

Ankit Kumar

+91 9306686556 | nasaankit01@gmail.com | Github: @anky209e | India | Website

ABOUT

Highly skilled AI/ML Engineer and Backend Developer with hands-on experience in building and deploying production-grade Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG) systems. Proven ability to translate complex data science problems into scalable software solutions, demonstrated through freelance consulting and impactful internships at organizations like ISRO, where I optimized autonomous systems using Reinforcement Learning. Seeking challenging roles as an **AI/ML Engineer**, **Data Scientist**, or **Software Developer**.

EXPERIENCE

Freelance AI/ML Engineer | Consultant

July 2024 – Present | Remote

- Designed and deployed 4+ production-grade LLMs and RAG systems for domain-specific applications, achieving 20% improvement in retrieval performance.
- Built and fine-tuned transformer-based Large Language Models (LLMs), including training a 124M parameter model from scratch.
- Designed and implemented end-to-end conversational AI solutions that scaled to handle 5,000 daily user interactions.
- Reduced harmful output instances by 95% through implementation of AI safety guardrails and advanced prompt engineering techniques.

ISRO - Indian Space Research Organisation | Project Intern

May 2024 – June 2024 | Chennai, India

- Optimized autonomous Lunar landing simulations using Reinforcement Learning agents, reducing trajectory error margin by 15% over baseline heuristics.
- Engineered Range Operations data pipelines to process 10,000+ data points per second, enabling real-time fault forecasting with 90% accuracy.
- Designed a low-level hardware processing layer in Python/C to increase telemetry throughput by 2X and condition GBs of multi-sensor telemetry for RL training loops.
- Collaborated closely with a team of distinguished Scientists to prototype mission-critical software frameworks.

Ecowiser | Backend Developer Intern

Nov. 2022 - Apr. 2023 | Hyderabad, India

- Designed and scaled production backend services for a core Chrome Extension and Web ecosystem supporting 10,000+ active users utilizing Django Rest Framework.
- Accelerated internal search query speeds by 30% through the integration of custom NLP tokenization into PostgreSQL Full-Text Search.
- Built an asynchronous, automated customer engagement pipeline utilizing Redis and Celery, processing over 5,000 scheduled feedback dispatches daily.

Velozity Global Solutions | Machine Learning Intern

June 2022 - Oct. 2022 | Chennai, India

- Developed and cross-validated predictive deep learning frameworks using PyTorch and TensorFlow across 15GB+ unstructured datasets.
- Designed and benchmarked LSTM-based architectures, improving accuracy by 14% in identifying long-term temporal dependencies within sequential data logs.

PROJECTS

Tensile Analysis Platform

Mar. 2026 – May 2026

- Designed and deployed a high-performance material characterization dashboard used by a University of Stuttgart research laboratory, automating 90% of complex stress-strain analysis.
- Built custom mathematical physics solvers, including an optimized Young's modulus window-search engine using cumulative sum statistics and a strict 0.2% yield strength offset script.
- Implemented fuzzy-token matching to integrate 100+ unstructured laboratory spreadsheets with raw data, embedding automated statistical analysis (ANOVA, t-tests, IQR outlier detection).

ModelSmith

June 2022 – July 2022

- Developed a unified model-evaluation web application in Django hosting 10+ distinct machine learning architectures, including custom ResNets and Generative Adversarial Networks (GANs).
- Authored thorough technical specifications and documentation detailing the underlying linear algebra, backpropagation mechanics, and production deployment parameters.

EDUCATION

Lovely Professional University

Bachelor of Technology in Computer Science and Data Science

Oct. 2020 – May 2024 | Punjab, India

- Relevant Coursework: Data Structures, Algorithms, Databases, Computer Systems, Machine Learning, Artificial Intelligence

CERTIFICATIONS

Reinforcement Learning (NPTEL)

Jan - Apr 2022 | NPTEL, IIT Madras

Python for Data Science (Elite)

Jan - Feb 2022 | NPTEL, IIT Madras

SKILLS

- Languages:** Python, SQL (PostgreSQL, MySQL), Ruby
- AI/ML:** PyTorch, TensorFlow, LangChain, LlamaIndex, Reinforcement Learning, Generative AI
- Frameworks:** Django (DRF), Flask, Pandas, NumPy, Scikit-learn, Fastapi, Streamlit, SciPy
- Tools/DevOps:** Docker, Git, Linux, Cloud Platforms, CI/CD, NLP