Vishnu Shreeram M. P.

Research Interests: Generative models, optimization & operations research, Applied ML, NLP

+91 9633695099 | 142201014@smail.iitpkd.ac.in
linkedin.com/vishnu-m-p | github | portfolio

Education

Indian Institute of Technology, Palakkad — Expected May 2026
*Bachelor of Technology in Data Science (**CGPA: 9.6 / 10**)* — *Palakkad, Kerala*

Publications & Honors

Tractable Sharpness-Aware Learning of Probabilistic Circuits — Hrithik Suresh, Sahil Sidheekh, **Vishnu Shreeram M. P.**, Sriraam Natarajan, Narayanan C. Krishnan. *Accepted for Oral Presentation at AAAI 2026*, Singapore.

Patent Filed — *Method, Apparatus, and System for Increasing Throughput in a Surface Mount Technology Assembly Line*. Patent Application No: 202541068350, Filed July 17, 2025.

Certificate of Academic Excellence — Awarded by IIT Palakkad for securing the **highest CGPA** among **first-year** (2022–23), **second-year** (2023–24) and **third-year** (2024–25) students in the Data Science Department.

GATE 2025 — Secured **All India Rank 73** in Data Science and Artificial Intelligence.

Experience

Laboratory of Statistical Artificial Intelligence and Machine Learning — May 2025 – July 2025
Research Intern — *Palakkad, Kerala*

- Co-authored **"Tractable Sharpness-Aware Learning of Probabilistic Circuits"**, accepted for an **Oral Presentation at AAAI-26** (Singapore); preprint available on arXiv.
- Implemented the paper's theoretical ideas into major PC libraries (**Einsum Networks** and **Pyjuice**), utilising second-order information to guide models toward flatter, more robust optima
- Conducted ablation studies on HPC infrastructure to validate the theory, showing that our method scales **linearly** with model depth while standard automatic differentiation scales exponentially.
- Ran experiments across multiple datasets to benchmark the hybrid method, proving it outperforms standard EM and SGD baselines in **data-scarce regimes**.

Laboratory of Statistical Artificial Intelligence and Machine Learning — May 2024 – April 2025
Part-time Research Intern — *Palakkad, Kerala*

- Optimized Bosch's SMT assembly line configuration involving 32 PCB types and 200+ unique components.
- Built a **hybrid optimization solver** combining **MILP (Gurobi)** and **Greedy algorithms** to handle combinatorial complexity, increasing **Components-Per-Hour (CPH)** by at least **7%**.
- Resulted in a **patent filing**, recognizing the solution's novel approach to balancing theoretical optimality with real-world manufacturing constraints.

Selected Projects

Interventional Sum-Product Networks (iSPN) with GNNs | *PyTorch, Causal Inference*

- Integrated Graph Convolutional Networks (GCNs) into the iSPN framework to encode **mutilated causal graphs**, enabling the model to learn from interventional data distributions.
- Improved the modeling of **causal interventions** in sparse networks by replacing standard MLP parameter generators with graph-structured encoders.
- Achieved a **26.15%** reduction in NLL on the ASIA dataset compared to baselines, empirically validating that structural inductive biases improve causal reasoning.

Viśva Mitra - Voice Enabled Agentic AI Assistant | *LangChain, MCP, Ollama, Docker, MLflow*

- Architected a modular agentic system using the **Model Context Protocol (MCP)** to facilitate communication between LLMs and isolated, containerized tool-execution environments.
- Integrated a multi-modal pipeline (Audio/Text) with **LangChain and Ollama** for intelligent tool-routing across 15+ actions, including hardware control and real-time search.
- Engineered a containerized sandbox environment (Docker) to ensure reproducibility and safety during autonomous agent tool-use, monitored via MLflow traces.

Technical Skills

Languages: Python, SQL (PostgreSQL)
ML & Data: PyTorch, TensorFlow, scikit-learn, NumPy, Pandas, Hugging Face Transformers
Systems & Tools: Docker, FastAPI, Git, Linux, Spark, Hadoop, Hive, AWS, Vector DBs (Pinecone, LanceDB), MLflow, MCP