

Xavier Bouthillier

MACHINE LEARNING RESEARCHER SCIENTIST

✉ xavier.bouthillier@umontreal.ca | 🏠 bouthilx.github.io | 📱 bouthilx | 📄 bouthilx

Skills Matrix

		ML Domains		ML & Soft. Eng.				ML Libs			Programming			Databases		Tools				Management & Development				
		NLP	Computer Vision	HPO/BO	Data Management	Parallelization	Reproducibility	CI/CD	System Monitoring	PyTorch	CUDA	Dask/Ray	Python	C/C++	JavaScript	MongoDB	PostgreSQL	Git/mercurial	github-actions/Travis	Docker/Singularity	Slurm/Moab/Condor	Team management	Hiring	Partnerships
Working experience	Res. Scientist	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
	Res. Developer			☐		☐	☐	☐			☐			☐				☐	☐			☐	☐	☐
	Res. Intern	☐				☐		☐			☐										☐			
	Res. Intern	☐									☐			☐							☐			
	Res. Assistant													☐		☐								
Studies	PhD	☐	☐	☐		☐	☐	☐			☐	☐			☐		☐	☐	☐	☐		☐		
	Master	☐	☐	☐											☐	☐								
	Bacc.															☐								
Open Source Projects	Mahler					☐	☐	☐	☐	☐					☐			☐	☐	☐				
	Kleió			☐		☐	☐								☐			☐						
	Theano							☐				☐	☐					☐						

Languages French (Native), English (Fluent), German (B1), Mandarin (Beginner)

☐ Clickable anchor links.

Working Experience

Mila

Montréal, Canada

Research Scientist, Product owner for multiple projects in the Innovation, Development and Technology team at Mila.

May 2022 - Present

Selected Projects and Corresponding Tasks

- Milabench: GPU Benchmark including DL experiments representative of research at Mila. (with Pierre Delaunay as Tech-Lead)
 - Supervise a new automated literature analysis using GPT-4o to gather statistics from Mila's publications.
 - Collaborate with professors and students to gather research pipelines for the benchmark.
 - Establish collaborations with Mila partners to ensure good support on a variety of GPUs.
- SARC: Monitoring and analysis of HPC clusters usage. (with Bruno Carrez as Tech-Lead)
 - Co-design an aggregation system for compute metrics across an heterogeneous group of clusters, leveraging SlurmDB and Prometheus.
 - Establish CI/CD best practices and requirements for an alert system.
 - Analyse system metrics aggregated with SARC to identify bad usages, then meet and help worst offenders to better use GPUs.
- External HPC resources: Responsible for the annual renewal of Mila's biggest compute allocation from DRAC.
 - Plan and organize the annual renewal submission, involving professors, professors' assistants and developpers.
 - Analysis of compute usage patterns and trends, and predictions of upcoming needs based on SARC data.
- Datasets and models: Dataset and model weights registry across HPC clusters. (with Satya Ortiz-Gagné as Tech-Lead)
 - Supervise benchmarking of data storage solutions for an heterogeneous group of clusters.

General Tasks

- Strategic planning, identification of major needs, and coordination of resources.
- Project management (6 to 14 projects in parallel, involving up to 9 developpers).
- Leading negotiations with several service providers in the ML field, in collaboration with legal and partnership teams.
- Involved in the hiring process.
 - Producing job description, filtering candidates and interviewing.
 - Developped an interview test (the Pr. Tournesol interview script) to verify familiarity with DL tools and optimization reflexes.
- Representing Mila in international visits, events and conferences.

Mila, Université de Montréal

Research Developer

Montréal, Canada

Oct. 2017 - May 2022

- Lead developer of Oríon, leading project directions and main goals.
- Co-design architecture of the framework.
- Reimplement hyperparameter optimization algorithms.
- Recruit and interview potential candidates for Mila IDT team.
- Mentor interns for software engineering and students for research projects.
- Organize and manage code development sprints involving up to 10 developers.
- Provide support to users.
- Setup and maintain CI with Travis, github-actions and codecov.
- Designed and implemented an experiment version control system.

Nuance Communications

Research Intern

Montréal, Canada

May. 2016 - Dec. 2016

- Developed a data augmentation algorithm to generate fake examples based on large medical documents.
- Reproduced the paper Hierarchical Attention Networks for Document Classification and improved the attention mechanism.
- Implemented a parallelized pipeline to efficiently convert large medical documents in deep hierarchical structures.

Nuance Communications

Research Intern

Montréal, Canada

Sep. 2014 - Jan. 2015

- Demonstrated higher performances with deep learning models on domain classification for queries of personal assistant queries.
- Developed a new convolutional model inspired by n-grams using PyLearn2.
- Implemented CUDA kernel for faster 1-d convolution, integrated with Theano.

LISA (Mila), Université de Montréal

Research Assistant

Montreal, Canada

2010 - 2012

- Learned website development from scratch, implemented jQuery plugins, and maintained website.

Education

Mila, Université de Montréal

PhD in Computer Science, Machine Learning

2014 - 2022

Master in Computer Science, Machine Learning

2013 - 2014

Bachelor in Computer Science

2009 - 2012

Freiburg Albert-Ludwigs Universität (Germany)

Exchange program during Bachelor's degree

2011-2012

CÉGEP Saint-Laurent

Natural Science

2008 - 2009

Music, Composition Profile

2005 - 2008

Open Source Projects

Mahler (bouthilx.github.io/projects/2-mahler)

Developer - Prototype stage

2019-2020

Mahler is a framework to provide more control over workflow, better resiliency and better automation in HPC

- Implemented an automated remote installations using Fabrik for multi-cluster setups.
- Implemented a singularity-based workflow to easily deploy experiments on different clusters.
- Implemented a Dispatcher that monitors GPU usage and oversubscribe them with additional workers if possible.
- Implemented a dashboard using Dash to provide visualization and control over the pool of workers and registered tasks.

Kleiô (bouthilx.github.io/projects/4-kleio)

Developer - Prototype stage

2018-2019

Kleiô is an experiment manager that provides full traceability.

- Implemented a new data architecture based on the concept of events sourcing.
- Implemented remote commands (cat, tail, head, ...) for logs of experiments.

Theano (github.com/Theano/Theano)

Supervisor

2015-2017

Theano is a Python library that allows you to define, optimize, and evaluate mathematical expressions efficiently

- Mentored students contributing for the Common-Code-Workflow

Publications

JOURNAL ARTICLES

Emonets: Multimodal Deep Learning Approaches for Emotion Recognition in Video

Samira Ebrahimi Kahou, Xavier Bouthillier, Pascal Lamblin, Caglar Gulcehre, Vincent Michalski, Kishore Konda, Sébastien Jean, Pierre Froumenty, Yann Dauphin, Nicolas Boulanger-Lewandowski

Journal on Multimodal User Interfaces 10.2 (2016) pp. 99–111. Springer, 2016

CONFERENCE PROCEEDINGS

Accounting for variance in machine learning benchmarks

Xavier Bouthillier, Pierre Delaunay, Mirko Bronzi, Assya Trofimov, Brennan Nichyporuk, Justin Szeto, Nazanin Mohammadi Sepahvand, Edward Raff, Kanika Madan, Vikram Voleti, Samira Ebrahimi Kahou, Vincent Michalski, Tal Arbel, Chris Pal, Gael Varoquaux, Pascal Vincent

Proceedings of Machine Learning and Systems 3 (2021). 2021

Unreproducible Research is Reproducible

Xavier Bouthillier, César Laurent, Pascal Vincent

International Conference on Machine Learning, 2019

Fast Approximate Natural Gradient Descent in a Kronecker Factored Eigenbasis

Thomas George, César Laurent, Xavier Bouthillier, Nicolas Ballas, Pascal Vincent

Advances in Neural Information Processing Systems, 2018

Efficient Exact Gradient Update for Training Deep Networks with Very Large Sparse Targets

Pascal Vincent, Alexandre De Brébisson, Xavier Bouthillier

Advances in Neural Information Processing Systems, 2015

Combining Modality Specific Deep Neural Networks for Emotion Recognition in Video

Samira Ebrahimi Kahou, Christopher Pal, Xavier Bouthillier, Pierre Froumenty, Çağlar Gülçehre, Roland Memisevic, Pascal Vincent, Aaron Courville, Yoshua Bengio, Raul Chandias Ferrari

Proceedings of the 15th ACM on International conference on multimodal interaction, 2013

WORKSHOPS

Improving Reproducibility of Benchmarks

Xavier Bouthillier

CI/ML Workshop at Advances in Neural Information Processing Systems, 2019

An Evaluation of Fisher Approximations Beyond Kronecker Factorization

César Laurent, Thomas George, Xavier Bouthillier, Nicolas Ballas, Pascal Vincent

Workshop at International Conference on Learning Representations, 2018

Orion: Experiment Version Control for Efficient Hyperparameter Optimization

Christos Tsirigotis, Xavier Bouthillier, François Corneau-Tremblay, Peter Henderson, Reyhane Askari, Samuel Lavoie-Marchildon, Tristan Deleu, Dendi Suhubdy, Michael Noukhovitch, Frédéric Bastien

AutoML Workshop at the International Conference on Machine Learning, 2018

REPORTS

Introducing Milabench: Benchmarking Accelerators for AI

Pierre Delaunay, Xavier Bouthillier, Olivier Breuleux, Satya Ortiz-Gagné, Olexa Bilaniuk, Fabrice Normandin, Arnaud Bergeron, Bruno Carrez, Guillaume Alain, Soline Blanc, ...

arXiv preprint arXiv:2411.11940 (2024). 2024

Survey of machine-learning experimental methods at NeurIPS2019 and ICLR2020

Xavier Bouthillier, Gaël Varoquaux

Research Report hal-02447823, 2020

Theano: A Python Framework for Fast Computation of Mathematical Expressions

The Theano Development Team, Rami Al-Rfou, Guillaume Alain, Amjad Almahairi, Christof Angermueller, Dzmity Bahdanau, Nicolas Ballas, Frédéric Bastien, Justin Bayer, Anatoly Belikov, ...

arXiv preprint arXiv:1605.02688 (2016). 2016

Exact Gradient Updates in Time Independent of Output Size for the Spherical Loss Family

Pascal Vincent, Alexandre Brébisson, Xavier Bouthillier

arXiv preprint arXiv:1606.08061 (2016). 2016

Dropout as Data Augmentation

Xavier Bouthillier, Kishore Konda, Pascal Vincent, Roland Memisevic

arXiv preprint arXiv:1506.08700 (2015). 2015

Teaching Experience

IFT6390 Foundations of machine learning, *Teaching Assistant, Université de Montréal*

Fall 2014

IFT6390 Foundations of machine learning, *Teaching Assistant, Université de Montréal*

Fall 2013

Talks

20 Nov 2024	Milabench: Benchmarking suite for GPU-intensive AI pipelines	tinyurl.com/4mcjwpx2	SC, Atlanta, USA
6 Nov 2024	Accounting for Variance and HPO in ML Benchmarks	tinyurl.com/3davxrus	UdeM, Montreal, Canada
9 Oct 2024	Variance et HPO dans les bancs de test en ML	tinyurl.com/3davxrus	UdeM, Montreal, Canada
3 May 2024	HPC and software engineering for HPC @ Mila		LRZ, Munich, Germany
29 Apr 2024	HPC and software engineering for HPC @ Mila		JSC, Jülich, Germany
26 Apr 2024	HPC and software engineering for HPC @ Mila		CINES, Montpellier, France
1 Feb 2024	Software engineering for HPC @ Mila		Vector, Toronto, Canada
29 Nov 2023	Variance et HPO dans les bancs de test en ML	tinyurl.com/3davxrus	UdeM, Montreal, Canada
29 Nov 2023	Accounting for Variance and HPO in ML Benchmarks	tinyurl.com/3davxrus	UdeM, Montreal, Canada
27 Sep 2023	Mila, IDT, Oríon and the community of HPO users		AutoML, Freiburg, Germany
1 Dec 2022	Accounting for Variance and HPO in ML Benchmarks	tinyurl.com/3davxrus	UdeM, Montreal, Canada
30 Nov 2022	Variance et HPO dans les bancs de test en ML	tinyurl.com/3davxrus	UdeM, Montreal, Canada
22 Sep 2022	Accounting for Variance in Machine Learning Benchmarks	tinyurl.com/vrzk3and	MLCommons, remote
22 Jul 2022	Did We Forget about the Canonical Source of Variance in ML?	tinyurl.com/mtnbdypa	DataPerf, ICML2022, remote
9 May 2022	Le petit scientifique numérique	tinyurl.com/ytjc6wj4	CScience, Félix-Leclerc, Remote
5 May 2022	Thesis Defense: Accounting for Var. and HPO in ML Benchmarks	tinyurl.com/47644yha	Udem, remote
13 Jan 2022	How do we get more people to use HPO? (Round table)		AutoML Seminar, remote
18 Nov 2021	Reliability of benchmarks and why HPO is important	youtu.be/ZRQF72IXiDc	AutoML Seminar, remote
14 Jul 2021	Practical approaches for efficient HPO with Oríon	youtu.be/H1jQBQlbQmA	SciPy 2021, remote
9 Jun 2021	Oríon - Librairie pour l'optimisation d'hyperparamètres	tinyurl.com/ycku5vhv	Mila, remote
19 May 2021	Black-Box Optimization using Dask with Oríon	youtu.be/W5oWdRiSSr8	Dask Summit 2021, remote
7 Apr 2021	Accounting for Variance in Machine Learning Benchmarks	youtu.be/-bukIBTkQiM	MLSys 2021, remote
16 Mar 2021	Practical Approaches for Efficient Hyperparameter Optimization	youtu.be/qq69vxF3LTI	AICamp, remote
12 Jan 2021	Oríon - A framework for distributed hyperparameter optimisation	tinyurl.com/yuy5w6ta	Mila & Vector, remote
12 Jan 2021	On a quest for more trustworthy scientific conclusions in ML	tinyurl.com/2hv27p5m	Parietal, INRIA, remote
5 Nov 2020	L'intelligence artificielle dans les bibliothèques	Table ronde	CBPQ, remote
17 Feb 2020	Reproducibility & Hyperparameter Optimization with Oríon	tinyurl.com/43zt5wej	Ai4Sim, Atos, remote
4 Dec 2019	Unreproducible Research is Reproducible	tinyurl.com/t7rpaaa	CHAI, Berkeley, USA
28 Nov 2019	Oríon: A Framework for Distributed Hyperparameter Optimisation	tinyurl.com/t4ubtb4	Mila, Montreal, Canada
21 Nov 2019	Reproducibility in ML, or why benchmarks are lotteries	tinyurl.com/u54a7o6	Mila, Montreal, Canada
21 Nov 2019	Oríon: A Framework for Distributed Hyperparameter Optimisation	tinyurl.com/scxq63b	Mila, Montreal, Canada
15 Nov 2019	Unreproducible Research is Reproducible	tinyurl.com/rkgp55k	Mila, Montreal, Canada
31 Jul 2019	Reproducibility in AI	tinyurl.com/roue8js	Stradigi AI, Montreal, Canada
13 Jun 2019	Unreproducible Research is Reproducible	tinyurl.com/w52vjjsy	ICML, Long Beach, USA
10 Apr 2019	Introduction à l'intelligence artificielle	For a specialized school for gifted children	UdeM, Montreal, Canada
14 Nov 2018	Intelligence artificielle: Une fabrique à outils	tinyurl.com/yx6thyjd	CPI, Montreal, Canada
12 Nov 2018	Introduction à l'intelligence artificielle	For high school <i>Forum des jeunes en science</i>	CEGEP Montmo., Laval, Canada
30 Aug 2018	Tutorial on Oríon	tinyurl.com/u4pjky	Mila, Montreal, Canada
14 Jul 2018	Oríon: Experiment Version Control for Efficient Hyperparameter Optimization		ICML, Stockholm, Sweden
26 Apr 2014	Introduction au langage de programmation Python		Google, Montreal, Canada

Extracurricular Activity

Reviewer

NeurIPS, ICML, ICLR, AISTATS, MLSys, Neural Computation, AutoML Workshop,
NAS Workshop, Reproducibility Challenges, MLSys'20 Artifact Evaluation,
Rethinking ML Papers workshop @ ICLR 2021

2015 - Present

Workshop Organizer

Retrospectives Workshop
Setting up ML Evaluation Standards to Accelerate Progress

Dec. 2019

Apr. 2022

Streaming Designer and Organizer

Mila, Université de Montréal

Mar. 2018 - Oct. 2018

- Designed from scratch the setup, procedure and guidelines for recording and streaming at Mila.
- Trained dozen of volunteers to record and stream about a dozen different reading groups at Mila.
- Received a price of 6000\$ to reward the quality and importance of the initiative and the work done.