

Srikar Sribhashyam

+91 7780590157 | srikarsribhashyam@gmail.com | [LinkedIn](#) | [Github](#)

EDUCATION

Birla Institute of Technology and Science Pilani, Hyderabad Campus

Hyderabad, India

M.Sc. Mathematics & B.E. Mechanical Engineering

2020 – 2025

- **Relevant Coursework:** Deep Learning, Quantum Computing, Linear Algebra, Object-Oriented Programming, Data Structures and Algorithms, Optimization, Graphs and Networks, Number Theory, Mathematical Modeling, Probability
- **Research Interests:** AI Memory Systems, Knowledge Graphs, Graph Representation Learning, Temporal Reasoning
- **Audited:** Yale CPSC 483/583: DL on Graph-Structured Data

WORK EXPERIENCE

SG Analytics (Straive)

Bengaluru

Data Scientist

July 2025 – Present

- Designed and executed **15+** hypothesis tests and developed **3 hierarchical Bayesian models** using MCMC (NUTS) for school funnel optimization, achieving robust posterior convergence ($\hat{R} < 1.012$), Pareto-LOO cross-validation diagnostics ($k < 0.6$), and reliable forecasting through posterior predictive validation. (Azure Databricks, Azure DevOps)
- Built a **smart PPT generator** capable of refreshing slides from templates, generating new decks, and supporting an **interactive chatbot with RAG induced memory** that can directly edit slides. Engineered **Multi-agent systems** (LangGraph) and **17 RESTful APIs** (FastAPI), reducing manual deck creation time by **85%**.
- Developed a Sales Executive **Agentic system** that ingested structured sales data, **Decision Trees** for week-level delta sales forecasting, and embedded **Causal Graphs** across multiple regions enhancing reasoning capability by **22%**.

Harman International

Bengaluru (Remote)

Data Science Intern (PPO)

Jan 2025 – June 2025

- Developed a **DOM Distiller**, a testing harness for the Genesis browser agent to extract structured web elements and validate agent functionality; built automated evaluation pipelines to improve reliability of structured extraction workflows.
- Architected a **Desktop Automation Agentic System** using **AutoGen**, GPT-4o, and Playwright; designed a **SWARM-based multi-agent architecture** with **4 autonomous agents** and **15+ integrated tools**, resulting in a **45% reduction in operational latency** and saving **150+ manual hours monthly** across legacy Windows applications.

SG Analytics

Bengaluru

Data Science Intern (PPO)

Aug 2024 – Dec 2024

- Designed a generalized **tabular data extraction pipeline** (PDF → Excel) for a leading ESG data provider, leveraging advanced computer vision models from FAIR and MSR (**Detectron2**, **DETR**) along with **GPT-4o/turbo** for intelligent table detection, contextual parsing, and hierarchical data structuring, reducing manual effort by over **80 man-hours** per week.

PROJECTS

mini_G: Graph Neural Network on Knowledge Graph (inspired by G-reasoner) | [GitHub](#)

Aug 2026

- **Query-conditioned DistMult NBFNet retriever** at toy scale in **PyTorch Geometric** (MessagePassing on a typed entity-document graph): Init → multi-hop DistMult → predict head (**0.45M params**), trained end-to-end for **multi-hop document ranking**.
- Benchmarked against dense retrieval on a curated multi-hop set (same frozen encoder) the GNN recovered ranking (**multihop MRR $\approx 0.14 \rightarrow 0.82$**), with a **no-edges ablation** to isolate structural message passing.

mem-void: Temporal Knowledge Graph Memory System for AI Agents | [GitHub](#)

June 2026

- Designing and building an open-source **temporal knowledge graph memory system** for AI agents using **Neo4j**, supporting **continuous event ingestion**, **entity resolution**, **temporal fact management**, and **graph-native retrieval**.
- Architected a **modular memory pipeline** inspired by Graphiti (open source), implementing **predicate-aware temporal invalidation**, **provenance tracking**, and **extensible retrieval** over evolving knowledge graphs.

LM-Hypotest: Deterministic Statistical Testing Library with LLM Interpretability | [PyPI](#) | [GitHub](#)

Apr 2026

- Designed and published a production-grade **Python statistical inference library** on PyPI implementing **deterministic hypothesis testing**, **automated assumption validation**, and **structured statistical result** along with option for **OpenAI compatible LLM summary** and suggestions.

Undergraduate Research Assistant | *Predictive Modeling of Tool Wear* | Prof. K. Kumar

Jan 2024 – May 2024

- Developed a **Bi-directional LSTM** pipeline augmented with **Variational Quantum Circuits (PennyLane)** for industrial tool wear prediction using **force and positional sensor data**. Engineered domain-specific features derived from machining literature, achieving a **4% reduction in MAE** and a final prediction error of **8.5 μm** on real-world machining datasets.

TECHNICAL SKILLS

Languages: Python (advanced), SQL, C++, Latex

Libraries: NumPy, Pandas, Scikit-learn, NLTK, HuggingFace, OpenCV

Frameworks: LangGraph, TensorFlow, PyTorch, Pytorch Geometric, PennyLane, Playwright, FastAPI, Pyspark

Databases: MongoDB, MySQL, Chroma DB, neo4j

Version Control: Git

Cloud (familiar with): Azure Databricks, Azure DevOps (ADO), Google Cloud Platform (GCP)