

RAMZI ZEINEDDINE

Cambridge, MA 02141 | (347) 477-8159 | zeineddine.r@northeastern.edu | [LinkedIn](#) | [GitHub](#)
U.S. Citizen | Open to Full-Time, Part-Time, and Internship Opportunities

EDUCATION

Northeastern University, Khoury College of Computer Sciences Boston, MA
Master of Science in Artificial Intelligence — Concentration: Machine Learning Sept. 2025 – Expected May 2027
Relevant Coursework: Foundations of AI, Algorithms, Deep Learning, Machine Learning, Natural Language Processing, Computer Vision

American University of Beirut Beirut, Lebanon
Bachelor of Science in Computer Science — GPA: 3.5/4.0 Graduated June 2024

TECHNICAL SKILLS

Languages: Python, Java, C++, C#, TypeScript, SQL, HTML/CSS, .NET

ML & AI: PyTorch, TensorFlow, Scikit-learn, XGBoost, NumPy, Pandas, SHAP, Optuna, OpenAI API, Gemini API, RAG, Vector Search, Prompt Engineering, LLM Evaluation, Neural Networks, CNNs, Feature Engineering, Cross-Validation

Frameworks & Tools: React.js, React Native, Next.js, Angular.js, Node.js, TailwindCSS, REST APIs, Git/GitHub, Docker, Firebase, GCP, Cursor, Claude Code

Databases: SQL Server, MySQL, MongoDB, PostgreSQL, Firestore

WORK EXPERIENCE

Software Design Consulting Group — Beirut, Lebanon (Remote) Sept. 2024 – Present
Full-Stack Software Engineer

- Translated ERP business workflows across finance, inventory, and reporting modules into scalable AngularJS, C#.NET, and SQL Server solutions, improving process visibility and operational efficiency.
- Integrated OpenAI and Gemini APIs into AI-enabled support workflows, applying prompt engineering and NLP techniques to automate repetitive tasks and improve workflow responsiveness.
- Partnered with stakeholders to clarify requirements, explain technical tradeoffs, test releases, and deliver scalable solutions aligned with operational needs.

Saugo360 — Beirut, Lebanon Jul. 2023 – Oct. 2023
Software Engineer Intern

- Built a CRM healthcare dashboard using React.js, Node.js, and Syncfusion API, consolidating clinician workflows into a unified web application.
- Converted Figma designs and workflow requirements into responsive React components, improving usability through routing, lazy loading, and code splitting.
- Collaborated with QA and engineering during Agile sprints to triage defects, validate releases, and support production-ready feature delivery.

PROJECTS

FlowSight AI — Multi-Agent AI Consulting & Workflow Automation Platform May 2026

- Built an AI platform that automates the end-to-end consulting engagement — data analysis, research, business strategy, solution architecture, ROI modeling, risk & governance, and executive reporting — using a 7-agent LangGraph pipeline to deliver client-ready recommendations in minutes.
- Grounded every recommendation in the client's own policies and workflows through cited RAG retrieval (Hugging Face Qwen3 embeddings, Chroma vector store), producing traceable, audit-ready deliverables stakeholders can defend to leadership.
- Cut analysis turnaround ~60% (from 5+ minutes to ~2) and lowered per-engagement AI spend through cost-optimized model tiering — Claude Haiku for analytic steps, Claude Sonnet for executive report synthesis.
- Delivered an enterprise-ready client experience with live engagement status updates, persistent analysis history, and a polished reporting dashboard, built on FastAPI, Supabase Postgres, and Next.js/TailwindCSS.
- Safeguarded deliverable quality and consistency with LangSmith tracing and evaluation plus a 50+ test suite using dependency-injected fakes, enabling zero-cost, key-free continuous integration.

Tech Stack: Python, LangGraph, Claude API (Haiku, Sonnet), Hugging Face, Chroma, FastAPI, Supabase, PostgreSQL, Next.js, TailwindCSS, LangSmith

2026 FIFA World Cup Match Predictor — Machine Learning Pipeline Mar. 2026 – May 2026

- Engineered 25+ leakage-safe feature deltas from multiple data sources including Transfermarkt market valuations, FIFA/Elo ratings, rolling form, head-to-head records, injury burden, and manager history.
- Benchmarked Logistic Regression, KNN, Random Forest, XGBoost, HistGBM, SVM, and Stacking Ensemble models using Optuna, SMOTE, Macro F1, Log Loss, AUC-ROC, and SHAP explainability.
- Ran a 10,000-iteration Monte Carlo simulation of the full 48-team tournament bracket using calibrated probabilities to estimate advancement and title probabilities from group stage through the final.

Tech Stack: Python, Scikit-learn, XGBoost, Optuna, SHAP, Pandas, NumPy, Plotly, Matplotlib, BeautifulSoup

Split It — AI-Assisted Bill Splitting Platform Dec. 2025 – Feb. 2026

- Built OCR-powered receipt parsing workflows using OpenAI Vision API to extract structured transaction data from semi-structured receipts, enabling real-time item classification and collaborative bill allocation.
- Architected Node.js, MySQL, REST APIs, Socket.io, and a Next.js web companion to support live multi-user sessions, shareable invite links, synchronized item ownership, and cross-platform guest access.
- Integrated Stripe, Apple Pay tap-to-pay, Expo notifications, and live session updates to automate payment collection, guest follow-up, and operational coordination across connected devices.
- Tech Stack:** React Native, Expo, Node.js, MySQL, Socket.io, Next.js, OpenAI Vision API, Stripe, REST APIs

BFit — AI-Powered Fitness & Nutrition Platform Jan. 2025 – Present

- Leading development of an AI personalization platform in React Native, using OpenAI Vision API and prompt-engineered workflows for meal analysis, nutrition insights, and user guidance.
- Built a RAG-based recommendation system using Pinecone and OpenAI API to retrieve personalized workout guidance, connecting user context to AI-generated recommendations through semantic search.
- Designed social accountability features including messaging, shared progress, activity feeds, stories, and community challenges to connect AI personalization with user engagement and behavior change.
- Implemented OAuth 2.0 authentication, Firebase sessions, encrypted MySQL user data, GCP storage, and in-app subscription infrastructure to support scalable, privacy-conscious mobile deployment.
- Tech Stack:** React Native, MySQL, Node.js, OpenAI API, Pinecone, TailwindCSS, Google Cloud Storage, Firebase, Firestore

Breast Cancer Multi-Class Classification Model Apr. 2024 – May 2024

- Built an end-to-end Python classification pipeline for breast cancer subtype prediction, applying SMOTE, RFE feature selection, dimensionality reduction, and leakage-aware evaluation on clinical data.
- Benchmarked Random Forest, SVM, Gradient Boosting, and tuned TensorFlow/Keras neural networks using GridSearchCV, Keras Tuner, and StratifiedKFold, producing explainability-ready evaluation summaries for stakeholder interpretation of model tradeoffs.

LEADERSHIP & ACTIVITIES

Computer Science Student Society — American University of Beirut. Organized technical workshops and networking events, fostering cross-departmental collaboration across the CS community. 2021 – 2024