Muskan Sindhu

Gurugram, India | +91 8797686834 | shunyamuskan@gmail.com

muskansindhu | muskansindhu | muskansindhu | muskansindhu.in

Education

Indian Institute of Technology, Madras — Tamil Nadu, India
B.S. in Data Science and Applications — *2021 – 2025*

Kalindi College, Delhi University — Delhi, India
Bachelor of Vocation in Web Designing — *2019 – 2022*

Experience

TAOlie.ai – Decentralized Compute Marketplace
Software Engineer — *July 2024 – Aug. 2025*

- Engineered a high-throughput distributed data ingestion pipeline processing **24M+ records/day**; collected, normalized, and streamed multi-source data into the Bittensor network.
- Designed a queue-backed **batch processing system** with async worker pools and adaptive rate limiting, reducing operational costs by **96%.**
- Built a **geo-aware** account warm-up and rotation framework with intelligent **request routing** to preserve account health and prevent platform bans across regions.
- Built a cross-platform **search engine** that rewrote user queries into platform-specific API calls, applying **LLM-based re-ranking** to improve contextual relevance.
- Engineered a **containerized** miner deployment system for **Bittensor mainnet**, automating wallet creation, subnet registration, and horizontal scaling workflows.
- Implemented a **Prometheus and Grafana** observability stack tracking incentives, request latency, and error rates.

Deepgaze Technologies
Backend Engineer (Intern) — *July 2023 – Sep. 2023*

- Developed **REST APIs** to connect React UI with in-house TensorFlow LSTM models for real-time ML inference.
- Built a deepfake source identification pipeline using reverse image search APIs to trace manipulated image origin.
- Contributed to data preprocessing pipelines and **RLHF** fine-tuning workflows to improve model training accuracy.

Projects

TicketShow | *FastAPI, Kafka, PostgreSQL, Elasticsearch, MinIO, Docker, React* — Source Code

- Built a **distributed** ticket-booking platform with 6 **microservices** behind an API Gateway; load-tested using k6 constant-arrival traffic, sustaining ∼50 RPS with p95 latency < 500ms.
- Decoupled booking, payment, and notification flows using an **Apache Kafka** event bus, enabling each service to be deployed and scaled independently.
- Eliminated double-booking under **high concurrency** via database-level locking; used **idempotency** keys to ensure payment retries never double-charge users.
- Shipped fuzzy full-text search over shows with **Elasticsearch**, synced in real time from the write database via Kafka events, achieving a **CQRS-like** read/write separation.
- Generated **QR-coded tickets** on booking confirmation; stored all binary assets (posters, tickets, invoices) on **S3-compatible object storage**, keeping media decoupled from app logic.

Clipnote | *Flask, PostgreSQL, Gemini AI, AWS S3, Docker, Nginx, Chrome Extension API* — Link | Source Code

- Built a **Chrome extension** enabling users to take timestamped notes (user or AI-generated) for YouTube videos.
- Developed Flask backend APIs, styled the UI with HTML/CSS, and deployed via **Docker** on an EC2 instance.
- Configured **AWS S3** for storing video transcripts to enable AI notes and optimize transcript API call frequency.

Technical Skills

Languages: Python, JavaScript, SQL, Bash
Frameworks and Libraries: FastAPI, Flask, Node.js, Express.js, Asyncio, SQLAlchemy, Pydantic, React
Databases: PostgreSQL, MongoDB, SQLite, Supabase
Messaging and Search: Apache Kafka, Elasticsearch
Cloud and DevOps: AWS (EC2, S3), MinIO, Docker, Nginx, Prometheus, Grafana, GitHub Actions, Cloudflare Tunnel
Developer Tools: Git, GitHub, Postman, Linux