


Yunuo Zhang

+1 (615) 877 5371
✉ yunuo.zhang@vanderbilt.edu
yz1056yz.github.io
YZ1056YZ
in [yunuo-zhang-497450152](#)
G [Google Scholar](#)



Education

- 2022–Present **Ph.D. in Computer Science**, *Vanderbilt University*, Nashville, TN, US
Advisor: Prof. Abhishek Dubey
- 2018–2021 **B.S. in Computer Science**, *Vanderbilt University*, Nashville, TN, US
Double Major: Computer Science, Mathematics

Research Interests

- **Sequential Decision-Making Under Partial Observability:** Advanced algorithms for POMDPs including particle-based belief representation, annealed importance sampling, and Monte Carlo tree search with applications to real-time planning under uncertainty.
- **Adaptive Planning in Non-Stationary Environments:** Online learning and planning methods that adapt to changing environmental dynamics, including in-context learning with temporal abstractions and policy-augmented search for evolving systems.
- **AI for Time-Critical Applications:** Deployment of decision-theoretic algorithms to safety-critical domains including UAV search and rescue operations, smart grid vehicle-to-building systems, and emergency response coordination.

Work Experience

- 2023–Present **DARPA ANSR Program**, *Graduate Research Assistant*, Vanderbilt University
- Developed "Shrinking POMCP" algorithm that reduces computational complexity by dynamically guiding UAV search toward non-sparse belief regions, achieving significant improvements over standard MCTS in time-constrained urban search and rescue scenarios. Integrated AirSim-ROS2 simulator with probabilistic world model and neuro-symbolic navigator for real-time path planning under partial observability, enabling safe navigation in complex 3D environments [\[paper\]](#)
 - Applied neuro-symbolic AI techniques combining neural perception with symbolic reasoning for UAV search and rescue, contributing to *Neuro-symbolic AI: Foundations and Applications* (Wiley-IEEE Press, 2025) [\[book\]](#)
 - Extending framework to multi-agent pursuit-evasion problems in urban environments with dynamic obstacles and adversarial targets (ongoing)
- 2023–2024 **Nissan Motor Corporation**, *Research Collaborator*
- Designed DG-MCTS algorithm for vehicle-to-building (V2B) energy optimization using real-world EV data [\[paper\]](#)
 - Optimized charging schedules reducing peak demand charges while meeting user requirements [\[paper\]](#)
 - Deployed solution for heterogeneous charger configurations and time-of-use pricing [\[paper\]](#)
- 2022–Present **Belief Representation for Partial Observability**, *Graduate Researcher*, Vanderbilt University
- Developed ESCORT algorithm using Stein-variational methods with correlation-aware projections for multi-modal belief distributions [\[paper\]](#)
 - Created AIROAS framework combining annealed importance resampling with observation adaptation for POMDPs [\[paper\]](#)
 - Advanced in-context planning (I-TAP) using latent temporal abstractions for high-dimensional continuous control (*under review, ICLR 2026*)

- 2022–Present **Adaptive Planning for Non-Stationary Systems**, *Graduate Researcher*, Vanderbilt University
- Created NS-Gym, first standardized toolkit for non-stationary MDPs with comprehensive benchmark environments [\[paper\]](#) [\[code\]](#)
 - Developed ADA-MCTS algorithm that balances pessimistic and optimistic planning based on epistemic and aleatoric uncertainty in changing environments [\[paper\]](#) [\[code\]](#)
 - Designed PA-MCTS that combines stale policies with online search to maintain performance when environmental dynamics change [\[paper\]](#) [\[code\]](#)

Publications

My research advances sequential decision-making under uncertainty through both theoretical contributions and practical deployments. I develop algorithms for fundamental challenges in partial observability and non-stationary environments, demonstrating their effectiveness in time-critical applications from autonomous search and rescue to smart grid optimization, while creating open-source benchmarks and toolkits that enable reproducible research in the community:

Algorithmic Innovations in Sequential Decision-Making

- [1] Yunuo Zhang, Baiting Luo, Ayan Mukhopadhyay, Gabor Karsai, and Abhishek Dubey. “ESCORT: Efficient Stein-variational and Sliced Consistency-Optimized Temporal Belief Representation for POMDPs”. In: *Proceedings of the Thirty-Ninth Annual Conference on Neural Information Processing Systems (NeurIPS)*. 2025.
 - Summary: Develops particle-based belief representation using Stein-variational methods with correlation-aware projections to accurately model multi-modal distributions in POMDPs.
- [2] Yunuo Zhang, Baiting Luo, Ayan Mukhopadhyay, and Abhishek Dubey. “Observation Adaptation via Annealed Importance Resampling for Partially Observable Markov Decision Processes”. In: *Proceedings of the 35th International Conference on Automated Planning and Scheduling (ICAPS)*. accepted as oral presentation. AAAI Press, 2025. URL: <https://arxiv.org/abs/2503.19302>.
 - Summary: Proposes annealed importance resampling to address particle degeneracy in online POMDP solvers when observations are highly informative, improving belief approximation accuracy.
- [3] Baiting Luo, Yunuo Zhang, Nathaniel S. Keplinger, Samir Gupta, Abhishek Dubey, and Ayan Mukhopadhyay. “In-Context Planning with Latent Temporal Abstractions”. Under review at the Fourteenth International Conference on Learning Representations (ICLR 2026). 2025.
 - Summary: Unifies in-context adaptation and online planning in learned latent temporal-abstraction space, enabling efficient decision-making under stochastic dynamics and partial observability.
- [4] Baiting Luo, Yunuo Zhang, Abhishek Dubey, and Ayan Mukhopadhyay. “Act as You Learn: Adaptive Decision-Making in Non-Stationary Markov Decision Processes”. In: *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*. AAMAS '24. Auckland, New Zealand: International Foundation for Autonomous Agents and Multiagent Systems, 2024, pp. 1301–1309. ISBN: 9798400704864.
 - Summary: Introduces ADA-MCTS algorithm that adaptively balances pessimistic and optimistic planning based on epistemic and aleatoric uncertainty estimates in non-stationary environments.
- [5] Ava Pettet, Yunuo Zhang, Baiting Luo, Kyle Wray, Hendrik Baier, Aron Laszka, Abhishek Dubey, and Ayan Mukhopadhyay. “Decision Making in Non-Stationary Environments with Policy-Augmented Search”. In: *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*. AAMAS '24. Auckland, New Zealand: International Foundation for Autonomous Agents and Multiagent Systems, 2024, pp. 2417–2419. ISBN: 9798400704864.
 - Summary: Introduces PA-MCTS algorithm that combines stale policy knowledge with online Monte Carlo tree search to handle non-stationary environments where relearning optimal policies is computationally expensive.

Real-World Applications & Benchmarks

- [6] Nathaniel S. Keplinger, Baiting Luo, Iliyas Bektas, Yunuo Zhang, Kyle Hollins Wray, Aron Laszka, Abhishek Dubey, and Ayan Mukhopadhyay. “NS-Gym: Open-Source Simulation Environments and Benchmarks for Non-Stationary Markov Decision Processes”. In: *Proceedings of the Thirty-Ninth Annual Conference on*

Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track. 2025. URL: <https://arxiv.org/abs/2501.09646>.

- Summary: Presents first standardized toolkit for non-stationary MDPs integrated with Gymnasium framework, enabling reproducible evaluation of algorithms under changing environmental conditions.

- [7] Rishav Sen, Yunuo Zhang, Fangqi Liu, Jose Paolo Talusan, Ava Pettet, Yoshinori Suzue, Ayan Mukhopadhyay, and Abhishek Dubey. "Online Decision-Making Under Uncertainty for Vehicle-to-Building Systems". In: *Proceedings of the ACM/IEEE 16th International Conference on Cyber-Physical Systems (ICCPS)*. ICCPS '25. California, USA: Association for Computing Machinery, 2025.

- Summary: Models vehicle-to-building energy optimization as an MDP and uses domain-guided Monte Carlo tree search to handle heterogeneous chargers and demand charges over long planning horizons.

- [8] Alvaro Velasquez, Houbing Song, Pradeep Ravikumar, S. Shankar Sastry, and Sandeep Neema, eds. Wiley-IEEE Press, 2025, p. 512. ISBN: 978-1-394-30237-6. URL: <https://www.wiley.com/en-us/Neuro-symbolic+AI%5C%3A+Foundations+and+Applications-p-9781394302376>.

- Summary: Contributed Chapter 4 on neuro-symbolic approaches for UAV navigation and search planning in urban rescue operations.

- [9] Yunuo Zhang, Baiting Luo, Ayan Mukhopadhyay, Daniel Stojcsics, Daniel Elenius, Anirban Roy, Susmit Jha, Miklos Maroti, Xenofon Koutsoukos, Gabor Karsai, and Abhishek Dubey. "Shrinking POMCP: A Framework for Real-Time UAV Search and Rescue". In: *2024 International Conference on Assured Autonomy (ICAA)*. 2024, pp. 48–57. DOI: [10.1109/ICAA64256.2024.00016](https://doi.org/10.1109/ICAA64256.2024.00016).

- Summary: Presents Shrinking POMCP that reduces computational complexity in UAV search and rescue by dynamically guiding exploration toward non-sparse belief regions.

Teaching and Service

Professional Service

2023 **Student Volunteer**, *SMARTCOMP 2023*

IEEE International Conference on Smart Computing

Teaching Experience

2022 Fall **Teaching Assistant**, *Database Management Systems*, Vanderbilt University

2021 Fall **Teaching Assistant**, *Algorithms*, Vanderbilt University

2020 Fall **Teaching Assistant**, *Discrete Structures*, Vanderbilt University

Core Coursework

Reinforcement Learning, Statistical Foundations of Deep Learning, Foundations of Machine Learning, Advanced Artificial Intelligence, Algorithms for Decision-Making, Automated Verification, Foundations of Hybrid and Embedded Systems

Skills

Programming Python, C++, Julia, LaTeX

Languages:

Frameworks AirSim, Isaac Lab, OpenAI Gymnasium, Numpy, Pytorch, Tensorflow
& Libraries: