



Srivatsava Daruru

Affiliate, Program-Koyejo, O.

Bio

BIO

Srivatsava Daruru is a researcher and machine learning leader whose work spans natural language processing, neuro-symbolic AI, and large-scale learning systems. He is currently Chief AI Officer at Exlens AI and was formerly Senior Manager of Machine Learning at ServiceNow, where he led research in retrieval-augmented generation (RAG), question answering, post-training optimization of large language models, and agentic workflows for conversational AI. His contributions shaped ServiceNow's generative AI strategy, including the company's first production-grade generative application, Genius Q&A.

Daruru's research interests focus on self-improving large language models, reasoning, and mathematical verification. He is currently working on VeriBench, an end-to-end benchmark for translating Python into Lean 4, and VeriCI, a continuous verification framework for CI/CD pipelines, as part of neuro-symbolic software reliability.

He has published at leading venues such as ACM SIGKDD and IEEE ICDM, with research spanning scalable clustering for terascale astronomy, parallel data mining, and large-scale telecom analytics. His Google Scholar profile reflects a consistent track record of contributions to data mining, NLP, and applied machine learning. In addition, he is the inventor on multiple patents in NLP, fact validation, and semi-automated data labeling.

Daruru holds an M.S. in Computer Science from the University of Texas at Austin and a B.Tech. (Hons) in Computer Science from IIIT Hyderabad.

About Me (Informal)

I am a scientist and engineer working at the intersection of large language models, reasoning, and verification. My long-term vision is to build AI systems that are not only powerful but also trustworthy, capable of explaining themselves and proving their correctness. I'm especially excited about self-improving LLMs, agentic workflows, and neuro-symbolic methods that combine data-driven learning with formal verification. Currently, I'm working on VeriBench and VeriCI, projects that push AI systems toward rigorous mathematical guarantees while remaining practical for real-world development pipelines.