Conditional Flow-VAE for Safety-Critical Traffic Scenario Generation

Zimu Gong*, Brian Zhaoning Zhang*, Chris Zhang*, Kelvin Wong, Raquel Urtasun

ICRA IEEE INTERNATIONAL CONFERENCE ON ROBOTICS AND AUTOMATION 2026 VIENNA

Safety Critical Scenarios

Motivation: Safety-critical scenarios are

- Essential for developing and testing AVs
- Rare to find in the real world.

Can we learn to generate realistic safety critical scenarios at scale?

Existing Approaches

1. **Manual Specification**
   - Tedious, hard to scale
   - Requires expert domain knowledge
2. **Adversarial Optimization**
   - Real traffic participants are not inherently adversarial
   - Difficult to capture realistic, near-miss scenarios
3. **Curation**
   - Limited to number of real examples

Latent Flow Matching

Our method learns a **latent** space over scenarios and a **flow** transport to **map** the **nominal** distribution to the **safety critical** distribution

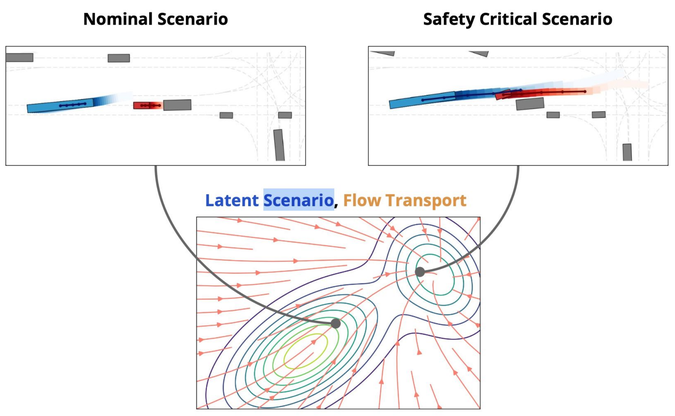

Our Approach: **FlowVAE**

We follow a **two-stage** training procedure. A VAE is trained on a **mixture of real and sim data** before a flow is learned on the **safety critical subset**

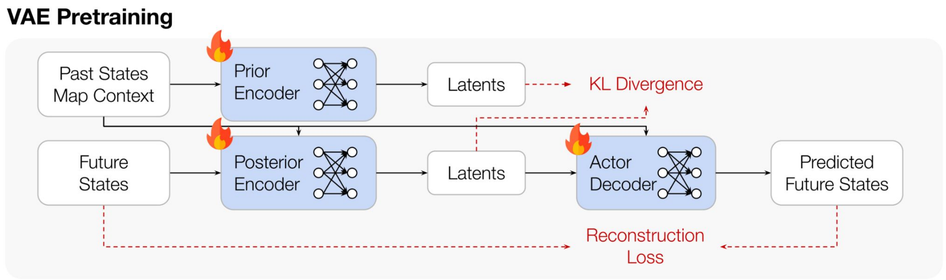

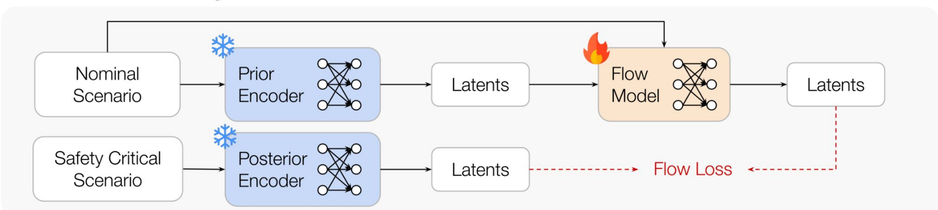

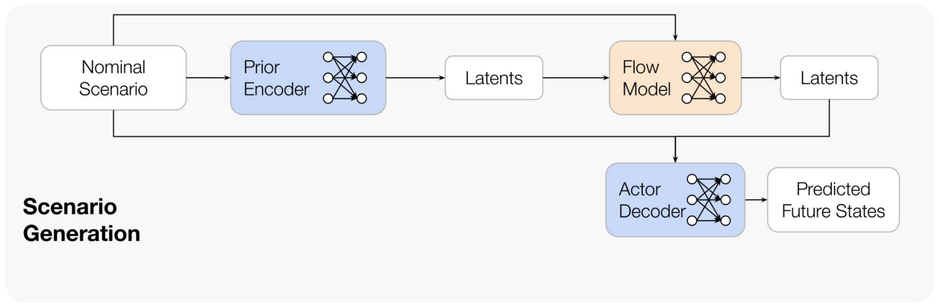

Realistic Safety Critical Generation

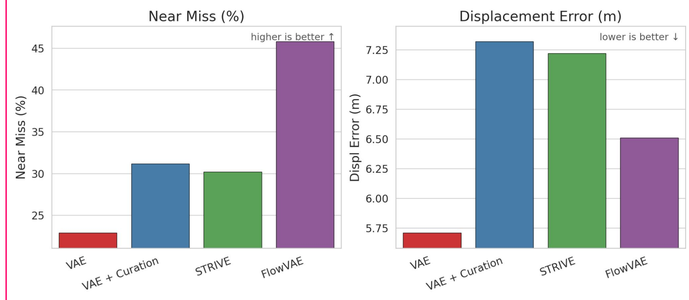

Compared curation or adversarial baselines, FlowVAE generates more near-miss scenarios, while staying true to the real safety-critical distribution

Simulation Data Transfer

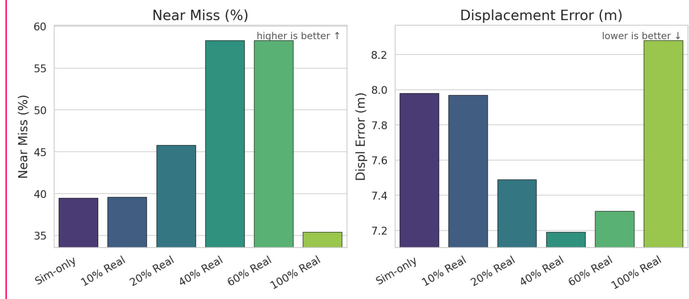

Introducing a moderate "goldilocks" amount of simulation data transfers to the real safety-critical distribution

Qualitative Comparison