Certified Human Trajectory Prediction

Mohammadhossein Bahari*, Saeed Saadatnejad*, Amirhossein Askari Farsangi,
Seyed-Mohsen Moosavi-Dezfooli, Alexandre Alahi

CVPR 2025

Motivation

- “Uncertain” error rates with noisy inputs

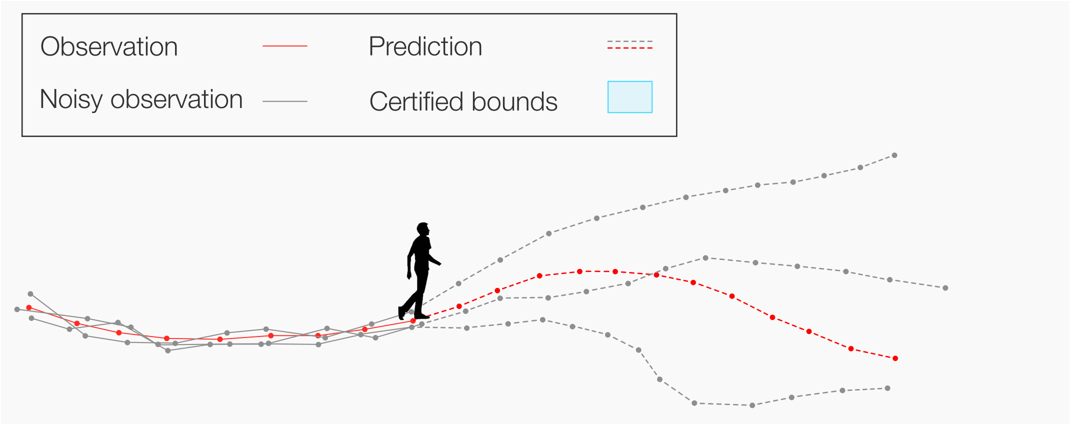

Motivation

- “Uncertain” error rates with noisy inputs
- A “Certified” trajectory prediction is required

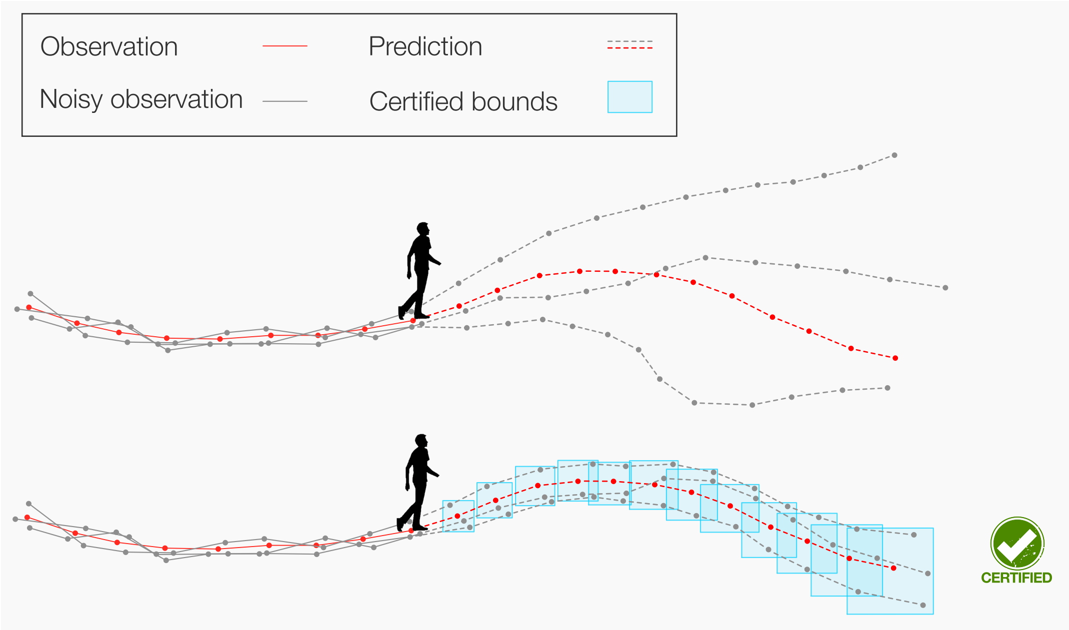

Related Work

- Trajectory prediction models are vulnerable to input noise [1,2].
- Existing robustness methods against input noise are heuristic and may fail on unseen data [3,4].
- We propose the first guaranteed robustness approach for trajectory prediction.

[1] Saadatnejad et al., Are socially-aware trajectory prediction models really socially-aware?, Transportation research part C, 2022
[2] Cao et al, Advdo: Realistic adversarial attacks for trajectory prediction, ECCV 2022
[3] Cao er al., Robust trajectory prediction against adversarial attacks, CoRL 2023
[4] Jiao et al., Semi-supervised semantics-guided adversarial training for robust trajectory prediction, ICCV 2023

Related Work

- **Randomised smoothing** guarantees that, under bounded input noise, model outputs remain within certified bounds.
- It has been previously applied in classification[1] and detection[2].
- We are the first to employ it for trajectory prediction by:
    - Adapting it to the multi-output regression task with domain-specific adjustments
    - Introducing a trajectory diffusion denoiser to mitigate performance drop
    - Proposing new certified evaluation metrics

[1] Salman et al., *Provably robust deep learning Via adversarially trained smoothed classifiers*, NeurIPS 2019
[2] Chiang et al., *Detection as regression: Certified object detection with median smoothing*, NeurIPS 2020

Certified Trajectory Prediction

Add n Gaussian perturbation to the input

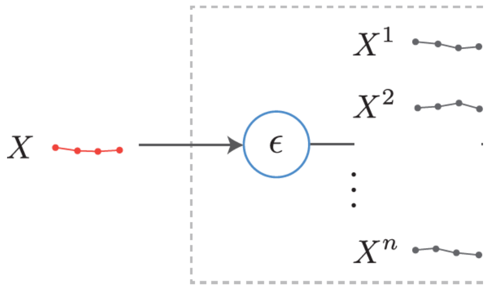

$$\epsilon \sim N(0, \sigma^2 I)$$

Certified Trajectory Prediction

Generate n predictions

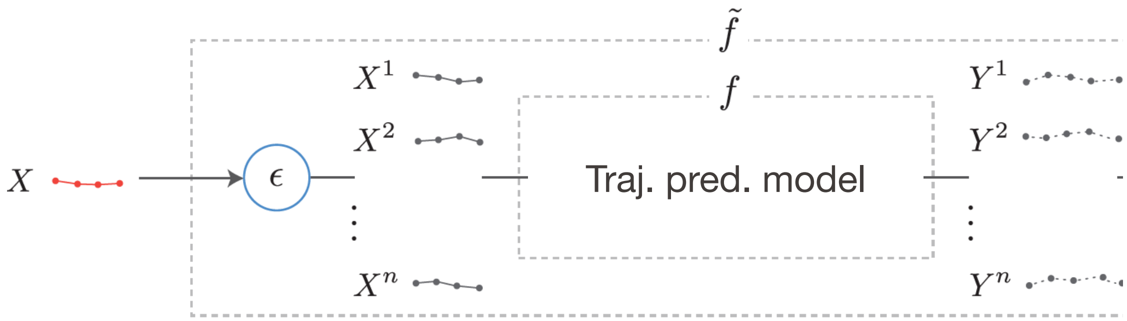

$$\epsilon \sim N(0, \sigma^2 I)$$

Certified Trajectory Prediction

Aggregate the predictions

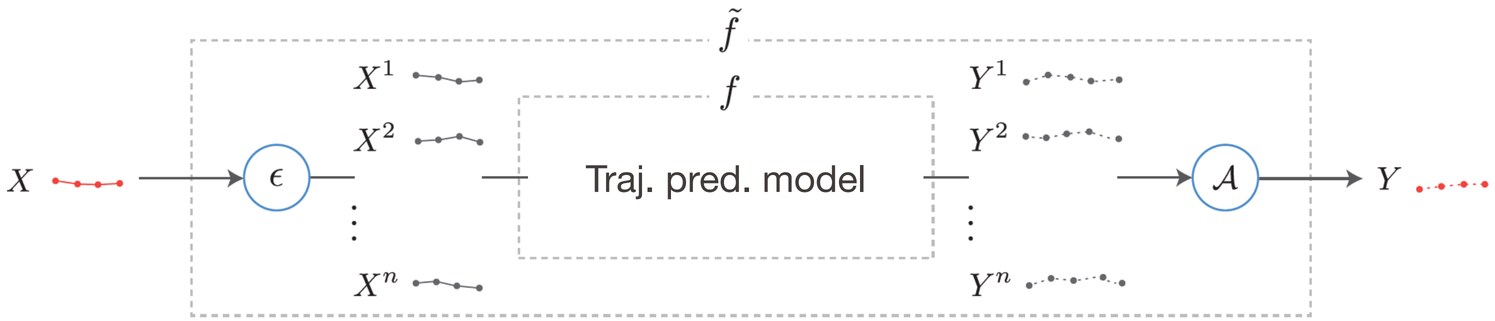

$$\epsilon \sim N(0, \sigma^2 I)$$

Certified Trajectory Prediction

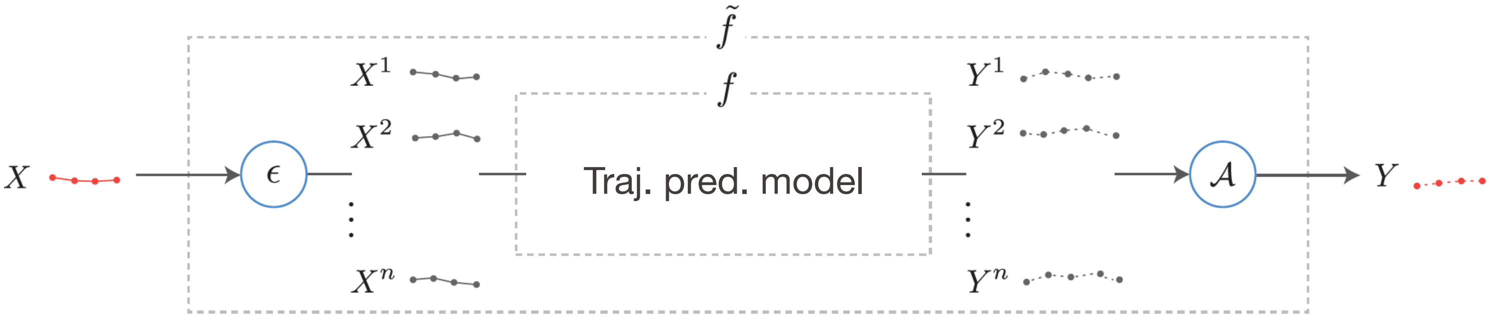

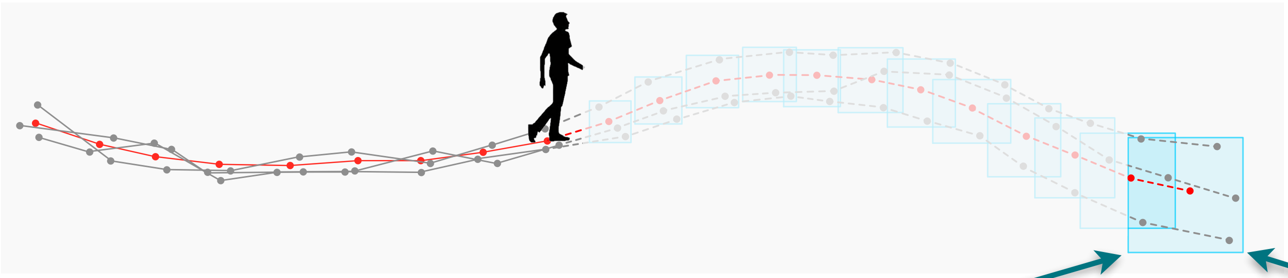

$$q_{\Phi(-\frac{R}{\sigma})}(X) \leq \tilde{f}(X+r) \leq q_{\Phi(\frac{R}{\sigma})}(X)$$

Diffusion Denoiser

- Problem: Randomized smoothing degrades performance due to injected perturbations
- Solution: an **unconditional denoising diffusion model** for trajectory denoising

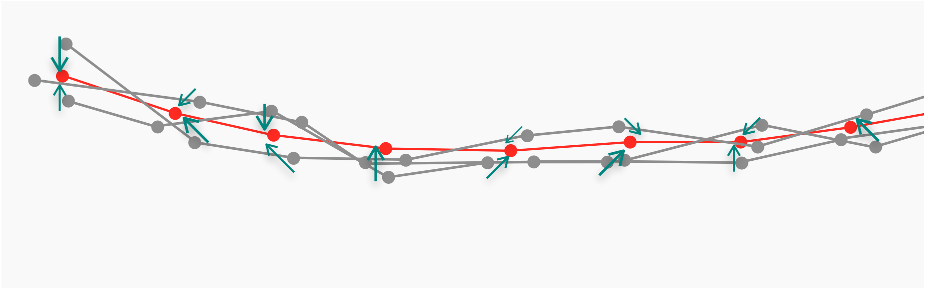

  - Training: Learns the trajectory distribution via diffusion, trained independently of the trajectory predictor
  - Inference: Given a noisy input trajectory, estimates the required denoising steps and then denoise to recover original trajectory

Diffusion Denoiser

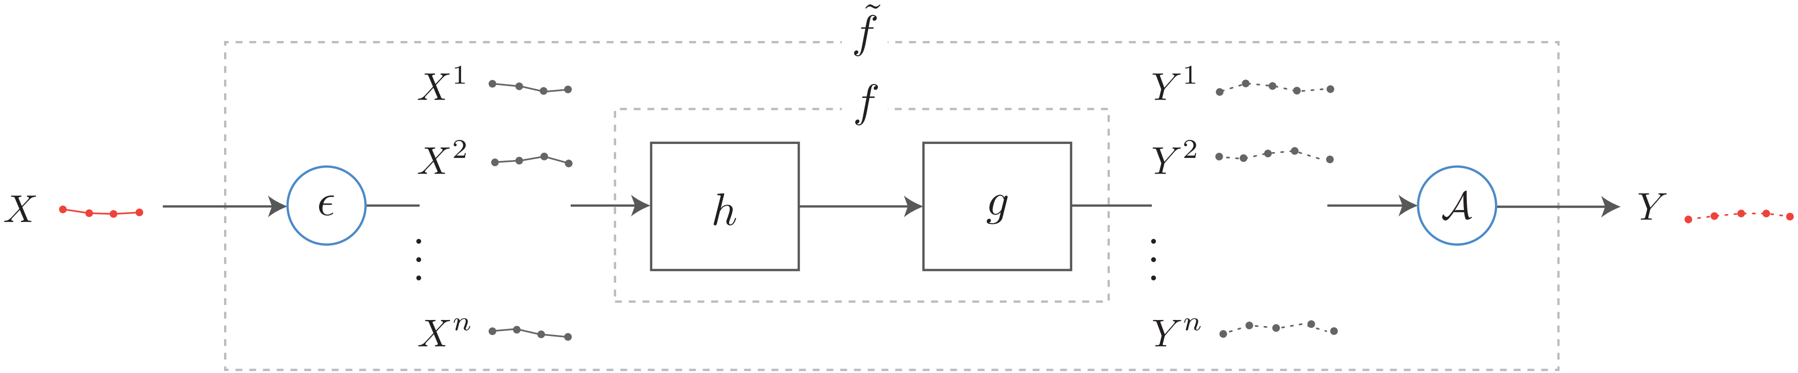

h : denoiser

g : original predictor

$\tilde{f}$: smoothed predictor

Certified Evaluation Metrics

- Common metric: Final Displacement Error (FDE)
- Our certified metrics:
  - Certified-FDE

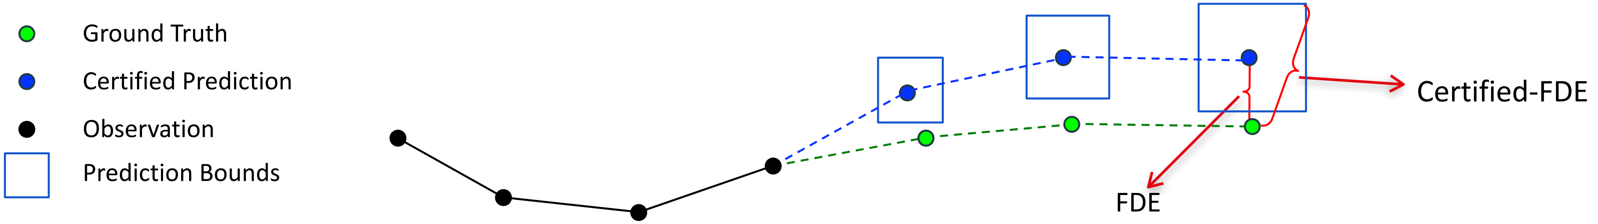

Certified Evaluation Metrics

- Common metric: Final Displacement Error (FDE)
- Our certified metrics:
  - Certified-FDE
  - Final Bound half-Diameter (FBD)

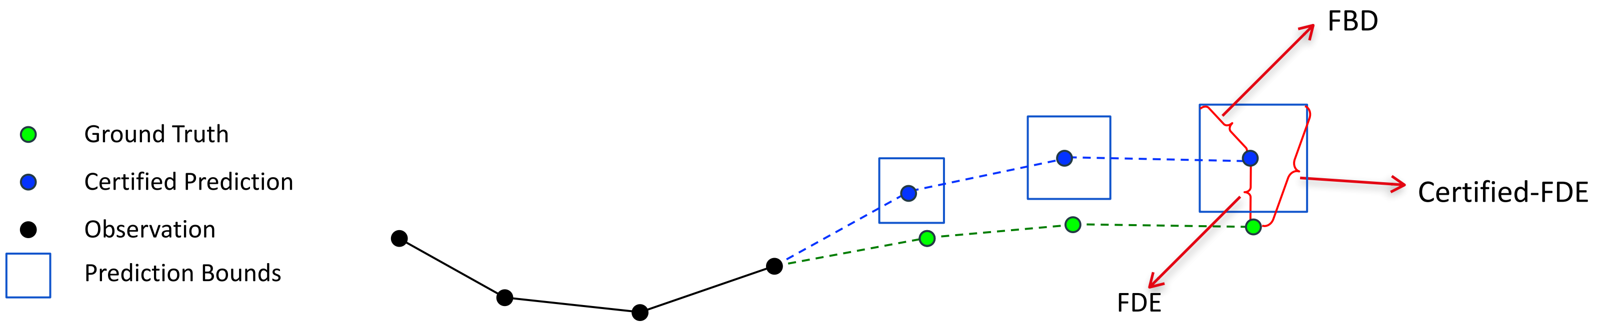

Quantitative Results

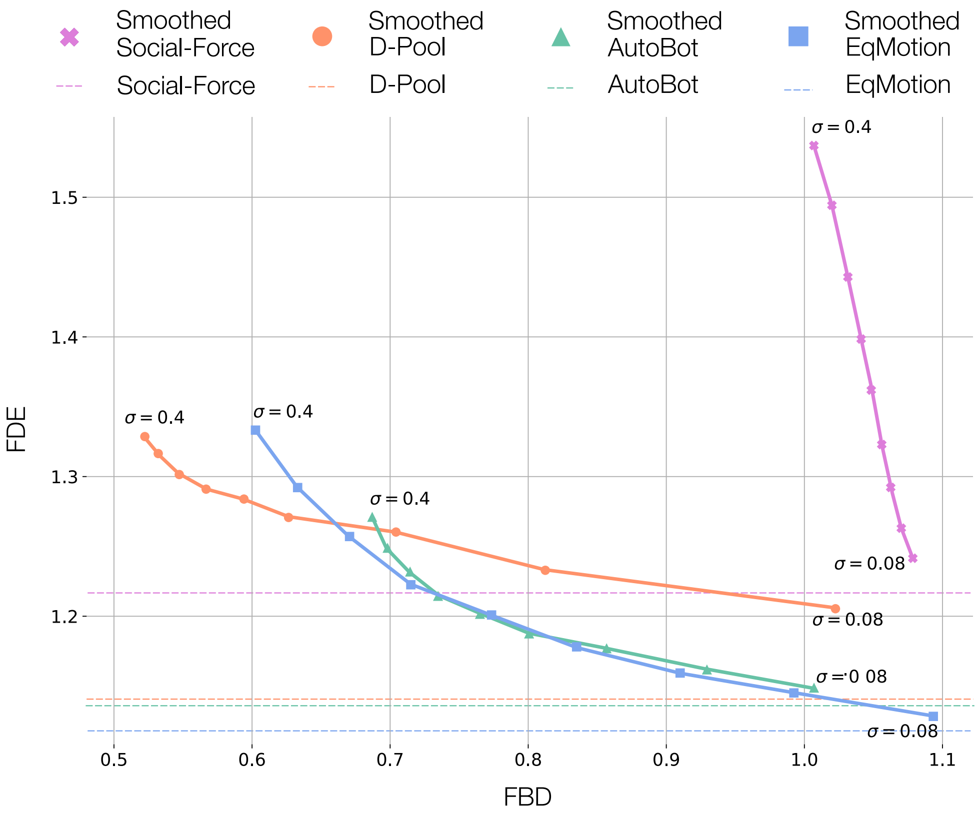

Qualitative Results

Original predictor

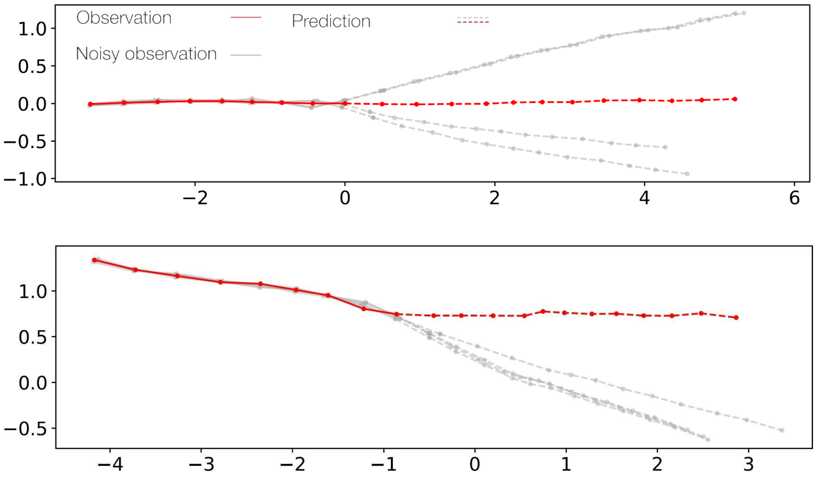

Smoothed predictor

CERTIFIED

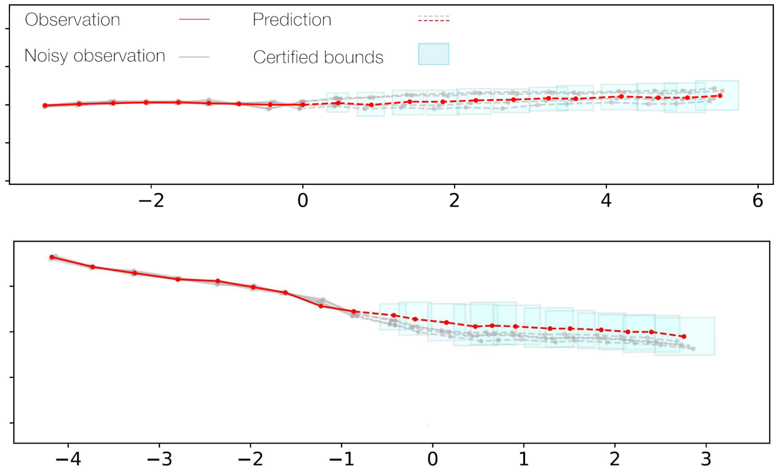

Results on Noisy and Adversarial Inputs

Original predictor

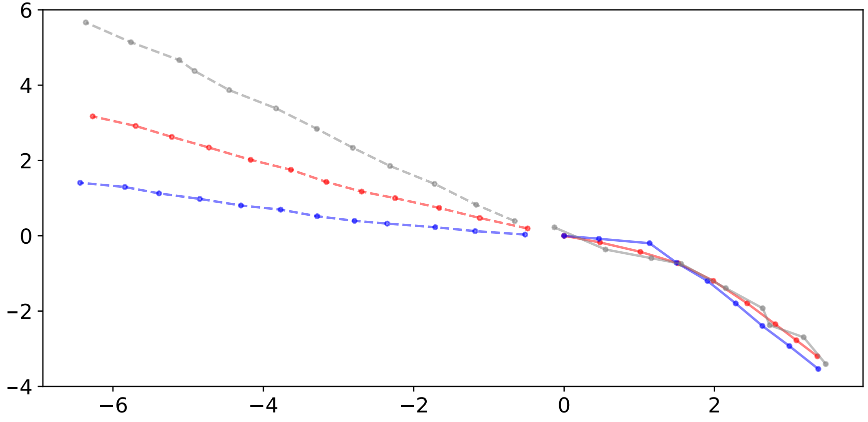

Smoothed predictor

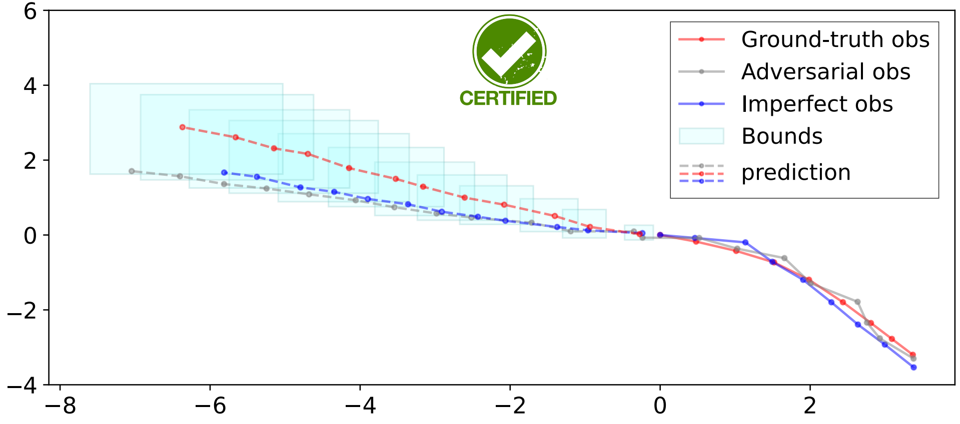

Summary

- **Problem:** Trajectory prediction models are vulnerable to input noise (e.g., sensor noise, random noise, adversarial perturbations)
- **Our contributions:**
  - **Certified Trajectory Prediction** guarantees output bounds under *any* input noise distribution
  - **Diffusion Denoiser** tightens certified bounds and improves predictive accuracy
  - **Certified Metrics** provide robustness evaluation for real-world deployment

s-attack.github.io