

ML engineer with 2+ years building production systems—pipelines, Dockerized inference services, and tracked experiments that teams actually use. Published a 1st-author conference paper, achieving a 56% inference speedup. Ranked Top 1% on Kaggle (4,082 competitors). Built a recommender that lifted client sales by 10% in 90 days. Founded Toolly — a live AI tools platform with 400+ tools and real users, shipped solo.

SKILLS

Programming & Data: Python · SQL · Java

Modeling & ML/DL: TensorFlow · scikit-learn · XGBoost · LightGBM · CatBoost

MLOps & Deployment: MLflow · ZenML · Docker · FastAPI · Flask · Streamlit

Generative AI: LLMs · RAG · LangChain · LangGraph · Prompt Engineering

PUBLICATION

Bangla Diarizz: Domain-Adapted Speaker Diarization for Bengali Long-Form Audio | 1st Author

BUET CSE Fest 2026 — DL Sprint 4.0 Bengali Speaker Diarization Competition

- Fine-tuned segmentation model on competition dataset; replaced speaker embeddings with WeSpeaker ResNet34-LM for targeted domain adaptation over language-agnostic baselines.
- Achieved 56% wall-clock speed improvement—reduced inference time from 1h 22m to ~36m across 14 test audio files. Competitive DER on Bengali-Loop benchmark.

EXPERIENCE

Founder & Full-Stack ML Engineer—Toolly · toolly.tech

Jun 2025—Present

- Architected and launched toolly.tech solo, a live AI tools directory with 400+ tools, 15 categories, a free community submission model, and an integrated Learn AI hub.
- Engineered full product stack: frontend (HTML/CSS/JS), tool submission pipeline, moderation system, search/filter logic, and analytics—206+ commits over 10 months.
- Built and deployed Toolly Studio — Streamlit + Bria AI image generation app with batch export and Docker deployment, bridging ML and non-technical users.

Freelance Machine Learning Engineer—Self-employed

- Built a hybrid recommendation system (collaborative filtering ALS + content-based TF-IDF embeddings) for a local business; Flask API deployment produced +10% sales lift in 90 days.

Kaggle Master | Open-Source Contributor — **Kaggle** / **GitHub**

- Ranked Top 1% — Predicting Road Accident Risk (29/4,082 participants); Top 2% — Predicting BPM of Songs (38/2,581 participants).
- Completed 22 competitions, hosted competitions to engage the data science community, and published ML notebooks and datasets.

PROJECTS

Prices Predictor System | [GitHub] | ZenML, MLflow, XGBoost, LightGBM, Docker, FastAPI

- Built end-to-end ML pipeline (ingest → preprocess → train → evaluate → register → serve) with ZenML orchestration and MLflow experiment tracking for reproducible production deployment.
- Dockerized FastAPI inference service with CV and hyperparameter tuning — production-grade, not a notebook.

Toolly Studio | [GitHub], Preview | Streamlit, Bria AI, Python, Docker

- Developed a demo-ready Streamlit app that generates professional product ads by integrating Bria AI image-generation APIs and templated prompts.
- Implemented client-preview and batch-export workflows and a one-command Docker demo to make the tool easy to share with non-technical users.

EDUCATION

B.Sc. Computer Science & Engineering — BNIST

Feb 2023 — Present

- **Relevant coursework:** Linear Algebra · Calculus · Probability & Statistics · Data Structures & Algorithms · Operating Systems · Database Systems · System Design Projects · Artificial Intelligence · Machine Learning · Data Science

Languages

- English (Fluent) · Bengali (Native) · Hindi (Conversational)