

Kumar Naman

Bhubaneswar · talktonaman@duck.com · LinkedIn · GitHub

PROFESSIONAL SUMMARY

Full-stack developer and AI engineer with 2.5+ years of experience building production AI systems, microservices platforms, and scalable infrastructure. Strong system design fundamentals — experienced in architecting distributed systems, agentic workflows, and data-intensive applications from the ground up. Specialized in integrating LLMs, computer vision, and speech models into real-world systems with robust deployment pipelines. Proven track record of delivering AI-powered solutions for industrial operations, mining safety, and government sectors.

EXPERIENCE

VyomChara Pvt. Ltd. — SDE

Hybrid

Jan 2024 – Present

- Developed end-to-end hyperspectral analysis pipeline for mineral classification in lab and field based Ore Body Mapping; integrated with a remotely controlled HSI kiosk machine using MQTT based IoT system.
- Built computer vision pipeline APIs for real-time vehicle/personnel/hazard detection in mining operations; integrated object storage and relational database for media/metadata management
- Architected an AI-powered case management platform with agentic investigation workflows — autonomous task delegation across specialized agents for document processing, evidence analysis, and legal vetting; multi-role approval pipelines with self-correcting response generation
- Developed an on-premise RAG system combining vector database retrieval with full-text indexing for text/image Q&A; built pluggable service layer supporting multiple cloud and local providers with multimodal analysis and web search grounding
- Built a real-time 3D AI avatar interface with speech-to-text and text-to-speech integration; optimized GPU memory management for large models to reduce latency
- Set up and managed a private AI HPC with Kubernetes clusters for ML workloads; designed backend systems supporting drone telemetry and low-latency video streaming over 5G networks
- Built deployment pipelines and dev environments for startups incubated under the TEXMiN foundation; engineered company website achieving 90+ Lighthouse score with 99.9% uptime

TEXMiN, IIT ISM Dhanbad — Software Engineering Intern

Dhanbad

May 2025 – Jul 2025

- Benchmarked multiple deep learning architectures for mineral segmentation, achieving 99%+ validation accuracy on benchmark datasets
- Engineered data preparation pipeline with dimensionality reduction, noise filtering, and data augmentation to address class imbalance in real-world datasets
- Operationalized models via REST APIs supporting data upload, classification runs, and high-throughput batch processing for industrial use cases
- Architected microservices backend for hyperspectral image analysis; containerized 5 specialized services with GPU detection and memory management
- Integrated object storage with presigned URLs and PostgreSQL for analysis session tracking across multi-tenant architecture; built React frontend with responsive scientific data visualization interfaces

TEXMiN, IIT ISM Dhanbad — Project Intern

Dhanbad

Aug 2023 – Jan 2024

- Analyzed 10+ academic papers on hyperspectral mineral classification, evaluating model performance on benchmark datasets
- Authored comprehensive comparison report of deep learning approaches for mineral classification, identifying top-performing architecture as 15% more accurate
- Mastered hyperspectral camera operation, capturing and processing high-resolution spectral data from 7+ mineral types
- Transformed raw spectral data into reflectance curves, developing preprocessing techniques that improved analysis accuracy by 30%

PROJECTS

Multilingual Conversational AI System

- Building a dual-pipeline system for code-switched speech with fine-tuned speech recognition and parameter-efficient language model for structured output generation
- Applied parameter-efficient fine-tuning techniques on consumer GPU hardware, reducing trainable parameters by 99%+ and memory footprint to under 8GB
- Generated synthetic training data using translation models; defined domain-specific accuracy benchmarks and hallucination checks

Enhanced SRGAN Pipeline

- Architected custom Super-Resolution GAN with residual blocks, achieving 4x image enhancement with PSNR score above 25
- Combined pixel-wise MSE, perceptual loss, and adversarial loss, outperforming traditional upscaling methods by 35%

3D CNN Mineral Classifier

- Developed CNN to classify multiple mineral categories with 91% F1-score, optimizing model using standard training techniques

TECHNICAL SKILLS

Languages:	Python, C++, JavaScript/TypeScript, HTML5, CSS3
Full-Stack:	React, Node.js, REST APIs, SSE, Microservices, Docker, Docker Compose, Kubernetes
System Design:	Distributed Systems, API Architecture, Data Modeling, Scalability Patterns, Event-Driven Design
ML & AI:	TensorFlow, PyTorch, scikit-learn, GANs, CNNs, YOLO, OpenCV, LangChain, Agentic AI, Multi-Agent Systems, LLM Integration, RAG Systems, ASR/TTS
DevOps & Cloud:	Git, GitHub Actions, AWS, Vercel, CI/CD Pipelines, Container Orchestration
Platforms:	Linux, NVIDIA Jetson, Raspberry Pi, PX4/ArduPilot UAV stacks

EDUCATION

Kalinga Institute of Industrial Technology (KIIT) — B.Tech CSE	Bhubaneswar, India
Sep 2022 – Jul 2026 (expected)	CGPA: 7.3/10.0