Yiyuan Pan

yiyuanp@andrew.cmu.edu · pyywill.github.io · github.com/PyyWill

Education

2025 – 2027 Expected
Carnegie Mellon University, Pennsylvania, USA
M.S. in Robotics · GPA 4.17/4 · Advisor: Changliu Liu

2021 – 2025
Shanghai Jiao Tong University, Shanghai, PRC
B.Eng. in Automation · GPA 3.92/4

Motivation

As learned end-to-end models become increasingly capable in planning and control, my research asks what structure robots should still represent explicitly. I study how to ground such structure in experience and compile it into task-aware representations, such as memory and constraints, to inform interpretable and compositional decision-making.

Selected Publications

* Equal contribution

RSS 2026 [1]
Emergent Neural Automaton Policies: Learning Symbolic Structure from Visuomotor Trajectories
Yiyuan Pan*, Xusheng Luo*, Hanjiang Hu, Peiqi Yu, Changliu Liu

NeurIPS 2025 [2]
Seeing through Uncertainty: Robust Task-Oriented Optimization in Visual Navigation
Yiyuan Pan, Yunzhe Xu, Zhe Liu, Hesheng Wang

NeurIPS 2025 [3]
Wonder Wins Ways: Curiosity-Driven Exploration through Multi-Agent Contextual Calibration
Yiyuan Pan, Zhe Liu, Hesheng Wang

AAAI 2025 Oral [4]
Planning from Imagination: Episodic Simulation and Episodic Memory for Vision-and-Language Navigation
Yiyuan Pan, Yunzhe Xu, Zhe Liu, Hesheng Wang

CoRL 2026 [5]
Inference-Time Robot Behavior Steering through Physically-Aware Reconfiguration of Task-Structure
Yiyuan Pan, Hanjiang Hu, Shangtao Li, Xusheng Luo, Changliu Liu

Under Review [6]
What Makes a Good Symbol for Hybrid Discrete–Continuous Planning?
Yiyuan Pan, Changliu Liu, Jiayuan Mao

AAAI 2025 Oral [7]
FLAME: Learning to Navigate with Multimodal LLM in Urban Environments
Yunzhe Xu, **Yiyuan Pan**, Zhe Liu, Hesheng Wang

TPAMI 2026 [8]
Dream to Recall: Imagination-Guided Experience Retrieval for Memory-Persistent Vision-and-Language Navigation
Yunzhe Xu, **Yiyuan Pan**, Zhe Liu

IROS 2026 [9]
Learning Safe-Stoppability Monitors for Humanoid Robots
Yifan Sun, **Yiyuan Pan**, Shangtao Li, Caiwu Ding, Tao Cui, Lingyun Wang, Changliu Liu

ICLR 2026 [10]
Seeing, Listening, Remembering, and Reasoning: A Multimodal Agent with Long-Term Memory
Lin Long*, Yichen He*, Wentao Ye, **Yiyuan Pan**, Yuan Lin, Hang Li, Junbo Zhao, Wei Li

arXiv 2026 [11]
VLESA: Vision-Language Embodied Safety Agent for Human Activity Monitoring
Hanjiang Hu*, **Yiyuan Pan***, Jiaxing Li, Xusheng Luo, Alexander Robey, Na Li, Yebin Wang, Changliu Liu

Experience

(i) Contextualizing Decisions: What to Remember

Sep 2023 – May 2025
Shanghai Jiao Tong University *Advisor: Hesheng Wang*

- Adapted a multimodal language model for urban navigation using cross-attention over recent observations and three-stage tuning on scene descriptions, route summaries, and actions. **[7]**

- Fused observations with recursively imagined scenes in topological memory **[4]**; queried environmental and behavioral memories across episodes with language-conditioned world-model predictions **[8]**.
- Encoded interaction histories as graphs of inferred peer intentions to calibrate intrinsic rewards, separating informative novelty from environmental stochasticity in multi-agent exploration. **[3]**

Mar 2025 – Jul 2025 — **ByteDance Seed** *Advisor: Hang Li*

- Designed entity-centric multimodal memory that incrementally links episodic events and semantic knowledge via persistent identities, supporting cross-modal retrieval and multi-turn reasoning. **[10]**

(ii) Assessing Decisions: What to Follow

Jun 2024 – Dec 2024 — **California Institute of Technology** *Advisor: Steven Low*

- Developed a learning-to-optimize framework that maps visual predictions into calibrated convex uncertainty sets for robust navigation, using planning feedback to align network training with downstream decision quality rather than prediction accuracy alone. **[2]**†

Sep 2025 – Present — **Carnegie Mellon University** *Advisor: Changliu Liu*

- Learned a neural monitor estimating safe-stopping probability under a fixed fallback policy, using importance sampling near the safety boundary to reduce false-safe predictions on humanoids. **[9]**
- Coupled intent–action prediction from egocentric video with a goal-conditioned safety filter trained through reinforcement learning, evaluating anticipated human actions against context-dependent safety requirements to support proactive intervention. **[11]**

(iii) Organizing Decisions: What to Symbolize

Sep 2025 – Present — **Carnegie Mellon University** *Advisor: Changliu Liu*

- Learned state-machine task models from visuomotor demonstrations through adaptive clustering and automata inference, without symbolic annotations; coupled their action priors with a residual policy for phase-conditioned action generation. **[1]**
- Composed user-preference automata with learned task structure and recomputed action priors through demonstration replay, grounding symbolic edits in demonstrated behaviors while steering a frozen residual policy at inference time. **[5]**

Mar 2026 – Present — **University of Pennsylvania** *Advisor: Jiayuan Mao*

- Formalized solver-sufficient symbols that expose discrete commitments and leave efficiently solvable subproblems; developed counterexample-guided refinement to discover compact symbol sets while preserving optimality in structured hybrid models. **[6]**
- Extended symbol discovery to neural energy landscapes using sampled optimization trajectories to identify barriers and refine partitions, enabling region-wise optimization and behavior steering without retraining. *[under review]*

Technical Skills

Programming	Python, C/C++, MATLAB
Machine Learning	PyTorch, Transformers, OpenAI / Azure API, W&B
Robotics	Kinova Gen3, Unitree G1, Galaxea R1 Lite · MuJoCo, Isaac Lab, Habitat-Sim
Formal Methods	Convex / combinatorial optimization, automata theory, temporal logic

† The framework began as my Caltech SURF report on power systems and was later instantiated as a robotics paper at SJTU.