GeRo: Generative Scenario Rollouts for End-to-End Autonomous Driving

Rajeev Yasarla[1] · Deepti Hegde[1] · Shizhong Han[1] · Hsin-Pai Cheng[1] · Yunxiao Shi[1] · Meysam Sadeghigooghari[2] · Shweta Mahajan[1]*
Apratim Bhattacharyya[1] · Litian Liu[1] · Risheek Garrepalli[1] · Thomas Svantesson[2] · Fatih Porikli[1] · Mohammad Ghavamzadeh[1] · Hong Cai[1]
[1] Qualcomm AI Research [2] Qualcomm Technologies, Inc.

NeurIPS 2026

Motivation

Driving logs show only one realized future. A planner must also anticipate how a maneuver unfolds and how nearby traffic may evolve over the next several steps.

Generative rollouts extend supervision across time: predicted ego and agent tokens expose future conflicts, while language grounds each action phase.

Safety feedback then teaches which generated futures support safer plans, beyond matching the logged trajectory.

Contributions & Novelty

1. Shared scenario token space

Joint ego plans, agent-motion forecasts, and ego-action Q&A connect visual context to language and behavior. Compact tokens enable future prediction without images.

2. Temporally grounded generation

Autoregressive rollout predicts successive scene states. Temporal consistency reduces drift; motion and planning losses keep generated futures grounded in behavior.

3. Safety and language feedback

GRPO ranks futures by collision avoidance, time-to-collision, and language alignment. The planner learns from future safety outcomes and action-grounded reasoning.

Inference-Time Rollouts

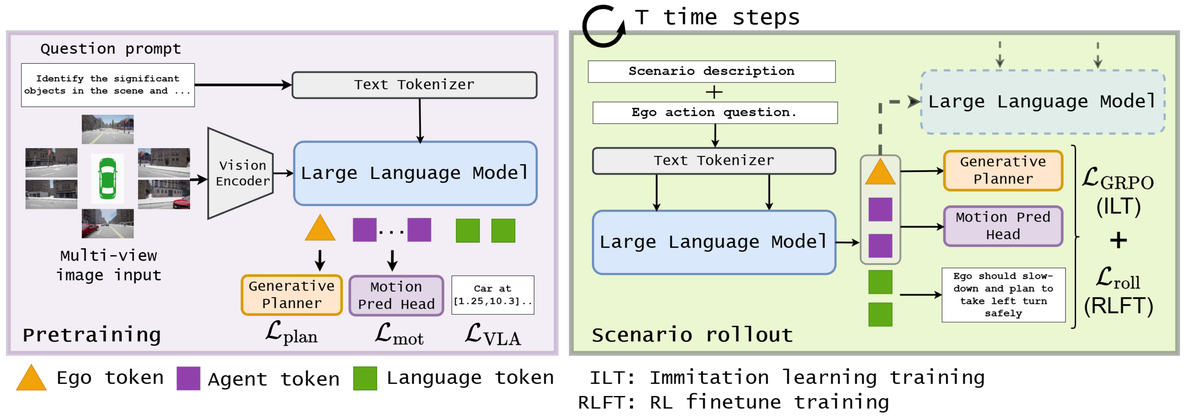
GeRo (ORION) · Bench2Drive

Optional lookahead improves driving at higher cost.

Rollout	DS ↑	SR ↑	s/frame ↓
Off (default)	81.9	60.1	0.306
On	**85.7**	**63.9**	**0.940**

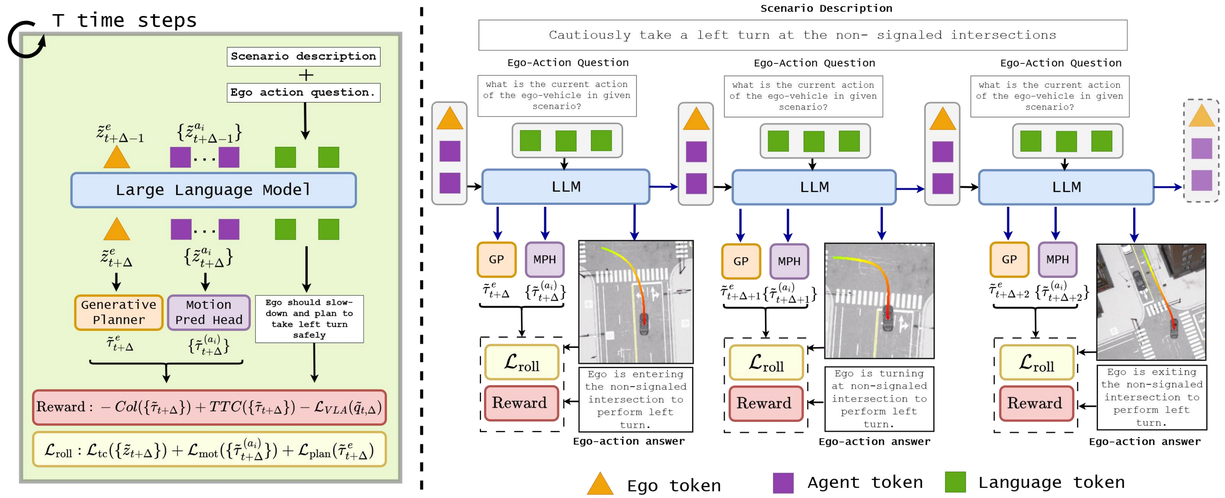
+3.8 DS · +3.8 SR points; 3.1x latency.

Proposed Method: GeRo Overview



Stage 1 learns shared tokens; Stage 2 generates ego plans, agent-motion forecasts, and action-grounded language.

Language-Conditioned Scenario Rollout



Closed-Loop Comparison

Bench2Drive · Navigation-command setting
DS: Driving Score; SR: Success Rate (%) ↑

Method / objective	DS ↑	SR ↑
UniAD-Base	45.81	16.36
VAD	42.35	15.00
GenAD	44.81	15.90
DriveTransformer-L	63.46	35.01
ReCogDrive	71.36	45.45
Qwen2.5	63.90	31.60
ORION	77.74	54.62
ORION (Qwen2.5)	76.88	53.90
GeRo (Qwen2.5)	**79.60**	**57.80**
GeRo (ORION)	**81.90**	**60.10**

+15.7 DS · +26.2 SR points
GeRo (Qwen2.5) vs. Qwen2.5

Ablation of Rollout Training

Cumulative objectives · Qwen2.5 backbone

Method / objective	DS ↑	SR ↑
Pretraining baseline	63.9	31.6
+ temporal consistency	67.8	40.1
+ ego planning + agent motion	74.4	50.4
+ collision + TTC feedback	78.5	55.7
+ ego-action Q&A + lang. reward	**79.6**	**57.8**

Each added component improves planning.
Final row also adds ego-action questions.

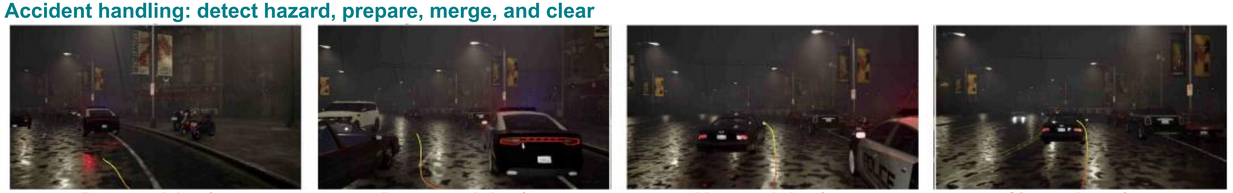
Qualitative Results: Language Grounding Across Driving Phases

STOP-sign intersection: slow down, yield, wait, and proceed



Approach · 2.8 m/s | Yield · 0.0 m/s | Wait · 0.5 m/s | Proceed · 5.4 m/s

Accident handling: detect hazard, prepare, merge, and clear



Detect · 1.0 m/s | Prepare · 2.1 m/s | Merge · 3.1 m/s | Clear · 4.7 m/s

* Work done at Qualcomm AI Research. Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.
ryasarla@qti.qualcomm.com