

ZEYU DONG

(+1) 631-428-2858 | Zeyu.Dong@stonybrook.edu | [LinkedIn](#) | [Website](#)

About Me: Fifth-year Ph.D. at Stony Brook University, focusing on autonomous driving. Specialization in deep learning, computer vision, vision-language model, reinforcement learning, large-scale training, and real-time inferencing.

Education

Stony Brook University

Doctoral Program in Applied Math & Statistics (GPA: 3.9/4.0)

Stony Brook, NY
Sep. 2020 – Dec. 2025 (Expected)

Southern University of Science and Technology

Bachelor's Degree in Pure and Applied Mathematics (GPA: 3.8/4.0)

Shenzhen, China
Sep. 2016 – Jul. 2020

Experience

Machine Learning Intern @ Rippling | LLM, PyTorch, PySpark, Databricks

Jan. 2025 – Mar. 2025

Developed regression models to evaluate, predict, and explain employee performance using LLM embedding from code, PR review, and etc., improving the F1 score by 20%. [\[page\]](#)

Distributed System Research @ Sunrise Technology Inc. | InfiniBand, RDMA, RoCE v2, C/C++

Aug. 2024 – Dec. 2024

Designed and deployed an RDMA-based framework for 100Gbps zero-copy data transfer, integrating GPUDirect for real-time ML data processing, enabling over 30TB daily streaming at Brookhaven National Lab.

Autonomous Driving Research @ Sunrise Technology Inc. | Computer Vision, Transformer, ROS, C/C++

Jun. 2021 – Aug. 2024

Developed a platform and ML infrastructure for autonomous vehicles, integrating Transformer-GRU models for sensor fusion, achieving <50ms inference latency. [\[page\]](#)

Research Projects

Generalization of End-to-end Autonomous Driving with LLM | VLM, Real-time Inference

Mar. 2024 – Nov. 2024

Designed a hybrid architecture combining VLMs with end-to-end driving models, leveraging pre-trained VLMs for generalization, and achieving ~50% failure rate reduction in real-world deployments. [\[paper\]](#)

Training Models to Assist Legacy Devices | PyTorch, Computer Vision

Aug. 2024 – Nov. 2024

Developed Learning to Help, a hybrid framework jointly optimizing cloud and edge models, improving image classification by 20% with minimal server interaction. [\[paper, code\]](#)

Sim2Real for End-to-end Autonomous Driving | Highlight: PyTorch, CARLA Simulator, Computer Vision

Sep. 2023 – Mar. 2024

Developed a training approach to transfer expert driving knowledge from simulation to real-world, addressing data scarcity with transformers and domain-randomized pre-training, achieving ~60% failure rate reduction on unseen real-world tasks. [\[paper\]](#)

Intelligent Control for Electron Orbit Feedback System | Highlight: PyTorch, FPGA, Deep RL

Nov. 2021 – Aug. 2023

Applied Deep RL to control high-dimensional system at NSLS-II, developing model-based policy gradient algorithms for adaptive control, deploying on FPGA with <100ns latency, improving electron orbit stability by ~30%. [\[paper\]](#)

Publications

- **Z. Dong**, Y. Zhu, Y. Li, K. Mahon, and Y. Sun, “Generalizing end-to-end autonomous driving in real-world environments using zero-shot LLMs,” in 8th annual conference on robot learning (CoRL 2024).
- **Z. Dong**, Y. Zhu, K. Mahon, and Y. Sun, “Transformer-Based Domain Knowledge Transfer for End-to-End Autonomous Driving.”
- **Z. Dong**, Y. Tian, Y. Sun, “Adaptive Model-Based Reinforcement Learning for Orbit Feedback Control in NSLS-II Storage Ring,” in 3rd ICFA Beam Dynamics Mini-Workshop on Machine Learning Applications for Particle Accelerators.
- Y. Wu, Y. Li*, **Z. Dong*** (*equal contribution), N. Sathyavageeswaran, and A. D. Sarwate, “Learning to Help in Multi-Class Settings,” The Thirteenth International Conference on Learning Representations (ICLR 2025).
- Y. Li*, **Z. Dong*** (*equal contribution), E. Luo, Y. Wu, S. Wu, S. Han, “When to Trust Your Data: Enhancing Dyna-Style Model-Based Reinforcement Learning With Data Filter,” arXiv: arXiv:2410.12160.
- Y. Li, **Z. Dong**, and S. Han. “Bayes-Optimal, Robust, and Distributionally Robust Policies for Uncertain MDPs,” University of Illinois at Chicago. doi: 10.25417/uic.27138990.v1.

Skills

Programming: PyTorch, TensorFlow, NumPy, SciPy, Pandas, C/C++, SQL, PySpark, Java, MATLAB, R, etc.

Techniques: Linux, Distributed Computing, RDMA, Networking, ROS, EPICS, FPGA, LaTeX, etc.

Awards

Stony Brook Resources Access Project (RAP) Grant

Dec. 2022 and Dec. 2024

- Presentation at the 3rd ICFA Beam Dynamics Mini-Workshop on Machine Learning Applications for Particle Accelerators.
- Presentation at 8th annual conference on robot learning.