

Priyanshu *Shekhar* Sinha

TECHNICAL LEAD, GENAI *AI/ML Engineer architecting enterprise RAG at ElastiQ AI.*

Bangalore, India · priyanshu1996@hotmail.com · +91 7349575896 · linkedin.com/in/pripy · github.com/sinha96 · sinha96.github.io

— SUMMARY

AI/ML Engineer with *seven plus years* shipping production-grade systems. Currently leading the design and development of an enterprise RAG platform at ElastiQ AI — multi-source ingestion, hybrid retrieval, cross-encoder re-ranking, and document-level access control synchronized across cloud providers. *Google Cloud Professional Machine Learning Engineer.*

— EXPERIENCE

Technical Lead — GenAI Jan 2025 — Present

Elastiq AI · Bangalore, India

Lead the design and development of an enterprise RAG platform spanning ingestion, retrieval, ranking, and access control across cloud providers.

- Architecting **scalable distributed-database** solutions via Citus DB research and implementation for enterprise-scale workloads.
- Mentoring cross-functional team on robust **data pipelines**, complex LLM fine-tuning challenges, and enterprise deployment strategies.
- Collaborating with senior leadership on **strategic product development** and onboarded the **first enterprise client** for the Text2SQL solution.
- Evaluating emerging protocols like **MCP** for enterprise connectors.

Senior Software Engineer Sep 2024 — Jan 2025

Elastiq AI · Bangalore, India

Foundational AI engineering on Discover — ElastiQ's flagship product — covering fine-tuning, retrieval, and agent design.

- Fine-tuned **LLaMA 3.1 70B** via advanced QLoRA and LoRA techniques, reaching **85% target accuracy**.
- Built the foundational **Text2SQL framework** from scratch and shipped AI agents for natural-language to SQL plus unstructured file processing.
- Integrated multiple **Graph DB** backends offering Classic RAG, Cypher-based RAG, and GraphRAG — extending Discover with structural retrieval.
- Engineered custom datasets and **context-aware chunking** algorithms for unstructured-document RAG performance.

Senior Data Scientist May 2023 — Aug 2024

Acuity Knowledge Partners · Bangalore, India

Senior Data Scientist at a research, analytics, and technology partner to the financial services sector — asset managers, investment banks, private equity, hedge funds, and consulting firms.

- Developed a **versatile analysis dashboard** for an investment firm enabling KPI monitoring and analysis experiments — **87% reduction** in client analysis time.
- Built a web dashboard integrated with an **ESG data engine** for a top-tier investment firm, shaping their investment-strategy decisions.
- Shipped an **end-to-end RAG system** for parsing financial reports.

Associate, Data Scientist Dec 2021 — Apr 2023

TheMathCompany · Bangalore, India

- Led a **5-person team** shipping a flexible **A/B Testing Framework** for a leading pharmaceutical client to measure promotion success against revenue lift.
- Co-led a **7-person team** on a sales-anomaly flag pipeline accounting for seasonality — drove customer outreach delivering **+\$125K revenue** in Q1.
- Built an **LSTM model** forecasting customer return-to-purchase windows for supply and demand balancing — **R² 0.9012**.

Data Analyst

May 2021 – Dec 2021

Recruitment Smart Technologies · Ahmedabad, India

- Reduced new-client integration from **10 days to 3 hours** (97% improvement). Operational reports from **5 days to 30 minutes**; opportunity reports from **1 day to 30 minutes**.
- Reduced advanced-analytics report generation from **7 days to 15 minutes** (99% improvement). Automated Power BI dashboard refreshes.
- Awarded *Star Performer of the Month*.

Data Scientist

Aug 2019 – Apr 2021

Sumyag Data Sciences · Bangalore, India

- Designed a multi-phase pipeline enriching extracted document data points — generated **500+ features** via NLP, data wrangling, and text mining.
- Co-developed a **matrix-multiplication ensemble** on the NumPy stack across **10+ unique sources** for entity classification, plus custom Bayesian models for word embeddings.
- Shipped end-to-end claims-document pipelines for automobile and healthcare enterprise clients; built a geometric model for tabular-entity extraction using asymmetric graphs.

Data Science Mentor (Freelance)

Dec 2020 – May 2023

Independent · Topmate

- Mentored college students and working professionals transitioning into data science — career strategy, project review, and mock interviews for analytics and ML roles.

— SKILLS

LANGUAGES	Python, SQL, Cypher
LLM AND GENAI	LLaMA 3.1 70B , QLoRA, LoRA, SFT, RLHF, Classic RAG, GraphRAG, Cypher-RAG, Text2SQL, MCP, LangChain, Transformers, Ollama
SEARCH AND RETRIEVAL	OpenSearch , BM25, kNN, Hybrid Search, Cross-encoder Re-ranking, ChromaDB, Embeddings, ACL Sync
MULTI-MODAL	Text, PDF, Audio, Video, Speaker Diarization, OCR
MACHINE LEARNING	PyTorch, TensorFlow, scikit-learn, MLFlow, Pandas, Polars, NumPy, PySpark, NLP, Deep Learning, MLOps
CLOUD	Google Cloud Platform , AWS (S3, EKS, ECS, SageMaker, SQS, SES, EC2), GCS, Azure Blob
DATABASES	PostgreSQL, Citus DB, Snowflake, Databricks, Graph DB, ChromaDB
TOOLING	Docker, Apache Airflow, Streamlit, ChainLit, Flask, Plotly, Power BI, Postman, Git, GitHub, GitLab
SOFT SKILLS	Communication, Project Management, Team Leadership, Mentorship

— EDUCATION

Post Graduate Program — Data Science Engineering

Great Lakes Institute of Management · 2019

Bachelor of Engineering — Electrical and Electronics

Sir M. Visvesvaraya Institute of Technology · 2014–2018

— CERTIFICATIONS

Google Cloud — Professional Machine Learning Engineer

Google Cloud · Active

Natural Language Processing Specialization

DeepLearning.AI · Coursera · 2021

Deep Learning Specialization

DeepLearning.AI · Coursera · 2018

Python — Data Structures

University of Michigan · Coursera · 2018

Spark and Python for Big Data with PySpark

Coursera

Programming for Everybody

University of Michigan · Coursera

— AWARDS AND RECOGNITION

- **Star Employee of the Month** — ElastiQ AI · March 2026
- **Star Performer of the Month** — Recruitment Smart Technologies · 2021
- **Arctic Code Vault Contributor** — GitHub Open Source Archive Program
- **Pull Shark** — GitHub Achievements