

# Geoffrey Meric

+33 6 51 26 20 96 | [gcs.meric@gmail.com](mailto:gcs.meric@gmail.com) | [linkedin.com/in/geoffrey-meric](https://linkedin.com/in/geoffrey-meric) | [gcsmeric.github.io](https://gcsmeric.github.io)

## EDUCATION

---

### McGill University

Sep 2020 - Dec 2024

**Bachelor of Science in Computer Science, GPA 4.00/4.00**

*Montreal, Canada*

Researched occupational gender bias in BART-NLI models, received Judie Rimer merit scholarship

## TECHNICAL SKILLS

---

**Languages:** Python, TypeScript, SQL, Bash **AWS:** ECS, SageMaker, Lambda, ElastiCache, CloudWatch, VPCs

**Frameworks & Technologies:** LangChain (Deep Agents), LangGraph, Weaviate, LaunchDarkly, Splunk, Dynatrace, Streamlit, asyncio, nginx, Gunicorn, Flask, numpy/pandas, Git, Docker, LLMs, Cursor, MCP, agent skill management

## EXPERIENCE

---

### Applied AI Engineer

Mar 2025 - Present

#### Autodesk

*Montreal, Canada*

- Distributed Backend Engineering (AWS), Science and LLMOps for Autodesk's Commerce & Support Agent
- Designing full-lifecycle feature improvement experiments, progressive rollout & success metric measurement strategies, configuring A/B tests using LaunchDarkly platform, performing causal inference analysis on outcomes
- Developing planning evaluation framework & named-entity recognition tool for ReAct Langchain Deep Agents
- Primary MLOps responder in US timezones: monitoring operations, handling error spikes/VDB outages/alerts, shipping hotfixes, and improving observability via CloudWatch, Dynatrace, and the Opik LLMops/eval platform
- Latency optimization: deploying ECS + DynamoDB context engineering infrastructure to parallelize context computation (e.g. conversation summarization), reducing answer latency by ~1s/10% in 15k daily conversations
- Backfilling LLM-as-a-judge conversation evaluation data for ~1M past conversations using Airflow & OpenAI Batch API, creating dashboard to semantically cluster user requests & source documents and inspect query performance (latency, errors, escalations, LLM-j metrics) filtered by product & request type and identify content gaps
- Creating personalization PoCs using inferred user intent from activity & recommender outputs to steer answering
- Developing a synthetic evaluation pipeline using LLMs to generate test queries and ground-truth answers from support docs, distill RAG outputs into fact-only representations with stylistic noise removed, and classify responses as semantically equivalent, incomplete, hallucinated, or contradictory while measuring document retrieval recall
- Addressing repeated client reinitialization causing high latency & resource leaks via singletons (10x Weaviate error reduction), rewrote high-latency event instrumentation services to asynchronously buffer events to batch send, resolved faulty data bugs causing A/B test SRM + UX inconsistencies and errors for all LATAM & Norway users

### Machine Learning Engineer

Sep 2024 - Dec 2024

#### Serifos Technologies

*Montreal, Canada*

- Developed online economic sentiment analysis tools for institutional clients using GPT and BART-NLI models
- Scraped and analyzed online forum discussions (X, reddit, etc.), using GPT and BART-NLI language models to identify salient topics and performed topic modeling to identify trending themes outside of predefined indicators
- Built Streamlit UI to visualize user sentiment over time across economic indicators, search comments via vectorDB

### Machine Learning Operations Engineering Intern

May 2024 - Aug 2024

#### Autodesk

*Montreal, Canada*

- Created a python-based programmatic evaluation system for Retrieval-Augmented Generation (RAG) models to diagnose model failures and identify trends in queries that yield poor quality responses, using LLMs & NLI models
- Improved conversation summarization by designing evaluation methods to compare summarization prompts

### Software Development Intern

May 2023 - Aug 2023

#### Autodesk

*Montreal, Canada*

- Built live job logging & monitoring tools for a cloud rendering Autodesk Maya plug-in & tripled its throughput

### Software Engineering Intern

May 2022 - Aug 2022

#### Procter & Gamble

*Geneva, Switzerland*

- Built customer data platform architecture visualization and CRM campaign operation automation dashboards

## PROJECTS & ACTIVITIES

---

**Competitive chess** | 2500+ elo bullet rating on lichess, youth regional champion & national top 15 chess player in France

**Substack Writing** | Writing about topics I find interesting on substack. Main focuses so far are financial systems & AI