Dev Chandra Adhikari

Kathmandu, Nepal
adkdev200@gmail.com

PROFESSIONAL SUMMARY

Highly driven Computer Engineering student at Thapathali Campus (IOE, Tribhuvan University) with deep hands-on expertise in Machine Learning, Deep Learning, Natural Language Processing (NLP), and Generative AI. Proven track record of designing and deploying sophisticated AI-powered systems — from intelligent recommendation engines and real-time computer vision pipelines to full-stack social media platforms and Retrieval-Augmented Generation (RAG) applications. Proficient in leveraging cutting-edge tools including Hugging Face Transformers, Model Context Protocol (MCP) for agentic AI workflows, and FAISS-powered vector search. Passionate about building scalable, production-grade applications that push the boundaries of what AI can deliver.

CORE SKILLS

Machine Learning & Natural Language Processing (NLP)

- Designed and deployed end-to-end ML pipelines — Recommendation Systems, Classification, Clustering, and Regression — using Scikit-learn, XGBoost, and custom ensemble methods with consistent top-tier performance.
- Built advanced NLP pipelines leveraging Hugging Face Transformers, NLTK, and spaCy — covering tokenization, sentiment analysis, named entity recognition, and semantic similarity at scale.
- Expert in feature engineering, hyperparameter tuning, cross-validation strategies, and model interpretability (SHAP).

Deep Learning & Computer Vision

- Architected and fine-tuned deep CNN models (ResNet, VGG, EfficientNet) achieving high accuracy on complex medical imaging and facial recognition tasks.
- Built Brain Tumor Detection, Doppelganger Finder, and real-time image processing systems from scratch with robust preprocessing, augmentation, and inference pipelines.
- Proficient in PyTorch and TensorFlow — model training, transfer learning, embedding extraction, and production deployment via REST APIs.

Generative AI, Agentic AI & Model Context Protocol (MCP)

- Integrated and orchestrated large language models (LLMs) via API — building context-aware, multi-step reasoning agents using Model Context Protocol (MCP) for seamless tool and data access.
- Engineered Retrieval-Augmented Generation (RAG) pipelines combining LLMs with FAISS vector search for intelligent document and multimedia Q&A systems.
- Experienced in advanced Prompt Engineering and Prompt Injection defense — building robust, safe, and high-performance AI applications.
- Developed Chat with YouTube Videos — a real-time multimodal RAG system that extracts, indexes, and answers questions from video content using vector embeddings.

Backend & Full-Stack Application Development

- Built and deployed production-ready REST APIs and WebSocket-based real-time systems using Django, Django REST Framework, and FastAPI.
- Developed a fully featured Social Media Application — including user authentication, posts, likes, real-time messaging, and notification systems.
- Expert in JWT authentication, Redis caching, PostgreSQL optimization, and cloud-ready deployment on both Windows and Linux environments.

PROJECT EXPERIENCE

Movie Recommendation System

Python • Scikit-learn • Collaborative Filtering • FastAPI

- Engineered a high-performance recommendation engine using collaborative filtering, matrix factorization, and content-based similarity — delivering highly personalized results.
- Designed scalable REST API backend with caching and real-time response for thousands of concurrent users.
- Applied advanced ensemble ML algorithms that significantly outperformed standard baseline models.

Brain Tumor Detection System

PyTorch • CNN • Transfer Learning (ResNet) • REST API

- Built a clinical-grade CNN classifier on MRI scans using ResNet transfer learning — achieving high diagnostic accuracy across multiple tumor categories.
- Designed complete data pipeline: preprocessing, augmentation, class balancing, and model evaluation.
- Deployed via REST API enabling seamless integration with medical imaging interfaces.

Doppelganger Finder

PyTorch • FaceNet • FAISS • Deep Embeddings • REST API

- Developed a real-time facial similarity engine using deep embedding extraction and FAISS-powered approximate nearest-neighbor search.
- Achieved near-instant matching across large face databases with cosine similarity ranking.
- Exposed via REST API supporting image upload and ranked similarity retrieval.

Chat with YouTube Videos — RAG System

LLM APIs • RAG • FAISS • Hugging Face • Vector Search

- Architected a full RAG pipeline that transcribes, chunks, and embeds YouTube video content into a FAISS vector store for intelligent conversational Q&A.
- Integrated Hugging Face Transformers for embedding generation and LLM APIs for response synthesis.
- Enabled multimodal context-aware conversation — users interact naturally with any YouTube video.

Social Media Application

Django • DRF • WebSockets • JWT • Redis • PostgreSQL

- Designed and built a full-featured social media backend supporting user profiles, posts, follows, likes, comments, and real-time messaging via WebSockets.
- Implemented secure JWT authentication, Redis-based session caching, and optimized PostgreSQL queries for high-throughput performance.

EDUCATION

Bachelor of Engineering in Computer Engineering

Thapathali Campus, Institute of Engineering (IOE) — Tribhuvan University, Nepal

TECHNICAL STACK

Python • PyTorch • TensorFlow • Scikit-learn • Hugging Face Transformers • NLTK • spaCy • Model Context Protocol (MCP) • Agentic AI • RAG • FAISS • LLM APIs • Prompt Engineering • Django • Django REST Framework • FastAPI • WebSockets • JWT • PostgreSQL • Redis • Windows • Linux