

Tanvirul Islam

Senior Software Engineer · AI / ML · MCP & AI Agents

Full-stack software engineer and AI/ML practitioner building production-grade AI products — AI agents, MCP, LLM fine-tuning, and model training — on scalable, well-engineered platforms.



📍 Uttara, Dhaka, Bangladesh 📞 +880 1988 067463 ✉ tanvir1ariyan@gmail.com 🌐 [linkedin.com/in/tanvirul-islam-software-engineer](https://www.linkedin.com/in/tanvirul-islam-software-engineer)

🔗 github.com/TanvirIslam-BD

PROFILE

Senior Software Engineer and AI/ML practitioner with 9 years of experience designing and delivering scalable web applications, robust backend systems, and intelligent, AI-powered products — spanning Model Context Protocol (MCP) integrations and autonomous AI agents. Deep expertise across Java, Python, TypeScript, and the Grails, FastAPI, Django, and Flask frameworks, with a strong track record of architecting maintainable systems and optimizing complex, high-traffic applications.

Specializes in applied machine learning, deep learning, and NLP, with hands-on experience across the full LLM lifecycle — instruction tuning, reward modelling, RLHF (PPO), and DPO — and in designing agentic systems and Model Context Protocol (MCP) integrations that connect LLMs to real-world tools and workflows. Passionate about turning cutting-edge AI research into dependable, high-performance solutions that deliver measurable impact.

EXPERIENCE

WebAlive

Sep 2023 – Present · 2 yrs 9 mos

Sr. Software Engineer

Remote · Melbourne, VIC, Australia

- Core contributor to [WebCommander](#), a CMS and web platform product — implementing core features such as product variations, product widgets, and order placement — working across Grails-based backend services and REST API implementation, with Java, Python, JavaScript, and React.
- Architected and built the WebCommander MCP Server, enabling websites to be deployed and managed directly through Claude chat, Claude Code, Figma Make, or any AI application via Model Context Protocol tools over STDIO and Streamable HTTP transports.
- Built plugin infrastructure for CMS plugin developers, including a [marketplace](#) plugin installation and management system, similar to the Shopify and Wix apps/plugins systems.
- Developed AI-powered product features — a generative AI email/template generator, an AI writing assistant, and an AI engine that converts design HTML into platform widgets — applying LLM and NLP techniques to real product workflows.
- Evaluated and supervised-fine-tuned foundation models on AWS Bedrock (including Llama 3.2 3B) for an AI-assisted store-migration field-mapping task, and trained models on an internal server.
- Resolved complex deployment and infrastructure challenges, improving reliability and performance of the content and page-deployment pipeline.

Bit Mascot

Feb 2017 – Sep 2023 · 6 yrs 7 mos

Software Engineer

Dhaka, Bangladesh

- Developed front-end editor and platform tooling (layer-tree UI, slider/hero management panels) and collaborated closely with cross-functional teams to ship production features.
- Designed and delivered full-stack web applications and backend services using Java, Grails, Python, and JavaScript across a wide range of client projects.
- Built scalable, maintainable systems and led the troubleshooting and performance optimization of complex, high-traffic applications.

CERTIFICATIONS


- Build AI Agents using MCP [↗ VERIFY](#)
- Generative AI Advanced Fine-Tuning for LLMs [↗ VERIFY](#)
- Building Generative AI-Powered Applications with Python [↗ VERIFY](#)
- IBM AI Developer Specialization [↗ VERIFY](#)
- Machine Learning with Python [↗ VERIFY](#)
- Developing AI Applications with Python and Flask [↗ VERIFY](#)

SELECTED OPEN-SOURCE PROJECTS

- **MCP Remote Suite** [↗ GITHUB](#)
Production MCP suite over Streamable HTTP — FastMCP server, roots security boundary, sampling, and an AI host.
- **Instruction Fine-Tuning of LLMs with LoRA** [↗ GITHUB](#)
Parameter-efficient instruction tuning with LoRA via TRL's SFTTrainer, evaluated with SacreBLEU.
- **Reward Model Training (RLHF)** [↗ GITHUB](#)
GPT-2 + LoRA reward model trained on chosen vs. rejected preference pairs with a Bradley-Terry pairwise loss.
- **Permission-Aware MCP Security System** [↗ GITHUB](#)
Server-enforced MCP permissions: risk classification, sandbox confinement, approvals, and audit logging.
+ 12 more open-source AI / ML projects — see the full list below.

TECHNICAL SKILLS

CORE FOCUS



AI / ML / MCP / AGENTS • GITHUB PROJECTS

A portfolio of 16 open-source projects spanning LLM fine-tuning, RLHF, MCP & agents, generative-AI applications, and applied ML.

LLM Fine-Tuning & RLHF

- **Instruction Fine-Tuning of LLMs with LoRA** [↗ GITHUB](#)
Parameter-efficient instruction tuning of a causal LM on CodeAlpaca-20k using LoRA via TRL's SFTTrainer, with completion-only loss masking and SacreBLEU evaluation.

- **Reward Model Training (RLHF)** [GITHUB](#)

Trains a GPT-2 + LoRA reward model on chosen vs. rejected preference pairs with TRL's RewardTrainer and a Bradley-Terry pairwise loss; evaluated by win rate.

- **RLHF with Proximal Policy Optimization (PPO)** [GITHUB](#)

Fine-tunes GPT-2 with PPO using a sentiment classifier as the reward model and a KL-controlled reference policy — the reinforcement-learning stage of RLHF.

- **Direct Preference Optimization (DPO) Fine-Tuning** [GITHUB](#)

Aligns GPT-2 directly on (prompt, chosen, rejected) preference pairs using Hugging Face TRL's DPOTrainer + LoRA, removing the need for a separate reward model.

MCP & AI Agents

- **MCP Remote Suite** [GITHUB](#)

Production-ready Model Context Protocol suite over Streamable HTTP, with a FastMCP server, filesystem roots security boundary, sampling, and an AI host that drives the tools.

- **Permission-Aware MCP Security System** [GITHUB](#)

Server-enforced MCP permission management with per-tool risk classification, sandbox path confinement, single-use human approval, and tamper-resistant audit logging.

- **AI Writing Assistant Agent** [GITHUB](#)

LangGraph agent that analyzes text and autonomously selects and chains writing-improvement tools, served through a Django web interface.

Generative AI & NLP Applications

- **PDF Q&A Chatbot with LLMs (RAG)** [GITHUB](#)

Document-aware chatbot that answers questions over PDF content using LangChain and Flask, applying retrieval-augmented generation.

- **NLP Chatbot with Transformers** [GITHUB](#)

Conversational chatbot built on Facebook's BlenderBot with Hugging Face Transformers and PyTorch, served via a Flask web app.

- **AI Photo Captioning App** [GITHUB](#)

Image captioning application using the BLIP vision-language model from Hugging Face Transformers, with a Gradio user interface.

- **Emotion Detection Web Application** [GITHUB](#)

Detects nuanced emotions from text using Watson NLP, deployed as a Flask web application with formatted, programmatic output.

- **AI Sentiment Analysis Web App** [GITHUB](#)

Sentiment analysis web application built on embeddable AI libraries and served with Python and Flask.

- **Transformer Text Generation with TensorFlow** [GITHUB](#)

Builds a Transformer model from the ground up in TensorFlow for human-like text generation, covering attention and contextual sequence modelling.

Applied ML Engineering

- **Fraud Detection with LLaMA Embeddings + Hoeffding Tree** [GITHUB](#)

Hybrid fake-order detector combining LLaMA-based text feature extraction with an incremental HoeffdingTreeClassifier for real-time, high-throughput classification.

- **ML Model Trainer API (Django)** [GITHUB](#)

Django REST API with JWT authentication for training machine-learning models and serving real-time predictions over HTTP.

- **AI Personalized Recommendation Engine** [GITHUB](#)

eCommerce recommender combining collaborative and content-based filtering (Surprise, Scikit-learn) for real-time, personalized product suggestions.

EDUCATION

B.Sc. in Computer Science & Engineering

Uttara University

2017