

Kushal Subedi

Marietta, GA, USA
kushalsubedi2@gmail.com | +977 9865359307
[LinkedIn](#) | [Github](#)

PROFESSIONAL SUMMARY

Data Engineer and Researcher specializing in real-time distributed systems, Multi-Agent Systems (MAS), and blockchain analytics. Proven ability to architect production-grade pipelines processing terabytes of data with sub-second latency. Expertise in Apache Kafka, Flink, and LLM integration. Currently conducting doctoral research on agent-based software quality assurance while maintaining a strong portfolio of deployed AI solutions.

TECHNICAL SKILLS

- **Languages:** Python, Go, SQL, Bash
- **Data Engineering:** Apache Kafka, Apache Flink, DBT, StarRocks (OLAP), Redis Streams
- **AI/ML & MLOps:** PyTorch, TensorFlow, LangChain
- **Infrastructure & DevOps:** Docker, Docker Swarm, GitOps, CI/CD (Jenkins, GitHub Actions), Prefect
- **Databases:** MySQL, PostgreSQL, MariaDB, SurrealDB, Neo4j (Graph), Vector Databases (FAISS)

WORK EXPERIENCE

Graduate Research Assistant **Jan 2026 – Present**
Kennesaw State University, CCSE *Marietta, GA*

- Conducting research in Distributed Artificial Intelligence (DAI) and Multi-Agent Systems (MAS), specifically investigating agent-based approaches to automated software quality assurance.
- Designing multi-agent frameworks where specialized agents collaborate to detect syntax errors, security vulnerabilities, and logic flaws in distributed codebases.
- Collaborating with faculty to simulate agent-based workflows for code review, aiming to improve defect detection rates compared to traditional static analysis tools.

Data Engineer **Feb 2025 – Dec 2025**
108 Capital Pvt. Ltd. *Kathmandu*

- Architected production-grade Kappa Architecture pipelines processing 10+ TB of Solana blockchain data, deploying self-hosted Apache Kafka and Apache Flink clusters with exactly-once semantics.
- Engineered a scalable historical data ingestion pipeline using OldFaithful archive and Subsqid, optimizing a self-hosted StarRocks OLAP cluster to achieve 10x query performance improvements.
- Built fault-tolerant real-time block streaming infrastructure leveraging gRPC and Redis Streams for deduplication, ensuring zero data loss during provider failures.
- Developed high-throughput transformation layers using Bento (Benthos) with custom Go plugins to parse raw block data and extract transaction instructions.
- Developed modular SQL transformation pipelines with dbt Core, optimizing query performance through strategic materialization configurations (tables, views, and incremental models).
- Managed end-to-end cloud infrastructure using Docker Swarm and GitOps workflows; established a full observability stack with Prometheus and Grafana.

Associate MLOps Engineer **Feb 2024 – Feb 2025**
TAI Inc. *Shantinagar, Kathmandu*

- Conducted in-depth research on Large Language Models (LLMs) and successfully integrated their capabilities into a document relation analyzer, enhancing semantic extraction accuracy.
- Engineered robust Python scripts for efficient graph database management and complex relationship manipulation within the Neo4j environment.

- Designed and implemented automated Jenkins pipelines and custom shell scripts to streamline data engineering processes, significantly enhancing operational efficiency.
- Leveraged Pandas and NumPy libraries for comprehensive data analysis, troubleshooting, and optimization of complex data operations to ensure system reliability.
- Developed comprehensive unit test suites using Pytest to validate data pipelines and ML models, ensuring code reliability and seamless CI/CD integration.

PROJECTS

Content-Aware Image Inpainting with U-Net

[Link to Project](#)

Stack: PyTorch, OpenCV, Azure App Service

- Implemented U-Net architecture with partial convolution layers, training on 50K+ images using mixed precision (AMP).
- Achieved 28.5 dB PSNR and 0.92 SSIM; deployed as a scalable RESTful API via FastAPI on Azure.

Videolize: AI-Powered Script-to-Video Generator

[Link to Project](#)

Stack: Stable Diffusion XL, SpeechT5, Gemini API

- **Winner (GenAI Category) at Hack-A-Week 2023.**
- Developed a pipeline converting text to video using diffusion models and SLERP interpolation for smooth frame transitions.

IOUMate: P2P Financial Ledger API

[Link to Project](#)

Stack: Python, Django REST Framework, PostgreSQL, Docker, Azure

- Engineered a RESTful backend for a peer-to-peer debt management system featuring a dual-party verification engine to ensure ledger integrity.
- Optimized relational database schemas in PostgreSQL to handle multi-user transaction states (Pending, Verified, Settled) with high consistency.

Smartbot: AI Agent for 29 Card Game

[Link to Project](#)

Stack: Reinforcement Learning, ISMCTS Algorithm

- **Winner of Smart Bots Coding Challenge 2023.**
- Engineered an Information Set MCTS agent handling imperfect information for 10K+ parallel simulations per move.

EDUCATION

PhD in Computer Science

Jan 2026 – Present

Kennesaw State University

Marietta, GA

Bachelors Degree in Computer Engineering | 3.84 GPA

2024

Pulchowk Engineering Campus, Tribhuvan University

Pulchowk, Lalitpur

CERTIFICATIONS

- **Deep Learning Specialization (Coursera):** Neural Networks, CNNs, Sequence Models, Hyperparameter Tuning.
- **Security Operations Fundamentals (Logpoint):** Threat Modeling, Incident Response, SIEM/SOAR.

HONORS & LEADERSHIP

- **Scholarships:** Golden Jubilee Scholarship (2019), Mahatma Gandhi Scholarship (2017) by Indian Embassy.
- **Instructor:** Taught Python/Django to 300+ students (PDSC) and ML Foundations to 100+ students (LOCUS).