

Sriram Gopalakrishnan

✉ marirsg2@gmail.com

AI/ML Researcher, Engineer, and Manager specializing in Generative and Classical AI, Human-AI Interaction, and Reliable Agentic Systems. Experienced in leading research and development of AI products for business intelligence and operations. Research expertise includes integrating insights from cognitive science and psychology into AI algorithms and decision-making systems, with a focus on effective human-in-the-loop interaction and robust, reliable agentic AI.

Education

Arizona State University

Ph.D. Computer Science

Tempe, AZ

December 2022

"Incorporating Human Cognitive Limitations Into Sequential Decision Making Problems and Algorithms"

Advisor: Prof. Subbarao (Rao) Kambhampati

Lehigh University

M.Sc. Computer Science with Thesis

Bethlehem, PA

May 2017

"Learning Hierarchical Task Networks Using Word Embeddings"

Advisor: Dr. Héctor Muñoz-Avila

Lafayette College

B.Sc. Electrical Engineering, B.A. in Computer Science

Easton, PA

May 2010

SPECIALIZATIONS AND SKILLS

- Research: Automated Planning, Reinforcement Learning, Human-AI Interaction and teaming, Large Language Models, Search and Optimization, and Natural Language Processing.
- Managed and mentored teams of employees and interns. Has led development across onshore and offshore teams and coordinating cross-team engineering initiatives. Contributed to defining recruitment and interview processes, including updating coding interviews to account for the impact of coding agents.
- Proven ability to translate senior leadership vision and customer feedback into successful products and outcomes.
- Completed CITI certification for responsible research. Conducted human subject experiments for Human-Aware-AI research.
- Packages/Tools: Pytorch, Ray, SciKit-Learn, Hugging Face, spaCy, CPLEX, and related packages.
- Worked with large code bases, agile software development process, as well as rigorous integration and test processes.

WORK EXPERIENCE

○ **AI Research Lead - AI Research Team (JP Morgan & Chase)**

February 2024 - Present

- Co-developed an algorithm to generate synthetic data for evaluating multi-hop logical reasoning of LLMs. Built multi-agent systems for this task to handle corpus of long documents with inter-document references.
- Helped set the research team-agenda as part of a core-team, and specifically worked on the agenda for human-AI interaction, data generation for AI, and continual learning.
- Built internal token-efficient, secure coding-agents for non-coders with an intuitive interface that runs directly on a browser leveraging Pyodide.
- Developed and programmed a process to generate datasets for optimizing retrieval models in the query-by-document search setting.
- Built AI tool used in production to optimize a large and critical company process for new products with classical information retrieval techniques and LLMs; applied principles from explainable AI and human factors research in product for improving user interaction.
- Co-developed and put into production a tool to help optimize the scheduling of tasks with an easy-to-use user interface for non-technical users to enter optimization constraints.
- Researched and published work on combining LLMs and automated planners for planning over multi-step LLM tasks, and travel planning.
- Managed employees and interns across research and product development, developing under-performing team members and enabling strong results within short timeframes.

○ **AI Research Senior Associate - AI Research Team (JP Morgan & Chase)**

August 2022 - February 2024

- Did research on constrained time series generation and retrieval that combined diffusion models and non-convex optimization. Research yielded 2 conference publications including a best industry paper award.
- Worked on safer algorithmic recourse. This research work with my intern yielded a conference publication.
- Developed tool to allow user to more easily specify a model of domain dynamics and generate synthetic data. This was used for evaluating the impact of internal programs.

○ **AI Research Summer Associate (intern) - AI Research Team (JP Morgan & Chase)**

May - August 2021

- Worked on algorithms for intelligently assigning and prioritizing tasks for human teams. Research yielded a conference publication. This work incorporated human factors into the algorithm to account for human limitations in context (task) switching.

○ **AI Research Summer Associate (intern) - AI Research Team (JP Morgan & Chase)**

June - August 2020

- Worked on the goal recognition problem for detecting client (financial) goals to offer better services. Presented work at a workshop on Finance and Planning.

- **Graduate Research Associate - Yochan Lab (Arizona State University)**
August 2017 - May 2022
 - Researched human-aware-AI in the context of sequential decision making.
 - Dissertation topic "Incorporating Human Cognitive Limitations Into Sequential Decision Making Problems and Algorithms".
- **Graduate Research Associate - Insyte Lab (Lehigh University)**
August 2015 - May 2017
 - Researched learning Hierarchical Task Networks (HTN) from action traces using clustering and action embeddings (thesis work)
- **Design and Development Engineer (Lutron Electronics)**
January 2012 - August 2015
 - Helped coordinate the embedded engineering team for a large-scale smart lighting and HVAC control system for home and commercial building automation.
 - Improved upon software architecture. Co-designed and implemented a new communication protocol for the system. This work yielded a patent.
 - Mentored new employees.
 - Represented the company at tradeshow, and conducted training classes for customers (installers).
- **Project Embedded Electrical Engineer (Lutron Electronics)**
June 2010 - December 2012
 - Developed new features for a lighting control system, running on a 32-bit microcontroller.
 - Used design patterns to write robust code. Designed comprehensive test case scenarios for verification.
 - Worked with the manufacturing team to troubleshoot production and hardware issues.
 - Helped coordinate the offshore development team for embedded software.
 - Supervised interns.

PUBLICATIONS

Google Scholar: <https://scholar.google.com/citations?user=So86Wl4AAAAJ>

- Daneshi, R., Patra, S., Dwarakanath, K., **Gopalakrishnan, S.**, Borrajo, D., Sreedharan, S. (2026). The curious case of planning for unreliable agents: Challenges and opportunities in orchestrating generative AI agents. In Workshop on Planning in the Era of LLMs at the International Conference on Automated Planning and Scheduling (ICAPS-26).
- Patra, S., Ramani, K., Borrajo, D., & **Gopalakrishnan, S.** (2025). Generating domain specific natural language SAT reasoning datasets. In Workshop on Efficient Reasoning at the Conference on Neural Information Processing Systems (NeurIPS).

- Borrajo, D., Canonaco, G., de la Rosa, T., Garrachón, A., **Gopalakrishnan, S.**, Kaur, S., & Veloso, M. (2025). GenPlanX: Generation of plans and execution. In Workshop on Bridging Language, Agent, and World Models at the Conference on Neural Information Processing Systems (NeurIPS).
- Morales, M., Pozanco, A., Canonaco, G., **Gopalakrishnan, S.**, Borrajo, D., & Veloso, M. (2025). On learning action costs from input plans. In the Proceedings of the European Conference on Artificial Intelligence (ECAI).
- de la Rosa, T., **Gopalakrishnan, S.**, Pozanco, A., Zeng, Z., & Borrajo, D. (2024). TRIP-PAL: Travel planning with guarantees by combining large language models and automated planners. *arXiv preprint arXiv:2406.10196*.
- Wu, H., Sharma, S., Patra, S., & **Gopalakrishnan, S.** (2024). SafeAR: Towards safer algorithmic recourse by risk-aware policies. In the Proceedings of the AAAI Conference on Artificial Intelligence.
Outcome of a research internship that I mentored
Oral Presentation (one of only 2% of accepted papers)
- Coletta, A., **Gopalakrishnan, S.**, Borrajo, D., & Vyetrenko, S. (2023). On the constrained time-series generation problem. In the Proceedings of the Conference on Neural Information Processing Systems (NeurIPS).
- Bamford, T*, Coletta, A.*, Fons, E.*, **Gopalakrishnan, S.***, Vyetrenko, S., Balch, T., & Veloso, M. (2023). Multi-modal financial time-series retrieval through latent space projections. In the Proceedings of the International Conference on AI in Finance (ICAIF).
*** Equal contribution**
Winner of the Best Industry Paper Award at ICAIF 2023
- Patra, S., Mahfouz, M., **Gopalakrishnan, S.**, Magazzeni, D., & Veloso, M. (2023). FinRDDL: Can AI planning be used for quantitative finance problems? In Workshop on Planning for Financial Services (FinPlan) at ICAPS.
- Thai, T., Verma, M., Soni, U., **Gopalakrishnan, S.**, Shen, M., Garg, M., & Scheutz, M. (2023). Methods and mechanisms for interactive novelty handling in adversarial environments. In the Proceedings of the 23rd International Conference on Autonomous Agents and MultiAgent Systems (AAMAS).
- Taitler, A., Gimelfarb, A., Jeong, J. M., **Gopalakrishnan, S.**, Mladenov, M., Liu, X., & Sanner, S. (2023). pyRDDLGym: From RDDL to Gym environments. In Workshop on Planning and Reinforcement Learning (PRL) at ICAPS.
- **Gopalakrishnan, S.**, & Borrajo, D. (2022). Assignment and prioritization of tasks with uncertain durations for satisfying makespans in decentralized execution. In the Proceedings of the 32nd International Conference on Automated Planning and Scheduling (ICAPS).
- **Gopalakrishnan, S.**, & Kambhampati, S. (2022). Minimizing robot navigation-graph for position-based predictability by humans. In the Proceedings of the 22nd International Conference on Autonomous Agents and MultiAgent Systems (AAMAS).

- **Gopalakrishnan, S.**, Verma, S., & Kambhampati, S. (2021). Synthesizing policies that account for human execution errors caused by state aliasing in Markov decision processes. In Workshop on Explainable AI in Planning (XAIP) at ICAPS.
- **Gopalakrishnan, S.**, Soni, U., Thai, T., Lymperopoulos, P., Scheutz, M., & Kambhampati, S. (2021). Integrating planning, execution, and monitoring in the presence of open-world novelties: Case study of an open-world Monopoly solver. In Workshop on Integrating Planning and Execution at ICAPS. (Associated AI agent was top performer in the DARPA SAIL-ON program for the Monopoly domain.)
- **Gopalakrishnan, S.**, Cohen, L., Koenig, S., & Kumar, T. S. (2020). Embedding directed graphs in potential fields using FastMap-D. In *Proceedings of the Thirteenth Annual Symposium on Combinatorial Search (SOCS)*.
- Borrajo, D., **Gopalakrishnan, S.**, & Potluru, V. (2020). Goal recognition via model-based and model-free techniques. In Workshop on Planning for Financial Services (FinPlan) at ICAPS.
- **Gopalakrishnan, S.**, Soni, U., & Kambhampati, S. (2019). Feature-directed active learning. In Workshop on Explainable AI in Planning at ICAPS.
- **Gopalakrishnan, S.**, & Kambhampati, S. (2019). TGE-viz: Transition graph embedding for visualization of plan traces and domains. Extended abstract presented at the International Conference on Automated Planning and Scheduling (ICAPS).
- Zha, Y., Li, Y., **Gopalakrishnan, S.**, Li, B., & Kambhampati, S. (2018). Recognizing plans by learning embeddings from observed action distributions. In the Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS).
- **Gopalakrishnan, S.**, Muñoz-Avila, H., & Kuter, U. (2018). Learning task hierarchies using statistical semantics and goal reasoning (journal version). *AI Communications*.
- Nguyen, C., Reifsnyder, N., **Gopalakrishnan, S.**, & Muñoz-Avila, H. (2017). Automated learning of hierarchical task networks for controlling Minecraft agents. In 2017 IEEE Conference on Computational Intelligence and Games (CIG).
- **Gopalakrishnan, S.** (2017). *Learning hierarchical task networks using semantic word embeddings* (Masters dissertation, Lehigh University).
- **Gopalakrishnan, S.**, Muñoz-Avila, H., & Kuter, U. (2016). Word2HTN: Learning task hierarchies using statistical semantics and goal reasoning. In 4th Workshop on Goal Reasoning at IJCAI.

PATENTS

- Granted Patents.....
- **Load Control system responsive to location of an occupant and mobile devices.** 2026.
US Patent No.: 12,602,020
<https://patents.google.com/patent/US12602020B2/en>
 - **Method and System for Multi-Modal Time-Series Retrieval Through Latent Space Projections.** 2025.

US Patent No.: 12,475,133

<https://patents.google.com/patent/US12475133B2/en>

- **Digital messages in a load control system.** 2024.

US Patent No.: 12,127,317

<https://patents.google.com/patent/US12127317B2/en>

Filed Patents.....

- **System and Method for Generating Data by Combining Large Language Models and Automated Planners.** 2024.

Application No.: 18/983,707 (Pub. No.: US 2026/0170578 A1)

<https://patents.google.com/patent/US20260170578A1/en>

- **Method and System for Constrained Time Series Generation.** 2023.

Application No.: 18/368,365 (Pub. No.: US 2025/0094522 A1)

<https://patents.google.com/patent/US20250094522A1/en>

- **Method and System for Assignment and Prioritization of Tasks in Decentralized Execution.** 2023.

Application No.: 18/118,529 (Pub. No.: US 2024/0303108 A1)

<https://patents.google.com/patent/US20240303108A1/en>