Ovi Poddar

Dhaka, Bangladesh | ovipdr.202575@gmail.com | linkedin.com/in/ovi-poddar | github.com/Ovi-Poddar
ovi-poddar.github.io/Portfolio

Research Interests

Generative & Agentic AI · Reliable & Trustworthy AI · LLM Evaluation · AI for Decision Support · Intelligent Software Engineering

Interested in developing reliable, domain-aware AI systems that combine LLMs, agentic workflows, structured knowledge, and rigorous evaluation to support decision-making in complex real-world settings.

Research Experience

Towards Equitable AI in Pathology: A Systematic Review of Fairness Challenges and Debiasing Methods in Histopathology Imaging

Collaborative Research Project | Manuscript in Preparation | March – May 2026

- Contributed to a **PRISMA 2020 systematic review** of **78 histopathology-AI studies** examining bias, fairness evaluation, and debiasing methods.
- Contributed to an **eight-category taxonomy of bias** and a proposed histopathology-specific fairness reporting framework for whole-slide image models.
- Investigated gaps in **fairness metrics, external validation, demographic metadata, and reproducibility** across pathology AI research.

Automated Classification of Code Review Feedback Feb 2022 – Apr 2023

Undergraduate Research | BUET

- Investigated deep-learning-based classification of code-review feedback using **CodeBERT embeddings and AST/change-based features**.
- Developed hybrid representations combining natural-language and program-structure information.
- Published in the **Proceedings of ESEM 2023**; Supervisor: Dr. Anindya Iqbal

Publications

Turzo, A. K., Faysal, F., **Poddar, O.**, Iqbal, A., & Bosu, A.
Towards Automated Classification of Code Review Feedback to Support Analytics.
ESEM 2023.

Paper | PDF

Professional Experience

Software Development Engineer, IQVIA – Dhaka, Bangladesh Feb 2024 – Present

- Contribute to enterprise **LLM-powered assistants for life-sciences analytics**, focusing on prompt engineering, agent behavior optimization, and LLM evaluation.
- Translate client feedback, business requirements, domain documentation, and data semantics into **system instructions, behavioral constraints, and agent configurations**.
- Design and iteratively refine prompts and agent behavior to improve **accuracy, consistency, and adherence to domain requirements**.
- Analyze LLM failure modes and edge cases and apply prompt- and configuration-level improvements.
- Work within **LangChain/LangGraph-based agent systems**, using **LangSmith** for observability, trace analysis, and evaluation.

Earlier in my tenure: Developed and maintained cloud storage, document management, authentication, and security components using **.NET Core, C#, Angular, PostgreSQL, and AWS S3**.

Software Engineer I, Enosis Solutions – Dhaka, Bangladesh June 2023 – Jan 2024

- Developed and enhanced an e-learning web application using **ASP.NET Core, C#, Angular, SQL Server, and Entity Framework Core**.
- Delivered feature enhancements, resolved complex bugs, and improved performance, scalability, and system reliability.

Selected AI Projects

Amazon Product Price Prediction — LLM & ML Engineering GitHub

LLM Engineering Capstone

Built an end-to-end text-to-price prediction pipeline from Amazon product metadata—data curation and Groq API-assisted preprocessing through classical NLP/ML baselines, deep learning, frontier LLM benchmarking, and parameter-efficient fine-tuning. Evaluated 15+ approaches; **QLoRA-tuned Llama 3.2 3B achieved the best MAE ($39.85)** among evaluated approaches.

- Curated product text from the **McAuley Amazon Reviews 2023** dataset; established baselines with vectorization, XGBoost ($68.23 MAE), and a custom **PyTorch** network ($46.49 MAE).
- Benchmarked zero-shot frontier LLMs and OpenAI fine-tuning (**GPT-4.1-nano**); fine-tuned **Llama 3.2 3B with QLoRA/PEFT** (4-bit) for the strongest result.
- Built a reusable evaluation framework with parallel inference, MAE benchmarking, and error analysis across model variants.

Agentic Chatbot — LangGraph Multi-Use AI Assistant GitHub

Built a stateful LangGraph workflow with specialized nodes for conversational reasoning, web search, and news aggregation, integrating Tavily, Groq LLMs, and Streamlit.

InfoScope AI — RAG & News Intelligence GitHub | YouTube

Built an end-to-end RAG pipeline using LangChain, FAISS, and OpenAI for citation-backed question answering over indexed articles and uploaded text.

Weather Agent — Multi-Agent MCP System GitHub

Developed a multi-agent system using LangGraph and Gemini, integrating custom Weather and Calculator MCP servers implemented in Go and Python.

Blog-Agentic — Multi-Agent Content Generation GitHub

Designed a multi-stage agentic pipeline with specialized agents for research, outlining, and drafting, coordinated through a central workflow.

Technical Training

LLM Engineering — Core AI Engineer Track | 2026

Hands-on training and implementation across LLM applications, Hugging Face/Transformers, RAG, vector databases, LoRA/QLoRA fine-tuning, deployment, and agentic AI. (GitHub)

Additional hands-on work: OpenAI Agents SDK, LangChain/LangGraph, CrewAI, MCP, tool integration, and multi-agent orchestration (GitHub).

Technical Skills

Generative AI: LLMs, Prompt Engineering, LLM Evaluation, RAG, Tool Calling, Multi-Agent Systems, Agent Orchestration, Fine-tuning

Frameworks & Tools: OpenAI Agents SDK, OpenAI APIs, LangChain, LangGraph, LangSmith, MCP, Hugging Face, Transformers, FAISS

Machine Learning: NLP, Machine Learning, Deep Learning, Scikit-learn, TensorFlow, NumPy, Pandas

Software Engineering: Python, C#, C++, .NET Core/ASP.NET, Angular, SQL

Cloud & Infrastructure: AWS, PostgreSQL, Docker, Git, Linux

Education

Bangladesh University of Engineering & Technology (BUET)

B.Sc. in Computer Science and Engineering, **CGPA: 3.61/4.00**

March 2018 – May 2023

Awards & Achievements

- **IQVIA Silver Award** — recognized for exceeding expectations, ownership, and supporting delivery and system stability.
- **BUET Dean's List** — recognized for academic performance (CGPA 3.75+ across all Level-2 semesters).
- **2,500+ competitive programming problems solved** on Codeforces, LeetCode, and other platforms.