# Erik Zhou



## Head of Embodied-AI Infrastructure · Verifiable Data Loops & Bounded RSI

National High-Level Young Talent (QM) · Ph.D., Beijing Jiaotong University · Visiting Researcher, University of York · Industry Professor, Guangzhou University & HIT School of AI

eluanshi7@163.com · linkedin.com/in/embodied-ai-guo-zhou · guozhou-9yi.pages.dev

Shenzhen, China · Chinese / English / Swedish (basic) · Updated Nov 2026

## Profile

From formal verification of railway signalling, through data loops for L4 autonomous trucks, to data and R&D infrastructure for embodied AI, my work keeps answering one question: **how does a learning system know where it is wrong, and improve in a verifiable way?** Senior Director of AI Infrastructure at Pudu Robotics, where I built an embodied skill training and deployment platform from zero, with an RSI data flywheel (field feedback → failure attribution → targeted re-collection → incremental retraining) as the product spine. Previously led HONOR's embodied-AI data engine. I combine formal methods, LLMs / agents and platform engineering to turn frontier directions into accountable systems.

## Research Directions

### Bounded, auditable recursive self-improvement (RSI)

Improve the *learning process*, not just a model: a Learning Scientist diagnoses capability gaps and proposes falsifiable hypotheses, a Learning Harness executes interventions, and scientific memory keeps transferable learning principles. Digital search is fast, the physical world is true — held-out real-robot evaluation acts as a **reality anchor** against self-deception in simulation.

### Verifiable neuro-symbolic world models

Learn predicates, preconditions and effects from interaction data to form a symbolic transition world model, with PDDL, SAT/SMT and model checking as reasoning back-ends; counterexamples from planners and verifiers become supervision. Machine-checkable **skill contracts** (pre-condition, skill, post-condition) link high-level reasoning to VLA execution so failures can be localized.

### Training-ready embodied data engines

Graded ground truth (robot-measured vs. video-reconstructed actions); auto-labeling from 3D geometry, VLMs and visual estimation; ontologies and data contracts unify task, object, action and failure semantics; failures and counterexamples drive active collection.

### AI infrastructure for the R&D loop

A Research IR makes experiments compilable and checkable; evaluation as first-class infrastructure returns structured failure taxonomies; an evidence graph makes every improvement replayable, auditable and attributable.

## Experience

Jul 2026 – now
Shenzhen

### Pudu Robotics — Senior Director, AI Infrastructure · Head of Department

**Challenge:** adding one robot skill meant five stages — collection, QA, training, evaluation, deployment — scattered across tools and teams, gated by scarce ML engineers, hard to replicate, and frozen once deployed. **Goal:** "demonstrate a few times on site and the robot learns a new skill", then keeps improving.

- **Built the embodied skill platform from 0 to 1**: set product direction and acceptance criteria around full-chain lineage and the RSI flywheel; benchmarked leading embodied R&D platforms to define an integrated robot + collection + edge + field-feedback strategy; lead data, training and algorithm teams (10+).
- **End-to-end pipeline in 6 weeks**: on-robot recording auto-uploaded (offline-tolerant) → cloud QA into datasets → post-training → versioned models traceable to dataset, base model and code → simulation / offline / real-robot evaluation → OTA delivery with hot-swap of the inference model; OTA shipped one release ahead of plan.
- **Post-training speed-up**: found 95% of time in compute; sequence trimming, multi-camera batching, selective recomputation, parameter sharding and compile caching cut an 8-GPU 30k-step run from 24.3 h to 17.8 h (−27%) and peak memory from ~85 to ~35 GiB per GPU.
- **Platform agent**: conversational query, failure diagnosis and job launch over 44 platform APIs with confirm-before-write guardrails; 701 / 732 passes on a 244-question regression suite run 3×.
- Driving the next platform: **agentic data / training / evaluation infrastructure (RSI)**, symbolic world models and evaluation technology; authored the Reality-Anchored Recursive Learning research plan.

| 6 weeks | −27% | 8.5 min | 95.8% |
|---|---|---|---|
| kick-off to end-to-end pipeline | 8-GPU 30k-step post-training time | OTA delivery of a 12.4 GB model | agent regression pass rate |

Sep 2025 – Jun 2026
Shenzhen

**HONOR** Senior Expert & Head of Embodied-Robotics Data Engine Infrastructure

**Challenge:** heterogeneous data across embodiments, scarce and costly real-robot data, human egocentric video without action ground truth, and labels (3D + temporal + physical) far beyond 2D — manual labeling had become the bottleneck for model iteration.

- **Proposed and delivered an agent-based next-generation data engine** built on three engines — VLM auto-labeling, 3D spatial understanding (reconstruction, depth and 6D pose) and multimodal retrieval / mining — with agents orchestrating cleaning, labeling, mining, retrieval and synthesis across embodiments.
- Led a 10+ team to deliver high-accuracy sub-task auto-labeling, human hand pose estimation, and data pipelines for spatial understanding, CoT and memory; **connected pre-training, post-training and RL data pipelines** for cross-embodiment VLM / VLA training.
- Established an **ontology-based embodied data system** and, with algorithm and cloud teams, a closed loop from multi-source collection to cloud processing and training on an agent-centric platform.

Jul 2023 – Aug 2025
Guangzhou

**GAC R&D Center** Expert in Data Loop, LLMs, Simulation & Safety · Project Director

- Led a 10+ team building the autonomous-driving data-loop pipeline and scenario database; data-driven scenario mining and controllable generation (multimodal LLM fine-tuning, world-model and diffusion research) closing the loop between training and simulation; led SOTIF / ISO 3450X work and two provincial R&D programs.

Mar 2020 – Jun 2023
Stockholm

**Scania · Autonomize** Founder & CEO of an AI & data venture for L4 trucks

- Founded Autonomize inside Scania's Innovation Factory with an international team; led scenario-based testing and safety assessment (SOTIF, ISO 26262 / 3450X) and research on formal verification of neural-network robustness; Swedish national AD assessment, standards reviews, Lund University industry supervisor.

2018 – 2020 **Prover Technology AB**, Stockholm · Formal Methods Engineer: predicate-logic development of safety-critical signalling software; Solidity verification, AI-accelerated SAT

2015 – 2017 **Thales, Shanghai** · Formal Development Engineer: InterVer (NuSMV interlocking verification), CTSV data verification

2011 – 2013 **Beijing Jiaoda Signal Technology** · On-board ATP Engineer: 11 safety analyses (FTA, FMEA…), TÜV Rheinland certification

## Education

2005 – 2016 **Beijing Jiaotong University** · B.Eng., M.Eng., Ph.D. · Thesis: SCBM-based safety analysis — automated qualitative and quantitative results from CPN and PRISM probabilistic model checking

2013 – 2014 **University of York**, Computer Science · Visiting Researcher with Prof. John McDermid FREng, safety of systems of systems

## Recognition

**National High-Level Young Talent (QM)**

**High-Level Talent, Shenzhen** · Registered High-Level Talent, Guangzhou

**Industry Professor**, Guangzhou University & HIT School of AI

F5000 Top Paper 2016 · IEEE ITSC Student Scholarship 2015

## Academic Service

Workshops: **SRToITS 2023** (SAFECOMP) · **MASE 2020** · **FT4DAS 2019**

Reviewer of international AD standards; mentor, Volkswagen Group elite program

## Selected Publications 15 journal & conference papers · 9 Chinese invention patents

1. **G. Zhou**, Z. Wang, J. Chen, et al. Safety-Critical Interactive Scenario Generation for Autonomous Vehicles Testing Using Diffusion-Based Method. *IEEE Trans. Industrial Electronics*, 2026.
2. H. Yao, …, **G. Zhou**, H. Zhang. Language-Guided Safety-Critical Mixed-Traffic Scenario Generation for Closed-Loop Automated Driving Validation. *IEEE Trans. ITS*, 2026.
3. H. Yao, …, **G. Zhou**, H. Zhang. Learning-Based Realistic Traffic Scenario Simulation via Multiagent World Models With Intention-Aware Interactive Transformer. *IEEE/ASME Trans. Mechatronics*, 2026.
4. H. Yao, J. Zhu, **G. Zhou**, et al. InstanceDrive: Autonomous Driving Video Generation with Enhanced Instance-Level Inputs and Similarity Rewards. *IROS*, 2026.
5. **G. Zhou**, H. Zhao. Modeling and Quantitative Safety Analysis of Chinese Train Control System of Systems. *IEEE ITSC*, 2015.