

Rahul Rakesh

437-983-4874 | rahul.rakesh1012@gmail.com | linkedin.com/in/rahulrakesh10 | rahulrakesh.dev

EDUCATION

University of Western Ontario

2023 – 2028

B.Sc. Honours Specialization in Computer Science (Co-op)

London, Ontario

– Relevant Coursework: Data Structures & Algorithms, Databases, Computer Architecture, Operating Systems

EXPERIENCE

Machine Learning Engineer

June 2026 – Present

Fly Rank

Remote

- Built a Python data pipeline processing a 28,795-row search dataset (90-day trailing aggregates), engineering features across 4 signals with leakage-safe, client-grouped transformations
- Benchmarked a Random Forest model against a baseline on a held-out 5,759-row test split; validated feature reliability using permutation importance
- Shipped an automated decision-support API output ("Action Playbook") that replaced manual review with data-driven prioritization for downstream teams

Freelance Web Developer

Sept. 2024 – Present

Self-Employed

Remote

- Built and maintained backend services and REST APIs for 5+ clients using Node.js, Express.js, and PostgreSQL, deployed to AWS with automated CI/CD pipelines
- Improved application performance with 50% faster load times through backend optimization and query tuning
- Owned projects end-to-end: requirements gathering, system design, automated deployment pipelines, and production releases on AWS

PROJECTS

Anomalock | *Python, FastAPI, PostgreSQL, Docker, scikit-learn*

[Live Demo](#)

- Built a FastAPI backend service scoring high-volume login events in real time against PostgreSQL-backed history, using identical feature logic in training and production to prevent skew
- Engineered a 5-feature data pipeline over a real-world 33M+ event dataset with leakage-safe transformations to flag compliance-relevant anomalies (impossible travel, brute-force patterns, device/location novelty)
- Trained and evaluated Isolation Forest and Random Forest models on a chronological split (137,164 rows), reaching 60.5% recall vs. 21.1% for Random Forest at a fixed 20% review budget; containerized with Docker and deployed to production

Fake Out | *Node.js, TypeScript, Socket.IO, Docker*

- Built a Node.js/TypeScript backend handling real-time multiplayer game state via Socket.IO, sustaining sub-100ms latency and 99.5% uptime under concurrent load
- Designed authoritative server-side state management with room-based event handling and heartbeat monitoring to prevent desynchronization across concurrent clients
- Containerized services with Docker and deployed via an automated CI/CD pipeline, supporting 1,000+ production sessions

VitaLink | *Django REST Framework, PostgreSQL, Gemini API | TerraHacks 2025*

- Built backend APIs (Django REST Framework) integrating Google Gemini for AI-assisted triage and a multi-factor hospital recommendation engine
- Implemented role-based access control across 3 tiers to enforce data privacy for patient records

TECHNICAL SKILLS

Languages: Python, Java, TypeScript, JavaScript, SQL, C/C++

Backend & APIs: FastAPI, Django REST Framework, Express.js, Node.js, REST APIs, Socket.IO

Cloud & DevOps: AWS (EC2, S3, Lambda), Docker, CI/CD (GitHub Actions), Git, Fly.io

Data: PostgreSQL, MongoDB, SQLite, ETL/data pipelines, feature engineering, scikit-learn, pandas

Certifications: Anthropic (Claude with the Anthropic API, Claude 101, AI Fluency) | Udemy (Intro to Cloud Computing on AWS)