

SAAD ZAIDI

AI & Software Engineer | BS Computer Science — FAST NUCES
saadzaidi129@gmail.com • 0330-2740209 • github.com/SaadZaidi1 • linkedin.com/in/saad-zaidi

PROFESSIONAL SUMMARY

Dean's List student at FAST NUCES with hands-on experience building intelligent ML systems and production-grade server-side services. Proven ability to design deep learning pipelines — from raw sensor data to deployable models — using PyTorch, scikit-learn, and Hugging Face, while delivering scalable RESTful endpoints, full-stack SaaS products, and containerized deployments. Shipped complete end-to-end solutions from database schema to payment processing to LLM-powered automation.

EDUCATION

FAST National University of Computer and Emerging Sciences — Bachelor of Science in Computer Science | Dean's List – 3rd Semester 2023 – 2027

TECHNICAL SKILLS

Languages: Python, JavaScript, C/C++, Assembly

Backend: FastAPI, Flask, RESTful APIs, JWT Auth, Celery, Redis, SQLAlchemy async, Alembic, Pydantic

ML / AI: PyTorch, TensorFlow, Keras, scikit-learn, XGBoost, Hugging Face Transformers, OpenAI API, SHAP, Streamlit

Databases: PostgreSQL, SQLite3, MySQL, Supabase

DevOps & Tools: Docker, Kubernetes, Railway, Render, Vercel, GitHub Actions, Linux, Git

Frontend: React, TypeScript, Vite, HTML, CSS, JavaScript

PROJECTS

Anchor AI — ADHD Accountability Agent | FastAPI, React, PostgreSQL, OpenAI, Celery, Redis, Paddle 2026

- Architected and shipped a full-stack SaaS ADHD productivity agent delivering personalized task management via WhatsApp — complete with Google OAuth, phone OTP verification, Paddle subscription billing, and a dark neural-themed React dashboard.
- Designed an 18-table PostgreSQL schema with async SQLAlchemy, Alembic migrations, and Supabase hosting; implemented row-level cascade deletes, partial indexes, and JSONB columns for flexible preference storage.
- Built a GPT-4o Mini LLM pipeline for intent classification, task extraction, and micro-step generation; engineered Celery Beat scheduling to fire timezone-aware morning check-ins, task reminders, and EOD summaries across all users simultaneously.
- Integrated Paddle Billing with HMAC-SHA256 webhook signature verification, idempotent event processing, and sandbox-to-production subscription lifecycle management (activated, paused, cancelled, resumed).
- Achieved 88% gross margin unit economics at scale (\$1.14/user/month cost vs \$9.99 revenue) by selecting WhatsApp Cloud API over SMS, reducing per-user messaging cost by 93%.

Intelligent CNC Process Monitoring & Defect Prediction | PyTorch, XGBoost, scikit-learn, SHAP 2026

- Developed a 3-layer ML pipeline: LSTM Autoencoder for anomaly detection (F1 = 0.9316, +43.3% over baseline), 1D-CNN + XGBoost for tool-wear classification (87.95% accuracy, F1 = 87.86%), and stacked BiLSTM + Bahdanau Attention for RUL prediction (R² = 0.5930, +21.9%, MAE = 0.1343).
- Extracted 1,000+ FFT, statistical, and mechanical features from 18 GB of sensor data across three CNC datasets; applied SMOTE oversampling for class balancing and SHAP TreeExplainer for full feature-level explainability.

Crowd-Funding Platform | Flask, FastAPI, PostgreSQL, Docker, JavaScript — **Backend Lead** 2025

- Architected a production-grade RESTful service layer with ACID-compliant transactions, role-based access control, and input validation — sustaining 100+ concurrent campaign operations with zero data integrity failures.
- Designed a RAG-based FAQ microservice using vector embeddings in FastAPI, containerized with Docker; reduced inbound support query volume by 40% through intelligent automated resolution.

LLM Engineering | Python, Jupyter, Hugging Face, OpenAI API 2025

- Conducted prompt-engineering experiments across GPT-4 and open-source LLMs, achieving 30%+ improvement in task-specific accuracy; delivered modular production-ready applications for text generation, summarisation, and classification.

CERTIFICATIONS

LLM Engineering — Udemy (Ed Donner) · 2024