2, we divide each digit image into four quadrants. We take the top-left quadrant as the condition \mathbf{X} and the remaining quadrants as the output \mathbf{Y} . We train a CVAE ($p_{\theta}(\mathbf{Y}|\mathbf{X} + \delta, \mathbf{Z})$) to model \mathbf{Y} by using clean data (\mathbf{X}) or noisy data ($\mathbf{X} + \delta$), where δ represents salt and pepper noise [30] or adversarial noise [8], resulting in Fig. 2 (a), (b), (c) respectively. The top-left region of each image in the first row is the conditional variables \mathbf{X} . The rest of rows are the generated images with different \mathbf{Z} . Each column in the same row uses the same \mathbf{Z} . We observe that the model trained on clean data successfully captures the conditional distribution (i.e., the generated image highly depends on \mathbf{X}) while the model trained with adversarial noise degenerates and ignores the condition (i.e., each row generates images of the same digit). This result shows that the conditional generative model fails to learn from \mathbf{X} . To provide a quantitative analysis, we measure the correlation between the label of the generated images and the label of their conditioned image quadrants, resulting in Fig 2 (d). More details on how to calculate the correlation are in the Appendix A. We observe that the correlation drops as the noise level increases for both adversarial noise and *salt and pepper* noise. Adversarial noise is more effective to degenerate the conditional generative model. Therefore, we conclude that (1) the noises in the conditional data lead to degenerated conditional generative model (i.e., from CVAE to VAE); (2) the level of degeneration depends on the noise levels.

Based on the analysis result, to better learn a robust trajectory prediction model, we need to bound $|f(\mathbf{X} + \delta) - f(\mathbf{X})|$ to reduce the noise level. Hence, we propose the following regularization loss \mathcal{L}_{reg} :

$$\mathcal{L}_{\text{reg}} = d(f(\mathbf{X} + \delta), f(\mathbf{X})), \quad (5)$$

where d is a distance function (e.g., we use L_2 norm as the distance metric).

In addition, because the clean data has a fixed distribution, simultaneously learning from the clean data during the adversarial training process anchors the conditional distribution on a stable clean data distribution. Specifically, we propose the following hybrid objective:

$$\mathcal{L}_{\text{clean}}(\mathbf{X}, \mathbf{Y}) = \mathcal{L}_{\text{total}}(\mathbf{X}, \mathbf{Y}), \quad (6)$$

where $\mathcal{L}_{\text{total}}$ could be the loss in Eq. 1 for CVAE-based model or Eq. 2 for cGAN-based model.

Protect benign performance using data augmentation. Adversarial training often leads to performance degradation on clean data [23]. However, trajectory prediction is a critical component for safety-critical AD systems and its performance degradation can result in severe consequences (e.g., collisions). Thus, it is important to balance the model performance on the clean and adversarial data when designing adversarial training algorithms.

To further improve the performance on clean and adversarial data, we need to address the overfitting problem of the min-max adversarial training [31]. Data augmentation is shown to be effective in addressing the problem in the image classification domain [32]. However, data augmentation in trajectory prediction is rarely studied and non-trivial. To design an effective augmentation algorithm, Rebuffi et al. [32] argues that the most important criterion is that the augmented data should be realistic and diverse. Thus, we design a dynamic-model based data augmentation strategy \mathbb{A} shown in the Appendix A. By using the augmentation, we can generate diverse, realistic multi-agent trajectories for each scene and construct \mathbb{D}_{aug} .

RobustTraj. In summary, our adversarial training strategy for trajectory prediction models is formulated as follows:

$$\begin{aligned} \delta &= \arg \max_{\delta \in \mathbb{S}} \mathcal{L}_{\text{adv}}(\mathbf{X} + \delta, \mathbf{Y}), \quad \text{where } \{\mathbf{X}, \mathbf{Y}\} \in \mathbb{D} \cup \mathbb{D}_{\text{aug}} \\ \theta, \phi &= \arg \min_{\theta, \phi} \mathcal{L}_{\text{total}}(\mathbf{X} + \delta, \mathbf{Y}) + \mathcal{L}_{\text{clean}}(\mathbf{X}, \mathbf{Y}) + \beta \cdot \mathcal{L}_{\text{reg}}(\mathbf{X}, \mathbf{X} + \delta), \end{aligned} \quad (7)$$

where $\mathbb{D}, \mathbb{D}_{\text{aug}}$ are the training data and augmented data; \mathcal{L}_{adv} is adversarial loss to generate effective adversarial examples in Eq. 4; $\mathcal{L}_{\text{total}}$ is the loss in Eq. 1 or Eq. 2 to train a robust model against adversarial examples; \mathcal{L}_{reg} and $\mathcal{L}_{\text{clean}}$ are loss shown in Eq. 5 and Eq. 6 to provide a stable signal for training. β is a hyper-parameter for adjusting the regularization.

5 Experiments and Results

5.1 Experimental setup

Dataset and models. We follow the setting in prior work [4, 3] and use the nuScenes dataset [33] for evaluation. For the trajectory prediction models, we select the representative conditional generative models based on CVAE (AgentFormer [4]) and cGAN (Social-GAN [1]). AgentFormer is a state-of-the-art model based on CVAE and Social-GAN is a classic model based on cGAN. We report the final results for all three models: AgentFormer (AF), mini-AgentFormer (mini-AF) and Social-GAN. More details are shown in the Appendix B.

Training details and hyperparameter choices. For the adversarial training, we choose a 2-step Projected Gradient Descent (PGD) attack for the inner maximization and choose $\beta = 0.1$. We train 50 epochs and 100 epochs for AgentFormer and Social-GAN respectively. For other hyperparameters during training, we follow the original settings for AgentFormer and Social-GAN. The details for choosing these hyperparameters can be found in the Appendix B. All experiments are done on the NVIDIA V100 GPU [34]. We consider various baselines, including naïve adversarial training (naïve AT) and four defenses proposed by Zhang *et al.* [11]: data augmentation with adversarial examples (DA), *train-time smoothing*, *test-time smoothing*, *DA + train-time smoothing* and *detection + test-time smoothing*.

Attack and evaluation metrics. For the adversarial attack, we choose a 20-step PGD attack (an ablation study on step convergence can be found in the Appendix B). Without loss of generality, we use L_{∞} as the attack threat model so that $\mathbb{S} = \{\delta \mid \|\delta\|_{\infty} \leq \epsilon\}$. We select $\epsilon = \{0.5, 1.0\}$ -meter, where the 1-meter deviation is the maximum change for a standard car without shifting to another lane [11]. We use four standard evaluation metrics for the nusenes prediction challenge [33]: average displacement error (ADE), final displacement error (FDE), off road rates (ORR), and miss rate (MR). We evaluate the model’s performance on both clean and adversarial data. For convenience, we use *ADE*, *FDE*, *ORR*, *MR* to represent the performance on the clean data and *Robust ADE*, *Robust FDE*, *Robust ORR*, *Robust MR* to represent the performance under attacks. We compute these metrics with the best of five predicted trajectory samples, i.e., $K = 5$.

Table 1: ADE and Robust ADE on different defense methods and models. The 1-st and 2-nd lowest errors are colored.

Model	mini-AF				AF				SGAN			
	ADE		Robust ADE		ADE		Robust ADE		ADE		Robust ADE	
Method	0.5	1.0	0.5	1.0	0.5	1.0	0.5	1.0	0.5	1.0	0.5	1.0
Clean	2.05	2.05	6.86	11.53	1.86	1.86	5.09	8.57	4.80	4.80	10.52	20.15
<i>Naïve AT</i> [8]	2.75	2.78	5.44	9.20	2.52	2.56	3.81	6.81	6.43	6.55	8.34	14.63
<i>DA</i> [11]	2.31	2.32	5.54	9.32	2.10	2.08	4.35	7.22	5.41	5.40	8.85	17.25
<i>Train-time Smoothing</i> [11]	3.14	3.07	5.67	9.31	2.11	2.13	4.19	6.79	5.50	5.47	8.74	16.51
<i>Test-time Smoothing</i> [11]	2.97	3.07	4.96	8.50	2.40	2.41	4.43	7.44	6.16	6.17	9.05	17.42
<i>DA + Train-time Smoothing</i> [11]	2.41	2.39	5.48	9.00	2.17	2.13	4.14	6.62	5.63	5.61	8.60	16.14
<i>Detection + Test Smoothing</i> [11]	2.31	2.28	5.91	9.85	2.08	2.03	4.45	7.59	5.35	5.37	9.28	17.39
<i>RobustTraj</i>	2.14	2.11	3.69	3.82	1.91	1.95	2.73	2.86	4.95	5.07	5.20	6.94

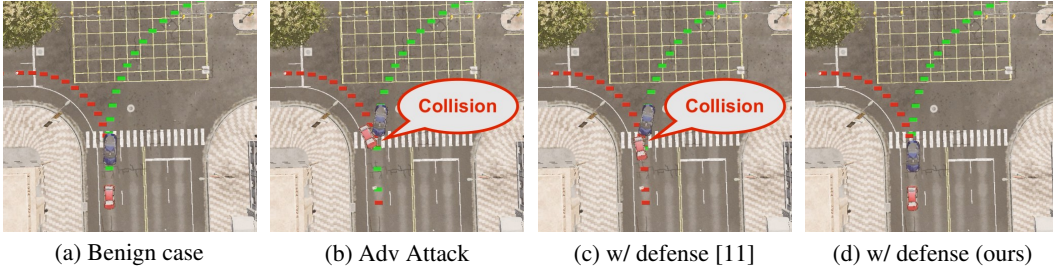


Figure 3: Impacts to a MPC-based downstream planner. (a) is under the benign case while (b), (c) and (d) are under the adversarial attacks. The blue car and the red car represent the AV and the adversarial agent respectively.

5.2 Main results

Here, we present our main results of *RobustTraj*. We compare it with the baselines including model trained with clean data (*Clean*) and naïve adversarial training (*Naïve AT*), and existing defense methods for trajectory prediction [11]. The results have been shown in Table 1.

We observe that our method achieves the best robustness and maintains good clean performance for most cases. For instance, with $\epsilon = 0.5$ attack on AgentFormer model, our method is able to reduce 46% prediction errors ($\frac{5.09-2.73}{5.09}$) under the attack at a cost of 2.6% ($\frac{1.91-1.86}{1.86}$) clean performance degradation on ADE, compared to the model trained with clean data at $\epsilon = 0.5$. Compared to the existing methods, our method also significantly outperforms in terms of the robustness. For instance, with $\epsilon = 1.0$ attack on AgentFormer model, our method achieves 45% better robustness with 9% better clean performance on ADE compared to the best results from existing methods [11].

Impacts to downstream planners. To further study the downstream impact of our robust trajectory model in the AD stack, we plug it into a planner. We select a MPC-based planner and evaluate the collision rates under the attack. To perform the attack on a closed-loop planner, we conduct attacks on a sequence of frames with the expectation over transformation (EOT) [35] method. We follow the setting from Zhang et al. [11] and choose $\epsilon = 1$. We choose AgentFormer model since it has the most competitive performance. As a result, we observe that, while AgentFormer model trained on clean data leads to 10 collision cases under attack, the robust trained model with the proposed *RobustTraj* is able to avoid all the collisions. As shown in Fig. 6, we demonstrate that the proposed *RobustTraj* is able to avoid the collisions (Fig. 6 (d)) while the *DA + Train-time Smoothing* method proposed by Zhang et al. [11] is not (Fig. 6 (c)).

5.3 Component analysis

In this section, we analyze the effectiveness of the three components. We use the mini-AgentFormer model since it has competitive performance and is lightweight for a fast adversarial training process.

Effectiveness of the *Deterministic Attack*. To demonstrate the importance of the *Deterministic Attack*, we compare it with competitive alternatives, *Latent Attack* and *Context Attack*, which also construct the deterministic path. However, they only attack a partial model as opposed to our end-to-

end full model attack. More details about these attacks are in the Appendix A. We evaluate their attack effectiveness by attacking a normally trained trajectory prediction model (without robust training). In Fig. 4, we demonstrate that *Deterministic Attack* is the most effective attack among all. Additionally, we embed them into the whole adversarial training pipeline and evaluate the adversarial robustness. The results are shown in Table 2. We observe that the model trained with the *Deterministic Attack* achieves the best robustness in terms of ADE. More results with other metrics and the other ϵ are in the Appendix B.

Effect of additional loss functions. We evaluate the performance of the models trained with additional loss terms: $\mathcal{L}_{\text{clean}}$ and \mathcal{L}_{reg} . In Table 2, we can see that the regularization term \mathcal{L}_{reg} improves robustness of the models and achieves better clean performance. It shows that the regularization of the introduced noises on conditional variables help the model to stabilize the training procedure. By adding the clean loss $\mathcal{L}_{\text{clean}}$, we observe that both the robustness and clean performance are improved further, which means the benign data indeed anchors the model output on clean data distribution and provides a stable signal for the better robust training for generative models.

Effect of domain-specific augmentation. To demonstrate the effectiveness of the domain-specific augmentation, We also combine it with all of the above components to validate its effect. The results are shown in Table 2. We observe that it achieves a better performance on clean and adversarial data.

6 Limitations

In this work, we identified the challenges of applying adversarial training on trajectory prediction models based on probabilistic generative models since they could cope with the natural uncertainty of motion forecasting. Though the probabilistic generative model is the main-stream architecture for the trajectory prediction task, there are other architectures (e.g., LSTM [1, 36], flow-based method [5, 6] and RL-based method [37]) for generating multi-modal predictions. Additionally, we only study the adversarial set with the threat model of \mathcal{L}_{∞} perturbation on trajectories instead of other types of threat models (e.g., optimization on the latent space [26], perturbation on raw sensor data [25]). Moreover, the primary goal of this paper is to study and improve the robustness of trajectory prediction models. To obtain salient and unambiguous insights, we minimize the conflating factors in our analysis without considering the perception model in our pipeline. We leave these as future work for building robust trajectory prediction models.

7 Conclusion

In this paper, we aim to study how to train robust generative trajectory prediction models against adversarial attacks, which is seldom explored in the literature. To achieve this goal, we first identify three key challenges in designing an adversarial training framework to train robust trajectory prediction models. To address them, we propose an adversarial training framework with three main components, including (1) a *deterministic attack* for the inner maximization process of the adversarial training, (2) additional regularization terms for stable outer minimization of adversarial training, and (3) a domain-specific augmentation strategy to achieve a better performance trade-off on clean and adversarial data. To show the generality of our method, we apply our approach to two trajectory




Figure 4: Performance of different attacks in mini-AgentFormer.

Table 2: ADE and robust ADE for different methods on mini-AgentFormer. The lowest error is in bold.

Method	ADE		Robust ADE	
	0.5	1.0	0.5	1.0
Clean	2.05	2.05	6.86	11.53
<i>Latent Attack</i>	2.55	2.70	4.10	4.71
<i>Context Attack</i>	2.47	2.59	3.94	4.78
<i>Deterministic Attack</i>	2.61	2.55	3.88	4.35
<i>Deterministic Attack</i> + \mathcal{L}_{reg}	2.29	2.31	3.76	4.28
+ $\mathcal{L}_{\text{clean}} + \mathcal{L}_{\text{reg}}$	2.23	2.19	3.71	3.83
+ $\mathcal{L}_{\text{clean}} + \mathcal{L}_{\text{reg}} + \text{Aug}$	2.14	2.11	3.69	3.82

prediction models, including (1) a CVAE-based model, AgentFormer, and (2) a cGAN-based model, Social-GAN. Our extensive experiments show our method could significantly improve the robustness with a slight performance degradation on the clean data, compared to the existing techniques and dramatically reduce the severe collision rates when plugged into the AD stack with a planner. We hope our work can shed light on developing robust trajectory prediction systems for AD.

References

- [1] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016.
- [2] Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2375–2384, 2019.
- [3] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *European Conference on Computer Vision*, pages 683–700. Springer, 2020.
- [4] Ye Yuan, Xinshuo Weng, Yanglan Ou, and Kris Kitani. Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [5] Nicholas Rhinehart, Kris M Kitani, and Paul Vernaza. R2p2: A reparameterized pushforward policy for diverse, precise generative path forecasting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 772–788, 2018.
- [6] Nicholas Rhinehart, Rowan McAllister, Kris Kitani, and Sergey Levine. Precog: Prediction conditioned on goals in visual multi-agent settings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2821–2830, 2019.
- [7] Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian Reid, Hamid Rezaatofghi, and Silvio Savarese. Social-bigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- [8] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [9] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [10] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.
- [11] Qingzhao Zhang, Shengtuo Hu, Jiachen Sun, Qi Alfred Chen, and Z Morley Mao. On adversarial robustness of trajectory prediction for autonomous vehicles. *arXiv preprint arXiv:2201.05057*, 2022.
- [12] Ahmadreza Jeddi, Mohammad Javad Shafiee, and Alexander Wong. A simple fine-tuning is all you need: Towards robust deep learning via adversarial fine-tuning. *arXiv preprint arXiv:2012.13628*, 2020.
- [13] Daniel Liu, Ronald Yu, and Hao Su. Extending adversarial attacks and defenses to deep 3d point cloud classifiers. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 2279–2283. IEEE, 2019.
- [14] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 582–597. IEEE, 2016.
- [15] Nicolas Papernot and Patrick McDaniel. Extending defensive distillation. *arXiv preprint arXiv:1705.05264*, 2017.
- [16] Ali Shafahi, Mahyar Najibi, Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! *arXiv preprint arXiv:1904.12843*, 2019.
- [17] Eric Wong, Leslie Rice, and J Zico Kolter. Fast is better than free: Revisiting adversarial training. *arXiv preprint arXiv:2001.03994*, 2020.

- [18] Cihang Xie, Mingxing Tan, Boqing Gong, Jiang Wang, Alan L Yuille, and Quoc V Le. Adversarial examples improve image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 819–828, 2020.
- [19] Cihang Xie and Alan Yuille. Intriguing properties of adversarial training at scale. In *International Conference on Learning Representations*, 2020.
- [20] Cihang Xie, Mingxing Tan, Boqing Gong, Alan Yuille, and Quoc V Le. Smooth adversarial training. *arXiv preprint arXiv:2006.14536*, 2020.
- [21] Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155*, 2017.
- [22] Yuzhe Yang, Guo Zhang, Dina Katabi, and Zhi Xu. Me-net: Towards effective adversarial robustness with matrix estimation. *arXiv preprint arXiv:1905.11971*, 2019.
- [23] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *International Conference on Machine Learning*, pages 7472–7482. PMLR, 2019.
- [24] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *International Conference on Machine Learning*, pages 274–283. PMLR, 2018.
- [25] Jingkan Wang, Ava Pun, James Tu, Sivabalan Manivasagam, Abbas Sadat, Sergio Casas, Mengye Ren, and Raquel Urtasun.
- [26] Davis Remppe, Jonah Philion, Leonidas J. Guibas, Sanja Fidler, and Or Litany. Generating useful accident-prone driving scenarios via a learned traffic prior. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [27] Saehyung Lee, Hyungyu Lee, and Sungroh Yoon. Adversarial vertex mixup: Toward better adversarially robust generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 272–281, 2020.
- [28] James Tu, Tsunhsuan Wang, Jingkan Wang, Sivabalan Manivasagam, Mengye Ren, and Raquel Urtasun. Adversarial attacks on multi-agent communication. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7768–7777, 2021.
- [29] Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- [30] Wikipedia contributors. Salt-and-pepper noise — Wikipedia, the free encyclopedia, 2022. [Online; accessed 16-June-2022].
- [31] Leslie Rice, Eric Wong, and Zico Kolter. Overfitting in adversarially robust deep learning. In *International Conference on Machine Learning*, pages 8093–8104. PMLR, 2020.
- [32] Sylvestre-Alvise Rebuffi, Sven Gowal, Dan A Calian, Florian Stimberg, Olivia Wiles, and Timothy Mann. Fixing data augmentation to improve adversarial robustness. *arXiv preprint arXiv:2103.01946*, 2021.
- [33] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- [34] Jack Choquette, Wishwesh Gandhi, Olivier Giroux, Nick Stam, and Ronny Krashinsky. Nvidia a100 tensor core gpu: Performance and innovation. *IEEE Micro*, 41(2):29–35, 2021.
- [35] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust physical-world attacks on deep learning visual classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1625–1634, 2018.
- [36] Hao Xue, Du Q Huynh, and Mark Reynolds. Ss-lstm: A hierarchical lstm model for pedestrian trajectory prediction. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1186–1194. IEEE, 2018.

- [37] Nachiket Deo and Mohan M Trivedi. Trajectory forecasts in unknown environments conditioned on grid-based plans. *arXiv preprint arXiv:2001.00735*, 2020.
- [38] Philip Polack, Florent Alché, Brigitte d’Andréa Novel, and Arnaud de La Fortelle. The kinematic bicycle model: A consistent model for planning feasible trajectories for autonomous vehicles? In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 812–818, 2017.
- [39] Matthew McNaughton, Chris Urmson, John M Dolan, and Jin-Woo Lee. Motion planning for autonomous driving with a conformal spatiotemporal lattice. In *2011 IEEE International Conference on Robotics and Automation*, pages 4889–4895. IEEE, 2011.
- [40] Eduardo F Camacho and Carlos Bordons Alba. *Model predictive control*. Springer science & business media, 2013.
- [41] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International Conference on Machine Learning*, pages 2206–2216. PMLR, 2020.

A Method and Implementations

A.1 Adversarial Attack on Trajectory Prediction

Latent Attack and Context Attack. Noticed that, besides *Deterministic Attack* introduced in the main paper, there are also two other less intuitive attacks. Since the prediction $\hat{\mathbf{Y}}$ is dependent on posterior distribution $q_\phi(\mathbf{Z}|\mathbf{X}, \mathbf{Y})$ and conditional variable $f(\mathbf{X})$, we can construct attacks based on that. *Latent attack* aims to increase the error of estimating $q_\phi(\mathbf{Z}|\mathbf{X}, \mathbf{Y})$, which is formulated as:

$$\delta = \arg \max_{\delta} \text{KL}(q_\phi(\mathbf{Z}|\mathbf{Y}, \mathbf{X}) \parallel q_\phi(\mathbf{Z}|\mathbf{Y}, \mathbf{X} + \delta)). \quad (8)$$

Context attack aims to increase the error of encoding the conditional variable $f(\mathbf{X})$, which is formulated as:

$$\delta = \arg \max_{\delta} d(f(\mathbf{X}), f(\mathbf{X} + \delta)), \quad (9)$$

where d is a distance function (e.g., L_2 norm). However, *latent attack* and *context attack* are effective due to two reasons. First, they are exploiting the vulnerability of a partial model. For example, *latent attack* only exploits the posterior estimation $q_\phi(\mathbf{Z}|\mathbf{X}, \mathbf{Y})$ and *context attack* only exploits the conditional encoder $f(\mathbf{X})$. Second, these attacks aim for a different goal. For the *latent attack* and *context attack*, the objectives are set for finding adversarial perturbations that maximize the difference of generated posterior distribution/context given \mathbf{X} and $\mathbf{X} + \delta$, due to lacking ground truth for intermediate latent variables. However, for the sample attack, the objective is directly set for maximizing the prediction errors from the ground truth (future trajectories), which is more effective.

Adversarial attack on consecutive frames. In order to fool a planner in a closed-loop manner to make consistent wrong decisions, we need to conduct adversarial attacks on consecutive frames. To attack L_p consecutive frames of predictions, we aim to generate the adversarial trajectory of length $H + L_p$ that uniformly misleads the prediction at each time frame. To achieve this goal, we can easily extend the formulation for attacking single-step predictions to attack a sequence of predictions, which is useful for attacking a sequence of decision made by AV planning module. Concretely, to generate the adversarial trajectories for L_p consecutive steps of predictions, we aggregate the adversarial losses over these frames. The objective for attacking a length of $H + L_p$ trajectory is:

$$\sum_{t \in [-L_p, \dots, 0]} \mathcal{L}_{\text{adv}}(\mathbf{X}_{\text{adv}}(t), \mathbf{Y}(t)), \quad (10)$$

where $\mathbf{X}_{\text{adv}}(t), \mathbf{Y}(t)$ are the corresponding $\mathbf{X} + \delta, \mathbf{Y}$ at time frame t .

A.2 Adversarial Training on Generative Models

Challenges. One challenge that hinders the adversarial training process is the noisy conditional data distribution disturbing the training process. One hypothesis we mentioned in the main paper is, the context encoding can magnify the bounded perturbation δ on history trajectory \mathbf{X} to an unbounded

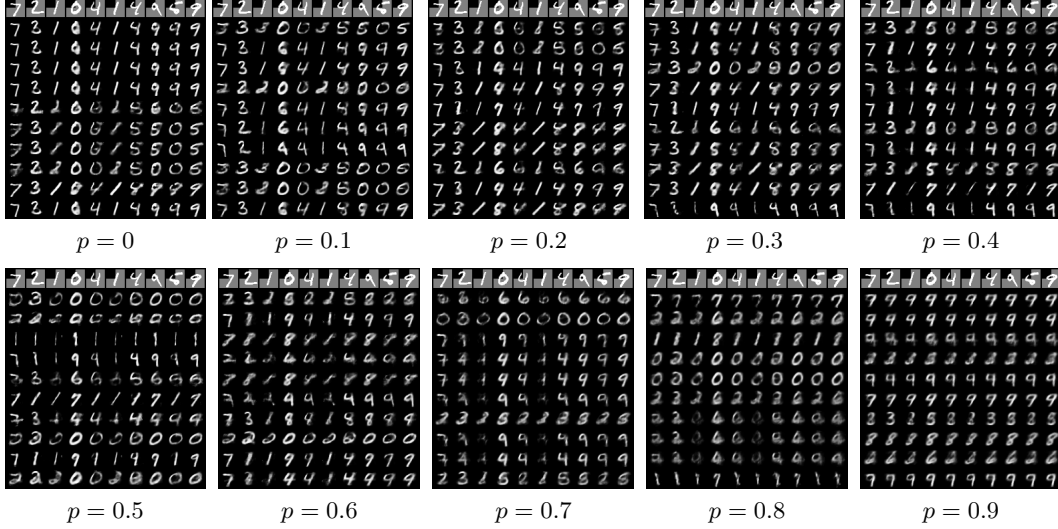


Figure 5: Visual examples of images generated from models trained with different levels of salt and pepper noises.

perturbation on the conditional variable $\mathbf{C} = f(\mathbf{X} + \delta)$, during the training process. For a neural network f which is not bounded on Lipschitz constant during the training procedure, given any constant η and an input \mathbf{X} , there exists a pair (δ, f) , that satisfies

$$\|f(\mathbf{X} + \delta) - f(\mathbf{X})\| \geq \eta.$$

Lemma A.2 can be easily derived by the definition of Lipschitz constant. Lipschitz constant L is defined as

$$L := \sup \frac{\|f(\mathbf{X} + \delta) - f(\mathbf{X})\|}{\|\delta\|}.$$

If L is not bounded, $\frac{\|f(\mathbf{X} + \delta) - f(\mathbf{X})\|}{\|\delta\|}$ is not bounded and so is $\|f(\mathbf{X} + \delta) - f(\mathbf{X})\|$ given a bounded δ . This means that, a bounded perturbation δ can potentially be magnified to be noisier on encoded conditional variable $\mathbf{C} = f(\mathbf{X} + \delta)$.

Analysis. In order to provide a quantitative analysis of the degeneration degree from conditional generative model to generative model (e.g., CVAE to VAE) with respect to the noise level, we propose a method to estimate the correlation between the degeneration and the noise level. Specifically, we trained a classifier with a 2-layer CNN achieving 99% accuracy on MNIST dataset. Then, given a conditional variable (the upper left quarter of an image of digit y), we generate images with the conditional generative models and use the classifier to calculate the confidence of the generated images labeled as digit y . We calculate the average confidence of 10 generated images on all 10,000 images in the MNIST test data set. The lower the score means the weaker the correlation between the generated images with the given conditions, or the stronger correlation between the degeneration and the noise level. We also provide visualization examples of generated images from models trained using data with different level of noises in Figure 5 and Figure 6. We can see that, with the noise level increases, the generated images are less dependent on the conditional variable (i.e., not being the same digit). In the extreme case of high level noise, for example when $p = 0.9$, the model generates images solely depends on the random prior value it samples and generates the similar images for each row. We can also see that, the adversarial noises are more effective compared to the salt and pepper noise. With a small amount of noise (i.e., $\epsilon = 0.1$), it can degenerate the conditional generative model to a generative model (i.e., CVAE to VAE here).

A.3 Data Augmentation with Dynamic Model

For the data augmentation strategy \mathbb{A} , we use a kinematic bicycle model [38] as our dynamic model to generate realistic trajectories that can be driven in the real world. Representing the behavior of actors as kinematic bicycle model trajectories allows for physical feasibility and fine-grained

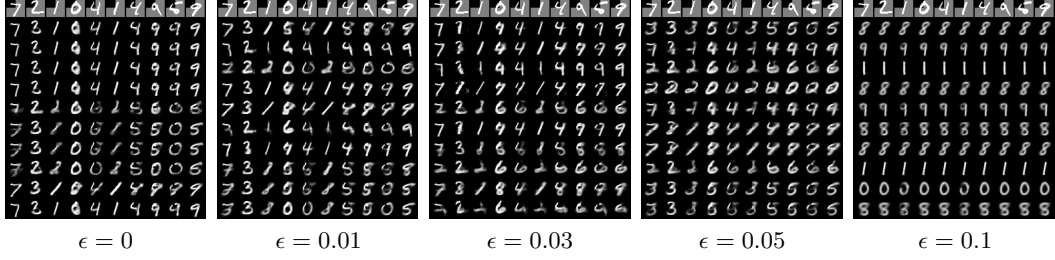


Figure 6: Visual examples of images generated from models trained with different levels of adversarial noises.

behavior control. To generate realistic trajectories, we first parameterize the trajectory $\mathbf{S} = \{s^t\}_0^T$ as a sequence of kinematic bicycle model states $s^t = \{p^t, \kappa^t, a^t\}$, where p represents the position, κ represents the trajectory curvature, and a represents the acceleration. Then, trajectories can be generated by controlling the change of curvature $\dot{\kappa}_t$ and the acceleration a_t over time, and using the kinematic bicycle model to update corresponding other states for each timestamp. To generate diverse trajectories, we set the objectives as biasing the trajectories to a given direction (e.g. forward, backward, left and right), while not colliding with other agents. To that end, we optimize a carefully-designed objective function \mathcal{L}_{dyn} over the control actions, i.e. $\dot{\kappa}_t$ and a_t for each agents. More specifically, the objective function consists of two components: $\mathcal{L}_{\text{dyn}} = \mathcal{L}_d + \gamma \mathcal{L}_{\text{col}}$, where \mathcal{L}_d is the deviation objective loss, \mathcal{L}_{col} is the collision regularization loss, and γ is a weight factor to balance the objectives. In each scene, we randomly pick a deviation objective loss \mathcal{L}_d from the set {moving forward, backward, left, right} for each agent. More specifically, the deviation objective loss \mathcal{L}_d is formulated as

$$\mathcal{L}_d = (\mathbf{X} - \mathbf{X}_{\text{aug}}) \bar{\mathbf{d}},$$

where \mathbf{X}_{aug} represents the generated trajectories by perturbing the trajectories in the dataset and $\bar{\mathbf{d}}$ represents the unit vectors for the target deviation directions in the set of {moving forward, backward, left, right}. And the collision regularization loss \mathcal{L}_{col} is formulated as

$$\mathcal{L}_{\text{col}}(\mathbf{X}_{\text{aug}}, \mathbf{X}) = \frac{1}{n-1} \sum_{i \neq \text{aug}}^{n-1} \frac{1}{\|\mathbf{X}_{\text{aug}} - \mathbf{X}_i\| + 1},$$

We also clip the maximum deviation of the positions so that the trajectories are constrained to be in the lane.

A.4 MPC-based Planner

Planner. In this work, to demonstrate the explicit consequences of the adversarial trajectory, we implement a simple yet effective planner that uses conformal lattice [39] for sampling paths and model predictive control (MPC) [40] for motion planning. We call this planner *MPC-based planner*.

Planning strategy. In this work, we consider a closed-loop planning strategy. Though for the closed-loop planning we have to replay the ground truth trajectories of other agents, we do notice reduced collisions and driving off-road consequences compared to open-loop planning and consider the closed-loop planning fashion meaningful.

B Experiment and Results

B.1 More details on Experimental Setup

Models. Since the adversarial training process is computationally heavy, we use a lightweight version of the AgentFormer in the analysis and ablation studies, namely mini-AgentFormer. In mini-AgentFormer, we (1) remove the map context and (2) reduce the transformer layer from two layers to a single layer. We report the final results for all three models: AgentFormer (AF), mini-AgentFormer (mini-AF) and Social-GAN.

Evaluating impacts to downstream planners. To demonstrate the impacts to downstream planners, we generate adversarial examples for consecutive frames on traffic scenarios in nuScenes dataset. With the MPC-based planner plugged in, we can demonstrate the consequences of the adversarial attacks on trajectory prediction models. We use the prediction results of AgentFormer trained with different methods due to its best performance among the three models. As we mentioned in the main paper, we show 10 cases where the AV collides with other vehicles under attack. We visualize 3 scenarios in the demo video of the supplementary material.

Hyperparameter choices. To select the hyperparameter β , we conduct adversarial training with different β for controlling the regularization. The results are shown in Table 3. We find that $\beta = 0.1$ achieves a good trade-off between robustness and clean performance. Therefore, we use $\beta = 0.1$ for the experiments in the rest of the paper.

β	0.01	0.1	0.5	1	10
ADE	2.19	2.29	2.37	2.39	2.57
Robust ADE	3.91	3.76	3.80	3.78	3.79

Table 3: Ablation study on different regularization loss weights.

To select the PGD attack step for evaluation, we conduct ablation experiments to show the convergence of different PGD steps. As shown in Figure 7, the attack converges at 20 steps. Thus, we select the 20-step PDG attack for the experiments in this paper




Figure 7: PGD step convergence for attack convergence with *Deterministic Attack*. Attack converges around 20 steps.




Figure 8: PGD step sizes ablation study. We find that except for 1 step PGD adversarial training, adversarial training with all the other step sizes achieves similar results.

To select the PGD attack steps of adversarial training, we conduct experiments on adversarially training the model with different PGD steps. Since we are using PGD attack with adaptive step sizes [41], attacks with any PGD steps are able to fully utilize the attack capability controlled by ϵ . In Figure 8, we show that except for the 1-step PGD attack, all other steps show the similar robustness and clean performance.

Evaluation with existing attack [11]. We also evaluate the robust trained model with the existing *search* attack [11]. The results are shown in the Table 4. We show that the existing attack increased less prediction error (e.g., 78% less for $\epsilon = 1.0$ on AgentFormer trained with clean data) due to the additional constraints of the attack. We demonstrate that the proposed *RobustTraj* achieves both best robustness and least clean performance degradation, compared to the baselines.

B.2 Main Results

We evaluate our methods and existing methods with four metrics in Table 5. We observe that the results are consistent where the proposed *RobustTraj* achieves the best results compared to the baselines and existing methods [11].

Table 4: Evaluation results of the proposed methods and existing methods on the *search* attack proposed by Zhang et al. [11]. mini-AF, AF and SGAN represent mini-AgentFormer, AgentFormer, and Social-GAN respectively. *DA* represents data augmentation with adversarial examples.

Model	Method	ADE		Robust ADE		FDE		Robust FDE	
		0.5	1.0	0.5	1.0	0.5	1.0	0.5	1.0
mini-AF	Clean	2.05	2.05	3.24	4.61	4.41	4.41	6.19	8.02
	<i>Naïve AT</i>	2.75	2.78	3.38	4.46	5.92	5.89	6.97	8.17
	<i>DA + Train-time Smoothing</i>	2.41	2.39	3.28	4.04	5.05	5.09	6.36	8.33
	<i>Detection + Test Smoothing</i>	2.31	2.28	3.26	4.41	4.96	4.91	6.41	8.27
	<i>RobustTraj</i>	2.14	2.11	2.50	2.51	4.36	4.35	5.07	5.11
AF	Clean	1.86	1.86	2.62	3.34	3.89	3.89	5.22	6.72
	<i>Naïve AT</i>	2.52	2.56	2.86	3.52	5.18	5.32	5.68	6.75
	<i>DA + Train-time Smoothing</i>	2.17	2.13	2.72	3.44	4.59	4.51	5.38	6.17
	<i>Detection + Test Smoothing</i>	2.08	2.03	2.58	3.29	4.43	4.26	5.49	6.22
	<i>RobustTraj</i>	1.91	1.95	2.14	2.21	4.02	4.01	4.31	4.31
SocialGAN	Clean	4.80	4.80	6.45	8.08	5.52	5.52	7.78	11.12
	<i>Naïve AT</i>	6.43	6.55	6.99	8.66	7.60	7.53	8.98	10.54
	<i>DA + Train-time Smoothing</i>	5.63	5.61	6.42	8.01	6.44	6.41	8.34	9.88
	<i>Detection + Test Smoothing</i>	5.35	5.37	6.34	8.72	6.12	6.07	7.77	10.21
	<i>RobustTraj</i>	4.95	5.07	5.01	5.49	5.72	5.73	6.68	6.40

In the demo video, we visualized scenarios where adversarial attacks on trajectory prediction models lead to collisions on both model trained on clean data and model trained with an existing defense [11], while model trained with *RobustTraj* is able to avoid the collisions.

Table 5: Additional evaluation results of the proposed methods and existing methods. mini-AF, AF and SGAN represent mini-AgentFormer, AgentFormer, and Social-GAN respectively. *DA* represents data augmentation with adversarial examples.

Model	Method	ADE		Robust ADE		FDE		Robust FDE	
		0.5	1.0	0.5	1.0	0.5	1.0	0.5	1.0
mini-AF	Clean	2.05	2.05	6.86	11.53	4.41	4.41	13.08	20.15
	<i>Naïve AT</i>	2.75	2.78	5.44	9.20	5.92	5.89	10.13	15.78
	<i>DA</i>	2.31	2.32	5.54	9.32	5.01	4.92	10.09	15.77
	<i>Train-time Smoothing</i>	3.14	3.07	5.67	9.31	6.77	6.61	10.51	17.48
	<i>Test-time Smoothing</i>	2.97	3.07	4.96	8.50	6.49	6.31	9.25	14.13
	<i>DA + Train-time Smoothing</i>	2.41	2.39	5.48	9.00	5.05	5.09	10.23	16.87
	<i>Detection + Test Smoothing</i>	2.31	2.28	5.91	9.85	4.96	4.91	11.49	17.57
	<i>RobustTraj</i>	2.14	2.11	3.69	3.82	4.36	4.35	7.10	7.59
AF	Clean	1.86	1.86	5.09	8.57	3.89	3.89	9.42	14.41
	<i>Naïve AT</i>	2.52	2.56	3.81	6.81	5.18	5.32	7.11	10.76
	<i>DA</i>	2.10	2.08	4.35	7.22	4.33	4.38	8.08	12.15
	<i>Train-time Smoothing</i>	2.11	2.13	4.19	6.79	4.40	4.46	8.01	11.13
	<i>Test-time Smoothing</i>	2.40	2.41	4.43	7.44	5.02	4.99	8.23	12.47
	<i>DA + Train-time Smoothing</i>	2.17	2.13	4.14	6.62	4.59	4.51	7.85	11.00
	<i>Detection + Test Smoothing</i>	2.08	2.03	4.45	7.59	4.43	4.26	8.01	12.74
	<i>RobustTraj</i>	1.91	1.95	2.73	2.86	4.02	4.01	5.22	5.48
SGAN	Clean	4.80	4.80	10.52	20.15	5.52	5.52	15.60	24.79
	<i>Naïve AT</i>	6.43	6.55	8.34	14.63	7.60	7.53	13.71	17.93
	<i>DA</i>	5.41	5.40	8.85	17.25	6.16	6.21	13.33	20.83
	<i>Train-time Smoothing</i>	5.50	5.47	8.74	16.51	6.27	6.31	14.03	19.48
	<i>Test-time Smoothing</i>	6.16	6.17	9.05	17.42	7.14	7.07	13.52	21.81
	<i>DA + Train-time Smoothing</i>	5.63	5.61	8.60	16.14	6.44	6.41	13.82	19.08
	<i>Detection + Test Smoothing</i>	5.35	5.37	9.28	17.39	6.12	6.07	13.36	21.59
	<i>RobustTraj</i>	4.95	5.07	5.20	6.94	5.72	5.73	8.97	8.89
Model	Method	MR		Robust MR		ORR		Robust ORR	
		0.5	1.0	0.5	1.0	0.5	1.0	0.5	1.0
mini-AF	Clean	0.33	0.33	0.77	0.93	0.08	0.08	0.28	0.45
	<i>Naïve AT</i>	0.45	0.45	0.60	0.70	0.11	0.10	0.22	0.34
	<i>DA</i>	0.37	0.38	0.61	0.72	0.09	0.09	0.22	0.35
	<i>Train-time Smoothing</i>	0.50	0.50	0.66	0.79	0.12	0.12	0.22	0.38
	<i>Test-time Smoothing</i>	0.48	0.50	0.56	0.65	0.11	0.11	0.20	0.32
	<i>DA + Train-time Smoothing</i>	0.39	0.40	0.65	0.77	0.09	0.09	0.22	0.37
	<i>Detection + Test-time Smoothing</i>	0.37	0.38	0.69	0.81	0.09	0.09	0.25	0.41
	<i>RobustTraj</i>	0.34	0.36	0.54	0.54	0.08	0.08	0.10	0.11
AF	Clean	0.29	0.29	0.66	0.88	0.04	0.04	0.16	0.30
	<i>Naïve AT</i>	0.39	0.38	0.51	0.69	0.06	0.06	0.13	0.22
	<i>DA</i>	0.32	0.32	0.56	0.74	0.05	0.05	0.14	0.26
	<i>Train-time Smoothing</i>	0.33	0.33	0.56	0.71	0.05	0.05	0.13	0.25
	<i>Test-time Smoothing</i>	0.37	0.37	0.58	0.77	0.05	0.05	0.14	0.26
	<i>DA + Train-time Smoothing</i>	0.33	0.33	0.54	0.70	0.05	0.05	0.13	0.24
	<i>Detection + Test-time Smoothing</i>	0.32	0.33	0.59	0.76	0.04	0.05	0.14	0.27
	<i>RobustTraj</i>	0.29	0.31	0.46	0.51	0.04	0.04	0.05	0.07
SocialGAN	Clean	0.40	0.40	0.85	0.99	0.14	0.14	0.52	0.60
	<i>Naïve AT</i>	0.53	0.53	0.63	0.77	0.19	0.19	0.39	0.44
	<i>DA</i>	0.44	0.44	0.72	0.85	0.16	0.16	0.44	0.51
	<i>Train-time Smoothing</i>	0.45	0.45	0.67	0.82	0.16	0.16	0.42	0.50
	<i>Test-time Smoothing</i>	0.51	0.51	0.74	0.85	0.19	0.19	0.45	0.52
	<i>DA + Train-time Smoothing</i>	0.47	0.46	0.66	0.81	0.17	0.17	0.41	0.49
	<i>Detection + Test-time Smoothing</i>	0.45	0.44	0.74	0.89	0.16	0.16	0.46	0.53
	<i>RobustTraj</i>	0.41	0.42	0.60	0.62	0.15	0.15	0.24	0.29