DIVYANSHI KASHYAP

Fredericton, NB, Canada · linkedin.com/in/divyanshik29 · github.com/divyanshik29
AI/MLOps Engineer · BSc Computer Science (Statistics Minor) · University of New Brunswick

EDUCATION

BSc Computer Science · Statistics Minor — Sep 2023 – May 2028
University of New Brunswick, Fredericton, NB, Canada

EXPERIENCE

AI/MLOps Engineer Intern · IGT × UNB — Jan 2026 – Present
Fredericton, NB · Co-op

- Architected a full end-to-end Medallion ML pipeline from scratch: Bronze layer for raw data ingestion via Snowflake; Silver layer via MySQL/EC2 hosting the silver_dim_game table encoding 22 structured game characteristics plus boolean bonus flags for comprehensive feature engineering; Gold layer for model training and evaluation on AWS SageMaker.
- Engineered three CatBoost models targeting distinct prediction axes — game shape classification, profitability scoring, and risk assessment — with dual-mode inference architecture supporting both batch and real-time serving paths from a single model artifact.
- Designed a YAML-based configuration registry enabling fully reproducible experiment definitions; implemented GroupKFold cross-validation with 30% holdout masking to prevent data leakage across correlated game-group boundaries; tracked all experiments, metrics, and artifacts in MLflow for auditability and rollback.
- Established GitHub Actions CI pipeline with automated model validation checks, ensuring pipeline integrity and catching regressions on every push before they could propagate to training runs.
- Caught and resolved a critical Italy data granularity bug: monthly revenue figures were being ingested and compared against weekly benchmarks, causing approximately 4× revenue inflation across ~10.9 million records — identified and corrected before model training began, preventing systematic downstream corruption across all profitability and risk models. This remains one of the most high-impact data quality catches in the project's history.
- Diagnosed and engineered workarounds for a corporate Network ACL blocking EC2 SSH access mid-sprint, maintaining Silver-layer pipeline continuity without escalating infrastructure dependencies or delaying model milestones.

AI Engineer & Product Lead · NOS — Jan 2026 – Present
Part-time · Remote

- Leading all AI architecture and product-level technical direction for TravCan — an end-to-end AI travel platform built on the Claude API — co-founded with Manjunath; responsible for the agentic system design, evaluation infrastructure, and all LLM integration decisions.
- Designed and implemented a 5-phase Progressive Disclosure agent system spanning the full travel user journey: Discovery (intent capture and destination scoping), Planning (itinerary construction), Booking (offer surfacing and comparison), Confirmation (transaction validation), and Post-Booking (follow-up, changes, and support) — with 14 purpose-built tool handlers wired to each phase.
- Built custom context-compression middleware achieving ~43% token reduction through selective state disclosure, phase-aware context pruning, and intelligent carry-forward of only decision-relevant signals — directly reducing API costs at scale without sacrificing conversational continuity.

- Implemented an LLM-as-judge evaluation framework that gates every model deploy behind automated multi-dimensional quality assessments — testing factual accuracy, tone consistency, booking intent precision, and tool-call correctness — preventing regressions before they reach production.
- Leading integration research and prototyping across Duffel (real-time flight inventory and booking API), Liteapi (hotel inventory), and Stripe authorize-then-capture flow to enable live transactional capabilities within the agent pipeline.
- Designing adversarial red-teaming coverage (NightShade methodology) targeting OWASP LLM01, LLM06, LLM07, and LLM08 attack vectors as the next security hardening milestone. Stack: Rust for performance-critical middleware, Next.js for frontend, Supabase for auth and persistent state.

ML Engineer — Cohort Member · BuildersLab Apr 2026 – Present
Remote · Apprenticeship

- Selected as a founding cohort member of BuildersLab's Machine Learning Engineer track — one of the inaugural members of the first-ever cohort of the program.
- Engaging in structured applied ML training through Kaggle competitions and end-to-end project-based problem solving, building production-oriented skills across feature engineering, model selection, cross-validation strategy, and evaluation — with a focus on developing the judgment to ship reliable ML systems, not just high-scoring notebooks.

Open Source Contributor · GirlScript Summer of Code 2026 May 2026 – Present
On-site · Fredericton, NB

- Selected as a GSSoC 2026 open source contributor; actively contributing to open-source repositories in Go and Rust — languages directly aligned with systems-level and high-performance backend infrastructure work.
- Engaged in the full open-source contribution lifecycle: issue triage, feature implementation, code review participation, and pull request collaboration within multi-contributor repositories spanning the summer cohort.

Teaching Assistant — Calculus · University of New Brunswick 2023 – Present
Fredericton, NB · On-campus

- Support undergraduate students across Calculus coursework through weekly office hours, structured assignment grading, and targeted exam preparation sessions — serving students across multiple cohorts simultaneously.
- Translate abstract mathematical concepts into accessible, intuitive explanations — a communication skill that directly informs how I present ML model behavior, statistical results, and technical trade-offs to non-specialist stakeholders in engineering contexts.

LEADERSHIP & COMMUNITY

President · WiCSE — Women in CS & Engineering, UNB 2024 – Present
Fredericton, NB · On-campus

- Lead the WiCSE student club at UNB — driving community programming, technical workshops, hackathon participation, and career development initiatives for women in computing and engineering.
- Co-building two campus-facing tech products under WiCSE: Freddy Trails (AI storytelling walking trail for Fredericton, Claude-generated landmark narratives, QR check-ins, business analytics dashboard, and a sponsor/vendor model — full PRD and pitch deck complete) and Freddy Finds (hyperlocal vendor discovery platform with committed vendors, launching post-Garrison Market).

– Running a structured mentorship program with two mentees: Faizan (ML engineering roadmap — weekly syncs, Kaggle projects, model deployment track) and Sammy (agentic AI engineering track — a former software developer in India transitioning into LLM-native product engineering).

Founder · The Circuit Fall 2025
Fredericton, NB · UNB

– Founded The Circuit, a student-led tech community at UNB focused on bridging the gap between classroom CS and real-world software engineering — bringing together students interested in building, shipping, and talking about technology.
– Launched Fall 2025; paused Winter 2026 due to concurrent full-time professional commitments across NOS and IGT co-op. Intended to resume post-co-op cycle.

PROJECTS

CalgaryPulse — AI Civic Intelligence Platform · Solo-Built

– Solo-engineered a full-stack AI civic tech platform tackling Calgary's 30.4% downtown commercial vacancy crisis and its $16B economic impact. Selected as a MindFuel Tech Futures 2026 finalist from a field of 31 projects across 7 Canadian provinces; seed funding approved; Prototype Showcase scheduled May 30, 2026.
– Designed a two-model AI architecture: Gemma 4 2B running locally for zero-cost, high-throughput civic intent classification (route all cheap inferences to a local model to eliminate API spend at volume) paired with Claude API for professional-tier demand intelligence, narrative generation, and complex civic reasoning that requires frontier-level output quality.
– Implemented CrewAI-style multi-agent orchestration with a dedicated Hermes Agent running nightly autonomous web scraping pipelines and signal enrichment against Calgary Open Data APIs — continuously updating the platform's understanding of property availability, zoning changes, and market signals without manual intervention.
– Developed the PulseScore compatibility algorithm weighting business-property fit across four axes: location desirability, space size match, budget alignment, and business use type — giving prospective tenants a quantified, explainable compatibility score for every available downtown property.
– Stack: React + Three.js (3D city visualization), FastAPI, PostgreSQL+PostGIS (spatial queries), Mesa ABM (agent-based economic simulation), NumPy, GeoPandas, Deck.gl, Docker, Railway. Scale: 500K+ property records, 750K simulation agents, <30s full-city simulation runtime, <100ms spatial query latency.

Monitor Lizard — Autonomous Co-op Job Tracking Agent · Production-Grade

– Built a fully autonomous, production-grade co-op job tracking agent on the OpenClaw framework. Nightly Noctis Mode pipeline scans 60+ job portals without manual intervention — running on a cron schedule during off-hours to surface new opportunities each morning.
– Implements A–F role scoring with a critic-review pass: an initial scoring agent evaluates each posting against a structured rubric, then a second critic agent challenges the score before it's finalized — reducing false positives and preventing good roles from slipping through underscored.
– ChromaDB vector memory enables semantic deduplication and cross-session recall; smart model routing optimizes cost vs. quality at each pipeline stage: Claude Opus for initial setup and configuration, Kimi 2.5 for daily scan runs, Claude Haiku for lightweight cron maintenance tasks. 172 tests. Submitted to OpenClaw Challenge.

Freddy Trails — AI Storytelling Walking Trail · WiCSE Initiative

- Designing an AI-powered walking trail experience for Fredericton that transforms city landmarks into interactive storytelling stops. Claude generates contextually rich, historically grounded narratives for each landmark stop, delivered to visitors via QR code scan on their phones.
- Platform includes a business analytics dashboard tracking visitor engagement per stop, session length, and trail completion — enabling a sponsor and local vendor monetization model. Full PRD and pitch deck complete; co-built with WiCSE VP.

Freddy Finds — Hyperlocal Vendor Discovery Platform · WiCSE Initiative

- Hyperlocal vendor and market discovery platform built for Fredericton — helping residents find local small businesses, artisan vendors, and pop-up markets that don't have major digital presence.
- Vendors have already committed ahead of launch; platform releasing publicly after the Garrison Market season opens. Co-built with WiCSE VP as a community-first, non-commercial platform.

Kaashvi — ReAct Desktop Agent · Personal AI Chief of Staff

- Built a desktop-native ReAct (Reasoning + Acting) agent that operates as a personal AI chief of staff — autonomously planning and executing multi-step tasks across the user's calendar, notes, and productivity tools without manual orchestration between apps.
- Integrated Google Calendar OAuth2 for full read/write calendar access and Notion API for knowledge base and task management — enabling Kaashvi to schedule meetings, create follow-up tasks, surface relevant notes, and cross-reference context across both systems in a single reasoning loop.
- ReAct loop architecture: agent reasons about the goal, selects a tool action, observes the result, and iterates — handling ambiguity, multi-step dependencies, and tool failures gracefully without requiring user re-prompting mid-task.
- Stack: Electron + React + Vite for a native cross-platform desktop experience with a clean, minimal UI that stays out of the way while the agent works.

NightShade — Adversarial LLM Red-Teaming Framework

- Designed a structured adversarial red-teaming framework targeting LLM-powered applications across the OWASP LLM Top 10 threat taxonomy — with specific focus on LLM01 (prompt injection), LLM06 (sensitive information disclosure), LLM07 (insecure plugin design), and LLM08 (excessive agency).
- Methodology directly applied to TravCan security hardening at NOS. PRD complete; open-source release scoped as the project matures.

CarbonLedger — C++17 Process Carbon Measurement Library

- Cross-platform C++17 library and CLI tool that instruments any process — measuring real-time CPU and memory consumption and converting those telemetry signals to CO2 equivalent estimates using the Green Software Foundation SCI formula.
- Architected around an ISampler interface with a MockSampler implementation for test-driven development; 85% coverage CI gate enforced; GitHub Actions matrix build across GCC, Clang, and MSVC compilers; AddressSanitizer and Valgrind for memory safety validation. Primary use case: EU CSRD software sustainability compliance reporting.

NyxLink — AI-Augmented URL Shortening Service · Production-Grade

- Built a production-ready URL shortening service that goes beyond basic shortening by integrating a phishing detection pipeline, real-time bot classification, full click analytics, and authenticated user account management — engineered to demonstrate MLOps principles, async system design, and cloud-native deployment practices at production scale.

- AI safety layer queries Google Safe Browsing API v4 on every shorten request — rejecting phishing, malware, and unwanted software destinations with HTTP 422 and storing a 0.0–1.0 safety_score per link record. Graceful degradation on API unavailability: logs failure, stores null score, and allows shortening without blocking the user.
- Real-time bot detection pipeline runs fully asynchronously via FastAPI BackgroundTasks — scoring every redirect using User-Agent string classification, headless browser signals, and inter-request timing anomalies. is_bot flag stored per click row; analytics dashboard filters human-only traffic to surface true engagement metrics without count inflation.
- Sub-millisecond redirect latency via Redis 7 caching (p50 <2ms cache hit, p99 <8ms). Async click logging ensures redirect latency is never blocked by analytics writes. Tiered rate limiting via SlowAPI + Redis: 10 req/min anonymous, 60 req/min authenticated, 300 req/min on the redirect endpoint — all returning HTTP 429 + Retry-After on breach.
- Full feature set: 7-character Base62 short codes, custom aliases with reserved-word blocklist, optional TTL with 410 Gone expiry, password-protected redirects, QR code generation (PNG, configurable size), paginated link management, and a rich analytics API covering clicks by day, referrers, countries, device breakdown, and bot vs. human split.
- Stack: FastAPI async (Python 3.11+), SQLAlchemy async + PostgreSQL 15, Alembic migrations, Redis 7, python-jose (JWT/HS256), Passlib (BCrypt), httpx, qrcode + Pillow, pytest + httpx AsyncClient, Locust load testing at 1,000 concurrent virtual users, Docker + Docker Compose, GitHub Actions CI, Railway deployment.

ACHIEVEMENTS & RECOGNITION

- **MindFuel Tech Futures 2026 — Finalist** 31 projects across 7 Canadian provinces; seed funding approved for CalgaryPulse. Prototype Showcase: May 30, 2026.
- **National Data Challenge — Top 18 Finalist** Competed nationally; placed in the top 18 submissions out of all participating teams.
- **GSSoC 2026 — Selected Open Source Contributor** GirlScript Summer of Code; contributing in Go and Rust.
- **BuildersLab — Founding Cohort Member** ML Engineer track; selected for the inaugural cohort of the program.
- **OpenClaw Challenge — Submission** Monitor Lizard autonomous job-tracking agent submitted to the OpenClaw open-source challenge.

TECHNICAL SKILLS

Languages: Python, Rust, Go, JavaScript/TypeScript, C++17, SQL

AI / LLM: Claude API, LLM-as-judge evaluation, CrewAI multi-agent orchestration, ChromaDB vector memory, prompt engineering, adversarial red-teaming (OWASP LLM Top 10)

ML & Data Science: CatBoost, AWS SageMaker, MLflow, GroupKFold CV, SHAP, Evidently AI, Gemma 4 (local inference), Kaggle competition workflows

Data Engineering: Snowflake, PostgreSQL/PostGIS, MySQL, EC2, Mesa ABM, GeoPandas, NumPy, Deck.gl, Medallion pipeline architecture

Dev & DevOps: Docker, Railway, GitHub Actions CI/CD, FastAPI, Next.js, React, Three.js, Supabase, React Native, OpenTelemetry

Systems: C++17 cross-platform library design, AddressSanitizer, Valgrind, GCC/Clang/MSVC matrix builds, ISampler TDD pattern

Frameworks & Concepts: Progressive Disclosure agents, Medallion architecture, SCI formula (Green Software Foundation), EU CSRD compliance, GroupKFold CV, dual-mode inference