

Athresh Karanam

Department of Computer Science • University of Texas at Dallas • Dallas, TX
athresh.karanam@utdallas.edu • athresh.github.io • github.com/athresh

EDUCATION

Ph.D. in Computer Science, University of Texas at Dallas May 2026

Dissertation: *Effective, Efficient, and Explainable Learning in Data-Scarce and Noisy Domains*

Advisor: Prof. Sriraam Natarajan

B.Tech. in Civil Engineering, Indian Institute of Technology, Guwahati May 2015

Advisor: Prof. Subashisa Dutta

RESEARCH INTERESTS

Deep Tractable Probabilistic Models, Human-Allied Learning, Explainable AI, Data- and Compute-Efficient Learning, Probabilistic Generative Modeling, Submodular Optimization, Causal Inference, Generalizable ML, AI for Healthcare.

PUBLICATIONS

Conference Papers

- [1] Prabhakar, N., Tenali, P., Abudeye Fernandez, W., Borah, S., **Karanam, A.**, Blasch, E., Sundaravadivel, P., Natarajan, S. On the Transferability of Agricultural Weed Detection Under Cross-Field Distribution Shift. *International Conference on Dynamic Data Driven Applications Systems (DDDAS 2026)*.
- [2] Prabhakar, N., Balaji, V., **Karanam, A.**, Kersting, K., Natarajan, S. Neurosymbolic Imitation Learning with Human Guidance: A Privileged Information Approach. *International Joint Conference on Learning & Reasoning (IJCLR 2026)*.
- [3] **Karanam, A.**, Mathur, S., Sidheekh, S., Natarajan, S. A Unified Framework for Human-Allied Learning of Probabilistic Circuits. *AAAI Conference on Artificial Intelligence (AAAI 2025)*.
- [4] **Karanam, A.**, Killamsetty, K., Kokel, H., Iyer, R. ORIENT: Submodular Mutual Information Measures for Data Subset Selection under Distribution Shift. *Advances in Neural Information Processing Systems (NeurIPS 2022)*.
- [5] **Karanam, A.**, Mathur, S., Radivojac, P., Kersting, K., Natarajan, S. Explaining Deep Tractable Probabilistic Models: The Sum-Product Network Case. *International Conference on Probabilistic Graphical Models (PGM 2022)*.
- [6] Mathur, S., **Karanam, A.**, Radivojac, P., Haas, D. M., Kersting, K., Natarajan, S. Exploiting Domain Knowledge as Causal Independencies in Modeling Gestational Diabetes. *Pacific Symposium on Biocomputing (PSB 2023)*.
- [7] Zečević, M., Dhimi, D. S., **Karanam, A.**, Natarajan, S., Kersting, K. Interventional Sum-Product Networks: Causal Inference with Tractable Probabilistic Models. *Advances in Neural Information Processing Systems (NeurIPS 2021)*.
- [8] **Karanam, A.**, Hayes, A. L., Kokel, H., Haas, D. M., Radivojac, P., Natarajan, S. A Probabilistic Approach to Extract Qualitative Knowledge for Early Prediction of Gestational Diabetes. *International Conference on Artificial Intelligence in Medicine (AIME 2021)*.

Workshop Papers

- [9] **Karanam, A.**, Natarajan, S. Test-Time Active Feature Selection through Tractable Acquisition Functions. *Tractable Probabilistic Modeling Workshop, UAI 2023*.
- [10] **Karanam, A.**, Mathur, S., Sidheekh, S., Natarajan, S. Bayesian Learning of Probabilistic Circuits with Domain Constraints. *Tractable Probabilistic Modeling Workshop, UAI 2023*.

Tutorials

- [11] Mathur, S., Sidheekh, S., **Karanam, A.**, Natarajan, S. Deep Tractable Probabilistic Models. *Tutorial at CODS-COMAD 2024*.

RESEARCH EXPERIENCE

Ph.D. Researcher, Starling Lab, University of Texas at Dallas

Aug. 2020 – Present

Advisor: Prof. Sriraam Natarajan

Dissertation research on deep tractable probabilistic models, human-allied learning, and explainability in data-scarce domains. Developed a unified framework for incorporating expert domain knowledge into probabilistic circuits (AAAI 2025); introduced CSI-tree representations for extracting human-interpretable context-specific independencies from learned models (PGM 2022); proposed ORIENT, a submodular mutual information-based data subset selection framework for domain adaptation with $2\text{--}3\times$ training speedups (NeurIPS 2022); and applied probabilistic models to clinical datasets for gestational diabetes prediction and knowledge extraction.

AI Research Intern, Altir LLC

Sept. 2025 – Nov. 2025

Developed an interpretable, uncertainty-aware LLM agent by augmenting large language models with local tractable probabilistic reasoning, improving reliability through faithful uncertainty quantification.

Machine Learning Intern, RelationalAI Inc.

Jan. 2020 – Apr. 2020

Built an interactive product recommendation tool using tractable probabilistic models; trained and evaluated boosted relational models on an online vendor catalog.

TEACHING EXPERIENCE

Co-Instructor — Advanced Machine Learning (Industry Course)

Aug. 2021 – Oct. 2021

Vistra Corporation / University of Texas at Dallas

Co-designed and co-taught an eight-week intensive ML course for engineers at Vistra Corporation, with Prof. Sriraam Natarajan and Dr. Harsha Kokel. The curriculum covered core ML algorithms, probabilistic reasoning, and neural networks, with emphasis on motivation, transfer, and deployment in high-stakes settings. Teaching working professionals sharpened my focus on connecting formalism to problems that matter to practitioners.

Tutorial Presenter — Tractable Probabilistic Models

Mar. 2023

Indian Institute of Technology, Madras

Delivered a three-day tutorial on tractable probabilistic models to researchers and ML practitioners, co-presented with Prof. Natarajan, Dr. Saurabh Mathur, and Sahil Sidheekh. With a technically advanced audience and dense material, the focus was on building a coherent conceptual framework and bridging to adjacent areas of specialization.

Graduate Research Assistant and Mentor, Starling Lab, UT Dallas Aug. 2020 – Present
Mentored undergraduate and graduate students in reading technical literature, forming research questions, debugging implementations, and presenting results.

PRESENTATIONS AND INVITED TALKS

Tutorial: Deep Tractable Probabilistic Models 2024
CODS-COMAD 2024. Co-presented with S. Mathur, S. Sidheekh, S. Natarajan.

Tutorial: Tractable Probabilistic Models Mar. 2023
Indian Institute of Technology, Madras. Co-presented with S. Natarajan, S. Mathur, S. Sidheekh.

ACADEMIC SERVICE

Reviewing

- Program Committee / Reviewer: AAAI (2022–2025), NeurIPS (2023–2026), ICLR (2024).
- Workshop Reviewer: Neuro Causal and Symbolic AI Workshop at NeurIPS (2022).
- Journal Reviewer: *Data Mining and Knowledge Discovery*; *Journal of Artificial Intelligence Research*.

PRIOR INDUSTRY EXPERIENCE

Software Developer, NetCracker Technology Pvt. Ltd. June 2015 – Jan. 2017
Developed features for ICOMS (Integrated Communications Operations Management System) across multiple product releases.

TECHNICAL SKILLS

Python, PyTorch, Java, MATLAB, R, Git, SQL.