

# Jeet Swadia

Boston, MA · [swadiajeet@gmail.com](mailto:swadiajeet@gmail.com) · (838) 333-0802 · [linkedin.com/in/jeetswadia](https://www.linkedin.com/in/jeetswadia)

## SUMMARY

---

Senior Data Scientist with 5+ years of experience owning end-to-end delivery of high-impact AI and data science solutions for Fortune 500 clients in regulated healthcare and insurance environments. Deep expertise in Generative AI, LLMs, agentic AI systems, and predictive ML modeling, with a strong consulting background translating ambiguous business problems into measurable AI-driven solutions. Expert in Python, SQL, and Azure cloud platforms (Azure ML, OpenAI, DevOps, AKS, Synapse). Proven ability to engage senior stakeholders, present findings with compelling storytelling, and drive adoption of data-driven strategies across complex organizations. Committed to Responsible AI: explainability, fairness, transparency, and compliance throughout all model development. M.S. in Electrical and Computer Engineering.

## TECHNICAL SKILLS

---

<b>Generative AI &amp; LLMs:</b>	Azure OpenAI, OpenAI API, AWS Bedrock, HuggingFace Transformers, LLM fine-tuning (LoRA/QLoRA), prompt engineering, RLHF/DPO, RAG architectures, LangChain, LlamaIndex, vector databases (FAISS, Pinecone), agentic AI systems
<b>ML &amp; Predictive Modeling:</b>	XGBoost, LightGBM, scikit-learn, PyTorch, TensorFlow, classification, regression, risk scoring, Bayesian modeling, A/B experimentation, statistical analysis, model evaluation and calibration
<b>Responsible AI:</b>	SHAP/LIME explainability, fairness and bias assessment, compliance guardrails for LLM outputs, audit logging, data privacy controls, model governance, regulatory compliance in regulated environments
<b>Azure Platform (Expert):</b>	Azure ML, Azure OpenAI, Azure Synapse, Azure Data Factory, Azure DevOps, AKS, Azure Monitor, Azure Blob, Azure Functions
<b>Cloud &amp; MLOps:</b>	AWS (S3, EC2, SageMaker, Lambda, Glue), GCP (BigQuery), Docker, Kubernetes, CI/CD (Azure DevOps, GitHub Actions), MLflow (experiment tracking, model registry), model monitoring, drift detection, automated retraining
<b>Data Engineering:</b>	PySpark (expert), Databricks (Unity Catalog, Delta Lake), large-scale ETL, data quality validation, schema governance, Airflow
<b>Consulting &amp; Communication:</b>	Executive stakeholder engagement, strategic AI advisory, project planning and risk management, data storytelling, presentation to C-suite, cross-functional collaboration, change management and AI adoption
<b>Languages:</b>	Python (expert), SQL (expert), Java, SparkSQL
<b>Certifications:</b>	Azure Data Engineer Associate (May 2023) · AWS Solutions Architect (Aug 2022) · TensorFlow Developer, Google (Jan 2021)

## PROFESSIONAL EXPERIENCE

---

### Senior Data Scientist — GenAI, Agentic AI & Strategic Consulting

May 2025 – Dec 2025

Client: MetLife · IBSE Model, Team Gatekeepers

- Owned end-to-end delivery of a high-impact AI solution for a Fortune 500 insurance client; collaborated with business stakeholders to define the problem, developed project plan, assessed risks, and delivered a production agentic AI system automating clinical document analysis for life insurance underwriting.
- Designed and implemented Generative AI applications integrating Azure OpenAI, HuggingFace Transformers, and LangChain; applied prompt engineering, LoRA/QLoRA fine-tuning, and RAG architectures to build context-aware multimodal document understanding pipelines.
- Built and deployed predictive and prescriptive ML models (XGBoost, LightGBM) with calibrated probability outputs; implemented SHAP/LIME-based explainability and compliance guardrails ensuring Responsible AI principles across fairness, transparency, and regulatory audit requirements.
- Integrated RLHF/DPO human feedback loops to align LLM outputs with domain expert expectations; built evaluation harnesses tracking model accuracy, safety, latency, and cost across all production workflows.
- Presented findings and solution architecture to senior leadership and compliance governance bodies using compelling data storytelling; simplified complex AI behavior into clear narratives driving stakeholder understanding and adoption.
- Deployed production AI solutions on Azure (ML, OpenAI, AKS, DevOps) with full CI/CD, model monitoring, drift detection, and automated retraining; reduced manual underwriting effort by 40% through reliable, scalable AI delivery.
- Identified opportunities to leverage Generative AI for broader business transformation; acted as trusted AI advisor to

compliance, legal, and product stakeholders throughout the engagement.

### **Lead Data Scientist — Strategic AI Consulting, Healthcare**

Mar 2024 – Apr 2025

*Client: Johnson & Johnson (Remote) · Voyager Metrics Portal*

- Led end-to-end delivery of an enterprise AI analytics platform for a Fortune 500 healthcare client as primary data science and consulting lead; defined business problem, designed solution architecture, developed project roadmap, and managed delivery across a 9-person cross-functional team in a global matrixed organization.
- Built and deployed scalable ML data pipelines on PySpark and Databricks integrating 50+ data sources; engineered systematic optimizations delivering 38% throughput improvement and maintained production reliability throughout.
- Presented strategic AI recommendations, performance metrics, and business impact analyses to Finance Directors and C-suite stakeholders; acted as trusted advisor translating data science findings into strategic decisions.
- Promoted Responsible AI practices: data governance, audit logging, access controls, and model transparency across all platform components; mentored engineers on ethical AI development standards.
- Documented best practices for AI application development; shared knowledge across the engineering and data science team to drive organizational capability building.

### **Senior Data Engineer — AI-Ready Data Infrastructure**

Aug 2022 – Feb 2024

*Client: Johnson & Johnson (Remote) · Kenvue Platform Migration*

- Architected PySpark ETL pipelines migrating 40+ regulated healthcare datasets to Databricks Lakehouse; implemented data quality validation, schema governance, audit logging, and compliance controls supporting downstream ML model training.
- Led cross-continent Agile team of 17+ engineers; enforced software engineering lifecycle best practices including version control, testing, code review, CI/CD, and production deployment standards.

### **Software Engineer**

Feb 2022 – Dec 2025

*Tata Consultancy Services (Remote)*

- Built scalable Java and Oracle Database solutions; earned AWS Solutions Architect certification; contributed to Agile cross-functional delivery teams.

### **Research Engineer — Applied AI & Statistical Modeling**

Jul 2020 – Jan 2022

*Binghamton University, Binghamton, NY*

- Conducted applied research on Reinforcement Learning and Deep Neural Networks for autonomous systems on AWS; designed statistical experiments, built evaluation benchmarks, and communicated findings to academic and applied audiences.
- Earlier (2018): Built a real-time CNN-based detection system (OpenCV, TensorFlow) on video data; designed data pipelines and model evaluation frameworks achieving 74% detection efficiency.

## **EDUCATION**

---

**M.S., Electrical & Computer Engineering** — Binghamton University, NY

Aug 2019 – May 2021

*Coursework: Reinforcement Learning, Detection Theory, Game Theory, Control Systems, Network Security*

GPA: 3.4/4.0

**B.Tech., Electrical Engineering** — Nirma University, Ahmedabad, India

Jul 2015 – May 2019

CGPA: 7.3/10