Surya “Sunny” Gundavarapu

Austin, TX • 909-234-5601 • surya_gundavarapu@yahoo.com
linkedin.com/in/suryagsr • github.com/sgundava

SUMMARY

ML/AI engineer with 8+ years shipping production systems at Meta and Prime Healthcare. Open-source author of tooling for ML reliability and meta-research auditing, with applied research in LLM hallucination detection and a national-level competitive ML track record.

EXPERIENCE

Meta Platforms Inc., Austin, TX *March 2022 – Present*

Business Support Engineer *August 2024 – Present*

- Own end-to-end user support journeys for global developer and business surfaces; design, build, and maintain support entry points that act as intake funnels between self-service and assisted channels.
- Drive reliability and developer-experience metrics (CSAT, case re-open rate, time-to-resolution) through systematic debugging, proactive monitoring, and infrastructure fixes.
- Provide integration support across cloud-based APIs, financial technology, and telecommunications systems; collaborate with Operations, Product, Program, and Engineering to scope features, estimate effort, and ship on roadmap.

Technical Analyst *March 2022 – August 2024*

- Built NLP-based classification and automation features in Hack and JavaScript that triage and auto-resolve support cases at Meta’s scale; leveraged state-of-the-art language models for intent and issue detection.
- Defined long-term automation strategy and roadmaps; partnered with analytics and product teams to size opportunities, scope features, and instrument automation-quality metrics.
- Developed, tested, and optimized automation rules to reduce cost-to-serve while improving operational efficiency and volume mitigation.

Prime Healthcare Management Inc., Ontario, CA *Sept 2018 – March 2022*

Business Analyst II *Jan 2021 – March 2022*

- Designed and built a React-based portal for physician-preference-item contract administration, with Python automation scripts for purchase-approval workflows.
- Automated purchasing across 45+ hospitals in Python, generating order feeds for recurring items that met strict pricing criteria and materially reducing manual processing errors.

Business Analyst *Sept 2018 – Jan 2021*

- Built and maintained a COVID data warehouse aggregating data from all 46 Prime Healthcare hospitals to support corporate decision-making and government reporting.
- Built an online dashboard to surface proactive savings opportunities in purchasing; built item crosswalks and contract-analysis tooling to generate value analyses for contracts under review.
- Developed automated report interfaces and data feeds supporting organizational decision-making.

RESEARCH & OPEN SOURCE

DConfusion *github.com/sgundava/DConfusion*

- Author and maintainer of an open-source Python library for confusion-matrix analysis and ML meta-research auditing. Features include metric reconstruction, probabilistic inference, statistical testing, and cost-sensitive analysis. Complementary to mlscorecheck; designed to catch errors and implausible results in published ML evaluations.

Data4Good 2025 — AI Hallucination Detection *Non-Academic Track Winner*

- Led team "Response Outliers (ResOut)" to first place in AI hallucination detection for educational Q&A; 99.11% balanced accuracy on held-out test data.
- Architected a dual-transformer system (DeBERTa-v3-Large + RoBERTa-Large) with agreement-based routing; LLM arbitration triggered only when models disagree (~1% of queries), reframing the LLM's task from ternary to binary classification and avoiding the failure mode of confidence-based routing on overconfident model errors.
- Presented at Johns Hopkins Carey Business School national championship, February 2026. Related paper submitted to INFORMS Journal on Data Science (IJDS).

The Scarcity Fallacy — Expertise Framing in LLMs *In progress, 2026*

- Independent research testing whether persona framing like "You are an expert in X" imposes a scarcity constraint LLMs don't have, suppressing cross-domain ideation on generative tasks. Task-dependent hypothesis: neutral for factual retrieval, beneficial for structured reasoning (via implicit CoT), harmful for creative ideation.
- 12-paper literature review complete; experimental design is 4 models × 4 tasks × 3 conditions × 15 generations (n=2,880) with Shannon entropy over a 15-domain taxonomy as the diversity metric. Pilot complete; refining measurement pipeline before full data collection.

PROJECTS

Prism — Multi-Agent AI Deliberation Tool *In development*

- Structured decision-making tool based on de Bono's Six Thinking Hats: AI personas supply objective perspectives in parallel while the user provides emotional context; a synthesizer integrates all inputs, including the user's own. Design thesis: LLMs lack authentic access to the user's emotional state, so that role stays with the human. Built on Next.js + TypeScript + Supabase + Vercel with a multi-provider AI abstraction layer (Anthropic + OpenAI), parallel streaming, @mention follow-ups, and on-demand synthesis with dissent flagging.

LinkedIn Resume Optimizer *In development, 2026*

- Two-pass LLM pipeline (PDF normalization → job-description-aware optimization) on FastAPI + React/Vite. Railway deployment.

LLM-Powered Q&A Application *2023*

- Built an early LangChain-based Q&A system using the OpenAI API and Python during the initial LLM adoption wave, demonstrating practical LLM integration before such patterns became standard tooling.

Cricket Management Simulation

- Agent-based modeling game simulating player, board, and fan stakeholder dynamics; tests the user's ability to balance competing expectations through emergent event simulation.

SKILLS

- **Languages:** Python, JavaScript/TypeScript, Hack, SQL, R
- **ML & AI:** PyTorch, Hugging Face Transformers, TensorFlow, LLM APIs, FastAPI, LLM evaluation, hallucination detection, NLP
- **Data & Infrastructure:** Spark, Hadoop, Hive, Docker, Kubernetes, Git, MongoDB, Tableau, PowerBI
- **Statistics & Methods:** A/B and hypothesis testing, time-series modeling, deep learning, classification/clustering, optimization, network analysis
- **Web:** React, Vue, Next.js, Node.js

EDUCATION

Purdue University — Krannert School of Management *2012 – 2018*

M.S., Business Analytics and Information Management (2018); B.S., Supply Chain Information & Analytics, B.S., Industrial Management, B.S., Economics (2017).