

Mathieu Duchesneau – AI & Software Engineer

Québec, Canada – (514) 246-6064

 duchesneau.mathieu@gmail.com

 mathieuduchesneau.ca

 [LinkedIn](#)

 [GitHub](#)

SUMMARY

AI and software engineer with 10+ years of experience building machine learning systems and the infrastructure that powers them.

Ph.D. graduate from Mila (Université de Montréal) with original research introducing the first neural method to achieve **large-scale O(1) exact retrieval** through binary encoders designed for direct hash-table indexing.

Co-founded a startup that deployed real-time electricity arbitrage using ML-driven decision models running on custom Linux infrastructure.

For a more interactive overview, I self-host a fine-tuned large language model on mathieuduchesneau.ca that answers questions about my work and background. The entire platform, from Kubernetes infrastructure to the front-end, is deployed and maintained by me.

WORK EXPERIENCE & EDUCATION

- | | | |
|-------------------|---|-------------------|
| 09/2017 – 01/2025 | Ph.D., Artificial Intelligence
<i>Mila, DIRO, Université de Montréal</i> | Montréal, Canada |
| | <ul style="list-style-type: none">• Thesis: Learning Equivalence Hash Functions<ul style="list-style-type: none">◦ Developed the first training objective enabling neural networks to encode elements drawn from billions of classes into binary representations that are unique for each class.◦ Built a prototype that performs full lookups in under 100 ms in a database containing a billion images, using a single GPU for encoding and a standard CPU for hash-table retrieval.• Extensive experience building and training models for computer vision and natural language processing.• Managed large-scale experiment pipelines on the Mila and Digital Alliance clusters via Slurm, running hundreds of training jobs. | |
| 08/2017 – 06/2019 | Linux System Administrator & AI Engineer
<i>Énergie Manifold Inc. (startup co-founder)</i> | Mascouche, Canada |
| | <ul style="list-style-type: none">• Built a fully automated electricity arbitrage system operating between NYISO (US-NY) and IESO (CA-ON). | |

- Designed and implemented **ML classifiers** to decide between “IESO → NYISO,” “NYISO → IESO,” and “no trade” using real-time grid loads, prices, and **40+ engineered features**.
- Deployed and maintained **on-premises Linux infrastructure** for real-time data ingestion, model inference, and full API integration with both grid operators.

09/2016 – 04/2022

Teaching Assistant

Montréal, Canada

DIRO, Université de Montréal

- Conducted Introduction to Theoretical Computer Science course over four semesters, grading assignments and delivering engaging presentations on finite automata, context-free languages, Turing machines, and complexity classes (such as P and NP), as well as language decidability.
- Conducted the Introduction to Data Science course, evaluating assignments and providing insightful presentations on classification, regression, clustering, and statistical concepts like confidence intervals and hypothesis testing.

09/2012 – 04/2015

B.Sc., Mathematics & Computer Science

Montréal, Canada

DIRO, Université de Montréal

09/2010 – 05/2011

D.E.C., Natural Science

Terrebonne, Canada

Cégep de Lanaudière à Terrebonne

SKILLS

AI ENGINEER

- Python
- PyTorch
- MLOps
- MLflow
- Hugging Face
- Llama.cpp
- Slurm
- Image processing
- Convolutional neural networks
- Natural language processing
- Recurrent neural networks

SOFTWARE ENGINEER

- Bash
- FastAPI
- Jinja2
- CSS
- JavaScript
- SQL
- CI/CD
- GitOps
- DevOps
- Docker
- Kubernetes
- GitHub actions

- Large language models
- Transformer
- Fine-tuning
- LoRA / QLoRa
- FluxCD
- Nginx
- Traefik
- Prometheus
- Grafana
- Calculus
- Linear algebra
- Probability & statistics
- Data structures
- Algorithmic complexity

PROJECTS

10/2025 – Present

Self-Hosted Website with LLM Assistant

Montréal, Canada

- Deployed a **full-stack Kubernetes platform** using FastAPI, Traefik, Prometheus, and Grafana for serving and monitoring.
- Built a complete **CI/CD pipeline** with GitHub Actions and FluxCD.
- Fine-tuned and hosted a **personalized LLM** using QLoRA, designed to answer questions about my work and background.
- Integrated MLflow for **continuous experiment tracking**, allowing new state-of-the-art models to be promoted automatically into production.
- The model is served live on mathieuduchesneau.ca, demonstrating full lifecycle AI deployment from training to production.

LANGUAGES

French (fluent)

English (fluent)

LINKS

- Thesis: <https://hdl.handle.net/1866/40944>
- Website: <https://mathieuduchesneau.ca>
- LinkedIn: <https://www.linkedin.com/in/mathieu-duchesneau-047a0b307>
- GitHub: <https://github.com/duchesneumathieu>