Prahitha Movva

prahitha.movva03@gmail.com ♢ LinkedIn ♢ GitHub ♢ Twitter ♢ Google Scholar

Objective

I work at the intersection of NLP, Vision and cognitively-inspired reasoning. The goal of my research is to develop robust and safe AI models that move beyond pattern recognition to achieve a deeper, context-sensitive understanding of the world. Specifically, I aim to establish diagnostic methods for multimodal reasoning, contamination-resistant evaluation frameworks, and trustworthy verification methods for AI systems before deployment.

Education

University of Massachusetts, Amherst — Amherst, MA, USA
MS in Computer Science, (GPA 3.84/4.0) — *Sep. 2023 – Dec. 2024*
Selected Courses: Reinforcement Learning, Advanced Algorithms, Machine Learning, Information Retrieval

Indian Institute of Information Technology, SriCity — Chittoor, AP, India
B.Tech in Computer Science, (GPA 8.7/10.0) — *Sep. 2018 – Jun. 2022*
Selected Courses: Artificial Intelligence, Deep Learning, Cryptography, Natural Language Processing, High Performance Computing, Digital Image Processing, Discrete Math and Probability Theory, Statistical Data Analysis

Publications

C=Conference, J=Journal, W=Workshop, S=In Submission, P=Preprint

[**W.1**] Ethan Seefried, Naga Harshita Marupaka, **Prahitha Movva**, Tilak Kasturi, Tirthankar Ghosal. (2025). **Enginuity: Building an Open Multi-Domain Dataset of Complex Engineering Diagrams**. In *NeurIPS 2025 AI4Science Workshop - Oral, $2000 grant (Top 5)*

[**W.2**] **Prahitha Movva** (2025). **Reasoning Riddles: How Explainability Reveals Cognitive Limits in Vision-Language Models**. In *COLM 2025 XLLM-Reason-Plan Workshop*

[**W.3**] **Prahitha Movva**, Naga Harshita Marupaka. (2025). **Enhancing Scientific Visual Question Answering through Multimodal Reasoning and Ensemble Modeling**. In *ACL 2025 Scholarly Document Processing (SDP) Workshop*

[**C.1**] **Prahitha Movva**, Hemanth Pasupuleti, Himangshu Sarma. (2022). **A self learning Yoga monitoring system based on pose estimation**. In *International Conference on Human-Computer Interaction (HCII 2022)*

[**P.1**] Prahitha Movva. (2022). **Survey on activation functions: A comparative study between state-of-the-art activation functions and oscillatory activation functions**.

Research Experience

Research Assistant — Sep. 2024 – Dec. 2024
University of Massachusetts Amherst — *Amherst, MA, USA*
Advised by Prof. Madalina Fiterau

- Proposed frequency masking as a self-supervision task for time series representation learning, modifying MOMENT architecture's random patch masking with wavelet-based frequency component masking.
- Conducted comprehensive ablation studies across 15+ masking configurations.
- Achieved state-of-the-art results on 96h forecasting and competitive performance across 192h-720h horizons.

Graduate Student Researcher (Report) — Feb. 2024 – May. 2024
Spotify — *Amherst, MA, USA*
Advised by Paul Bennett, Praveen Chandar, and Ghazal Fazelnia

- Investigated personalization spectrum for LLM-based sequential recommendation by fine-tuning TASTE (T5-based) models at user-level and cohort-level granularities.
- Implemented LoRA-based fine-tuning, achieving 14.85% improvement in Recall@10 for high-activity user cohorts.
- Designed and evaluated multiple user clustering strategies, analyzing trade-offs between personalization granularity and computational efficiency.

Deep Learning Research Intern (Certificate) — Feb. 2021 – May. 2021
UK-India Education and Research Initiative — *India (remote)*

- Worked with Flutura Decision Sciences and Analytics; won the runners-up award.
- Built a U-Net + SE-ResNet + WGAN-GP model to generate defective printed circuit board (PCB) images.
- Improved defect representation quality, surpassing StarGAN baseline.

Work Experience

Research Engineer Intern Dec. 2025 – Apr. 2026
Predii *San Jose, CA, USA*

- Iterated a two-step extraction pipeline for structured JSON from warranty PDFs; fixed systematic failures across batches - compound-part normalization, coverage inheritance
- Built and deployed the clustering and LLM-as-a-judge evaluation pipeline as a FastAPI microservice; cut manual annotation from 1500+ docs to 36 gold references, with MongoDB integration and multi-env subprocess isolation

Data Science Intern Jun. 2024 – Aug. 2024
Santander Bank *Boston, MA, USA*

- Architected an Airflow pipeline reducing runtime by 85% and saving $250K annually.
- Validated 40+ SAS queries, improving accuracy of customer segmentation for millions of users.
- Owned several tasks in savings and checking products, & compiled contact strategies for customized user segments.

DSSG Fellow (Poster, Report) Jun. 2023 – Sept. 2023
German Research Center for AI (DFKI) *Kaiserslautern, Germany*
Advised by Prof. Sebastian Vollmer

- Built a geospatial data pipeline using kNN & IDW to augment high-res weather data (0.1°x0.1°) with survey data.
- Implemented DBSCAN, DTW, and correlation analysis to study poverty-weather patterns across 6 socioeconomic indicators.
- Designed a scalable Streamlit pipeline enabling UNICEF policy teams to extend analyses across multiple countries.

Data Science Analyst Jan. 2022 – Jun. 2023
PharmEasy *India (remote)*

- Increased sales 5% by developing automated competitor price tracking system & optimizing pricing strategies.
- Developed ARIMA and SARIMA models for demand prediction, reducing stockouts and improving inventory efficiency.
- Boosted PL product ratings 11% through a Python/SQL pipeline tracking 51M monthly user interactions.
- Developed & owned reporting for a rewards program with Python, SQL & Microstrategy, saving 40 hours of work.
- Collaborated with business stakeholders to improve customer experience by identifying opportunities that drive fulfillment

Summer Intern (PRs) May. 2021 – Jul. 2021
Outreachy Intern at SymbiFlow (now ChipsAlliance) *Poland (remote)*

- Selected for Outreachy's internship (2.1% acceptance rate) to contribute to sv-tests, an open-source hardware verification tool.
- Implemented advanced search & filtering features with JavaScript, jQuery, and Jinja templates, improving debugging workflows.
- Contributed frontend redesigns and performed end-to-end testing with Selenium, making the system more accessible to developers.

Projects

Cluedo AI: Evaluating LLM Social & Deductive Reasoning Ongoing
A multi-agent platform for benchmarking LLM reasoning and Theory of Mind using the classic deduction game Clue. Includes a full game engine, memory system, and evaluation pipeline for comparing LLM strategic behavior.

- Designed a multi-LLM competitive environment (GPT, Claude, Gemini, Llama, etc.) for testing social reasoning and hidden-information deduction.
- Built agent memory systems that track beliefs, eliminated cards, and hidden states across turns.
- Developed a turn-based decision-making pipeline using structured YAML prompts for suggestion logic, challenges, and memory updates.
- Implemented visualization tools (Node.js + Socket.IO) to watch LLM agents play Clue in real-time.
- Ran large-scale evaluations and prepared batches for Dria and Predibase for fine-tuned GRPO training

Realtime TTS Aug 2025
A real-time Text-to-Speech (TTS) WebSocket service that provides streaming audio synthesis with character-level timing alignment. Built using the Kokoro TTS engine.

- Implemented character-level timing alignment system providing precise temporal mapping between input text and synthesized speech, enabling accurate subtitle generation.
- Engineered LaTeX mathematical notation parser and smart text chunking algorithm that splits input at sentence boundaries, improving synthesis quality for technical content while maintaining natural speech flow.

Statistical Time Series Analysis and Forecasting Dec 2020

Implemented a time-series analysis and forecasting pipeline using Python, applying foundational statistical tests, exploratory analysis, and predictive models on multivariate time-series data.

- Evaluated the dataset using normality tests, Durbin-Watson (collinearity checks), Granger causality, and Augmented Dickey-Fuller (stationarity checks).
- Developed predictive models using vector autoregression (VAR/OLS) and ARIMA models for forecasting time-series behavior.

Academic Service

NeurIPS, COLM, ICCV

- Reviewer for Main Conference Ethics, AI4Science and Math-AI Workshops (NeurIPS), XLLM-Reason-Plan Workshop (COLM), LIMIT Workshop (ICCV)

Teaching, Mentoring, and Other Service

Teaching and Academic Support

- *Grader, CMPSCI 532: Systems for Data Science (UMass Amherst)*: Evaluated graduate-level coursework in distributed systems and data pipelines; supported student learning through detailed feedback.
- *Student Mentor, Command Tech*: Mentored high-school girls on coding fundamentals, project development, and college application strategies.

Technical Mentoring and Community Engagement

- *Community Champion for the Applied AI Research Track, LevelUp Labs*: Curating and delivering biweekly applied AI research talks to senior industry practitioners (staff, principal, CEO level); publishing articles and walkthroughs on applied AI on platforms like Substack and O'Reilly Radar.
- *ML Industry Program Lead, Cohere Labs*: Designed and led sessions connecting industry engineers with early-career ML researchers; supported participants with project development and technical guidance.
- *Mentor at HackDavis, Hoya Hacks, and multiple other hackathons*: Assisted 60+ teams with debugging, idea refinement, prototyping, and pitch preparation.
- *Community Volunteer, Rewriting the Code (RTC)*: Engaged in support and mentorship for women in tech through community participation.

Open Source Contributions and Mentorship

- *Open-Source Mentor (PClub Summer of Code, Script Winter of Code, and others)*: Mentored 13+ contributors on projects including Rotten Scripts, ChessJS, and a Spotify Recommendation Engine; guided students on version control, code reviews, and project planning.
- *Frontend Developer, Learning Equality (Menon Labs)*: Built UI components and filtering logic for open-source educational tools.
- *Contributor, Ignitus-client*: Developed React (TypeScript) components using atomic design principles to enhance modularity and reusability.
- *Web Developer, Gateway Regional Council (Develop for Good)*: Collaborated with a 5-member team to build an accessible, user-centered WordPress interface for a regional nonprofit.

Event Organizing and Technical Logistics

- *Sessions and Speakers Coordination Lead, Winter Analytics Summit 2026*: Managed keynote outreach, session scheduling, and event logistics; worked with the technical team on ticket pricing and demand estimation.
- *Hackathons and Workshops Lead, IOTA*: Organized hands-on workshops on web development and blockchain; coordinated hackathon logistics and student engagement at IIITS.
- *Tech Coordinator, SheCanSTEM Hackathon*: Managed technical infrastructure (Discord, website, registration systems) for a large-scale hackathon supporting women in STEM.
- *Challenge Master, The WireUs Challenge*: Led a team of 3 and created 40 programming challenges for 500+ participants as part of 100 Days of JS.

Awards

NeurIPS AI4Science Dataset Proposal Competition Winner (2025): Received a $2000 award by Renaissance Philanthropy for the Enginuity dataset proposal (oral presentation)

COLM DEI Scholarship (2025): Awarded $700 travel stipend to attend COLM 2025 (Montreal)

KubeCon + CloudNativeCon Europe (2023): Awarded $2500 travel stipend and complimentary registration

PSI:ML Summer School, by Microsoft (2023): Fully funded participant in the ML summer school in Serbia

PyCon Dublin (2022): Awarded $1500 travel stipend and complimentary registration to attend PyCon Ireland 2022

Capture-the-Flag (CTF) Wins (2020-2021): Won multiple (COMPFEST CTF 2021, Pragyan CTF 2021, REST-CON CTF 2020, Shakti CTF 2020 etc) CTFs as part of teams Cracken and warlock_rootx

Hackathon Wins (2021): Won 5 hackathons and $8k+ in awards for projects focused on accessibility and inclusivity