

Bowen Feng

Ph.D. Candidate in Robotics, Princeton University

bf2192@princeton.edu | Google Scholar | github.com/fengb2-coder | fengb2-coder.github.io

Research. I work on **generalist robot learning**: building robots that actively reason, learn, and act proactively in an open, changing world, anticipating the emergent behaviors of the environment and the agents around them. My work distills the reasoning of vision-language models into fast, steerable action policies, and builds latent world models that plan and explore under uncertainty and partial observability. Applications include language goal-conditioned planning and navigation, autonomous driving, mobile manipulation, and multi-agent interaction.

Education

Princeton University

Ph.D. Candidate, Mechanical and Aerospace Engineering · Advisor: Prof. Felix Heide

Princeton, NJ

Aug 2022 – present

Northwestern University

M.S. in Robotics

Evanston, IL

Sep 2020 – Dec 2021

Rose-Hulman Institute of Technology

B.S. in Mechanical Engineering, Minor in Mathematics

Terre Haute, IN

Sep 2016 – May 2020

Selected Research Publications

* / † equal contribution (co-first author), ‡ joint supervision

[1] Training-Free Adaptation of Generalist Robot Policies via Distributional Policy Steering.

Bowen Feng et al.

Under review, 2026.

[2] Training Generalist Navigation Policies via Action Self-Distillation.

Bowen Feng[†], K. Wang[†], Y. Yang[†], A. Joshi, C. Allen-Blanchette, D. Shah[‡], F. Heide[‡].

[Conference on Robot Learning \(CoRL\), 2026.](#)

[3] EpicWorldModel: Exploration-driven Planning with Latent World Models.

Bowen Feng^{*}, J. Ost^{*}, M. Mei, A. Majumdar, F. Heide.

[Advances in Neural Information Processing Systems \(NeurIPS\), 2026.](#)

[4] VERDI: VLM-Embedded Reasoning for Autonomous Driving.

Bowen Feng^{*}, Z. Mei^{*}, J. Ost, F. Ghilotti, B. Li, R. Girgis, A. Majumdar, F. Heide.

[arXiv:2505.15925, 2025.](#) (Princeton × Torc Robotics)

[5] Behavior-Inspired Neural Networks for Relational Inference.

Y. Yang^{*}, Bowen Feng^{*}, K. Wang, N. E. Leonard, A. B. Dieng, C. Allen-Blanchette.

[Int. Conf. on Artificial Intelligence and Statistics \(AISTATS\), 2025.](#) PMLR 258.

Research Experience

Computational Imaging Lab, Princeton University

Graduate Researcher, with Prof. Felix Heide

Princeton, NJ

2024 – present

Distributional Policy Steering (DPoS): training-free VLA adaptation

Under review, 2026

- Introduced a training-free, inference-time method that steers a frozen vision-language-action policy with VLM spatial reasoning: an MMD-based energy attracts sampled action chunks toward a reasoned target distribution while repelling candidates from one another for better coverage of the action space.
- Improves over the $\pi_{0.5}$ base policy by **+44.4 pp** on 9 real-world out-of-distribution tasks and **+19.9 pp** on LIBERO-PRO, the highest success rate among leading baselines, with no additional training.

Action Self-Distillation (ASD): language-steerable navigation

CoRL 2026

- Proposed a self-distillation framework in which a *self-teacher* VLM sees the ground-truth trajectory, infers the goal and writes perception/planning reasoning, and a *student* VLM learns to produce the same reasoning and actions from the raw camera view; the teacher tracks the student via EMA.

- Removes the need for human goal or reasoning annotations; trains on ~ 100 h of navigation data, versus 9,500 h for the OmniVLA baseline.
- Deployed **zero-shot on a real FrodoBot rover across 13 tasks**: 84% indoor success vs. 54% for OmniVLA (94% with a Qwen3.5-VL backbone) and 95% outdoor; +30 pp over the strongest baselines indoors and on multi-stage instructions.
- Outperforms supervised fine-tuning on frontier-model reasoning labels (80% vs. 36% success on ablation tasks).

EpicWorldModel: exploration-driven planning with latent world models

NeurIPS 2026

- Replaced deterministic JEPA world models with stochastic JEPAs trained by flow matching, predicting multiple plausible futures when goal-relevant content is occluded.
- Used flow predictive variance as an exploration signal inside CEM planning, trading off goal-reaching against exploring uncertain regions; up to 22% higher success than LeWorldModel.

VERDI: distilling VLM reasoning into modular end-to-end driving

Princeton $\times$ Torc Robotics · arXiv 2025

- Aligned the perception, prediction, and planning modules of an end-to-end driving stack with VLM text features at training time, embedding commonsense reasoning with *no* VLM cost at inference.
- Up to 11% better ℓ_2 and driving performance than e2e methods without embedded reasoning on nuScenes, Bench2Drive, and closed-loop HugSim, at real-time speed.

Princeton University

Graduate Researcher

Princeton, NJ

Aug 2022 – 2024

- Coupled graph neural networks with nonlinear opinion dynamics to infer interpretable, overlapping relations among interacting agents, improving long-horizon trajectory prediction (*AISTATS 2025*).
- Combined nonlinear opinion dynamics with control barrier functions for safe, efficient two-agent coordination at traffic intersections.

Interactive & Emergent Autonomy Lab, Northwestern University

Graduate Researcher, with Prof. Todd Murphey and Dr. Matthew Elwin

Evanston, IL

Mar 2021 – Mar 2022

- Used a robot arm to bend flexible wire into target shapes as tools, via optimal control and active learning over the wire's equilibrium-shape manifold.

Selected Projects

- **Deep Q-learning for urban driving (CARLA)**: end-to-end collision-free driving learned without hand-coded traffic or lane-following rules.
- **Neural Circuit Policies on a Clearpath Jackal**: CNN + NCP navigation (obstacle avoidance, line following) with far fewer control neurons and less data.
- **Mobile manipulation, KUKA youBot (simulation)**: trajectory planning, odometry, and feedforward + PID control for pick-and-place.
- **Automated 3D modeling**: Sawyer arm and TurtleBot3 cooperatively scan objects with a RealSense camera for CAD-free reconstruction.

Skills

Learning: robot foundation model post-training, test-time inference & steering, VLM fine-tuning & distillation (Qwen-VL), latent world models, reinforcement learning, imitation learning, PyTorch, Python

Robotics: real-robot deployment (FrodoBot rover, Clearpath Jackal, Sawyer, TurtleBot3), ROS, optimal control, control barrier functions, motion planning

Simulation & data: CARLA, HugSim, Bench2Drive, nuScenes