Mila strategic research coalitions
Overview

Oct 2026

Mila

Mila

Mila: A unique AI research institute

Founded by Professor Yoshua Bengio, Mila is a global hub of scientific advances inspiring innovation and the growth of AI for the benefit of all.

Mila is the world's largest academic research lab bringing together a deep learning research community.

01
Talent and community

02
Cutting-edge AI research

03
Adoption and Innovation

04
Governance and global reach

Université de Montréal · McGill · Québec · Canada



Mila

Our scientific director

Dr. Hugo Larochelle

- Trained by AI pioneers Yoshua Bengio & Geoffrey Hinton
- Former head of Google Brain Montreal; co-founder of Whetlab (acquired by Twitter)
- Associate Professor at Université de Montréal

Mila

Research Community

Mila brings together the world's top researchers in machine learning

150+

World-renowned professors
in ML, DL and RL

1500+

Researchers and students
MSc, PhD, Post-doc

160+

Experts and staff
Applied AI experts, IT and administration talent

140+

Industry partners

Microsoft Google SAMSUNG Meta

Mila

World-renowned Professors



Yoshua Bengio

UdeM

AI Safety, Medical Machine Learning, Representation Learning, Reinforcement Learning, Causality, Generative Models, Molecular Modeling



Aaron Courville

UdeM

Representation Learning, Reinforcement Learning, Efficient Communication in General-Sum Games, Generative Models, Multi-Agent Systems, Game Theory, Natural Language Processing



Pierre-Luc Bacon

UdeM

Reinforcement Learning, Representation Learning, Neural Differential Equations, Transformer-Based Models



Laurent Charlin

HEC/UdeM

Reinforcement Learning, AI for Science, Generative Models, Probabilistic Models, Graph Neural Networks, Recommender Systems, Natural Language Processing



Doina Precup

McGill

Medical Machine Learning, Reinforcement Learning, Probabilistic Models, Molecular Modeling, Reasoning



Joëlle Pineau

McGill

Medical Machine Learning, Reinforcement Learning, Natural Language Processing

Université de Montréal · McGill · HEC Montréal · Polytechnique Montréal · Université Laval · Université de Sherbrooke

Mila

Scientific Leadership

53
CIFAR Research Chairs

730+
Ongoing research projects

2590+
Scientific papers published since 2018, 35% with industry

ADVANCED COMPUTING INFRASTRUCTURE

Advanced computing infrastructure

Significant contributions to the development and engineering of open source tools and IT infrastructure management

Mila

AI Safety Coalition

2.0 · STRATEGIC COALITIONS

PROPOSAL SUMMARY

How do we design AI systems that reliably follow their developers' alignment instructions without causing harm, whether from malicious requests or uncontrolled goals? The coalition studies the evaluation, interpretability and understanding of unwanted behaviours (including in multi-agent settings), and the design of safer models, harnesses and safeguards.

STRATEGIC RATIONALE

Current training and deployment recipes fail on several alignment and control fronts: jailbreaks, deception, oversight evasion, AI-driven cyberattacks, sycophancy. From a scientific perspective, we need to better understand these alignment, robustness and control issues and find ways to address them. This will help mitigate a broad spectrum of risks and negative societal effects of AI advances and deployment, while ensuring that we can continue to explore the many upsides of AI in a safe and reliable manner.

Leadership & Contributors

11 CORE PROFS

SPONSOR

Yoshua Bengio

PRIMARY CONTRIBUTORS

Guy Wolf
Matt Kusner
Simon Lacoste-Julien

SECONDARY CONTRIBUTORS

Gauthier Gidel
Reihaneh Rabbany
Aaron Courville
AJung Moon
Pierre-Luc Bacon
Siva Reddy
Christopher Pal

Mila

Sound Scientific Discovery Coalition

2.0 · STRATEGIC COALITIONS

PROPOSAL SUMMARY

Can machines verify science at an expert level? The coalition will formalize verification as auditable reasoning over claims and evidence, build verifiers with calibrated confidence that know what they cannot verify, test whether AI discovery loops built on verified claims improve or degrade, and study how scientists trust and override AI verifiers.

STRATEGIC RATIONALE

Generation is outpacing validation. Core profs rated limited metacognition (72%), hallucination (60%) and shallow reasoning (56%) as top AI limitations, and lack of reliability as the leading adoption barrier (60%). An open verifier supports strained peer review, offers a new training signal and positions Mila as a neutral third-party auditor.

Leadership & Contributors

8 CORE PROFS

SPONSORS

Laurent Charlin

Christopher Pal

PRIMARY CONTRIBUTORS

Joelle Pineau

Siva Reddy

Hugo Larochelle

SECONDARY CONTRIBUTORS

Guillaume Lajoie

Jackie Cheung

Aishwarya Agrawal

Mila

Environment and Energy Coalition

2.0 · STRATEGIC COALITIONS

PROPOSAL SUMMARY

Using AI to quantify the world's biodiversity at unprecedented scale; empowering sustainable agriculture, forestry and nature-based climate solutions through remote sensing and smart management; reducing uncertainty and inequity in weather and climate data with AI-integrated forecasting; and accelerating the energy transition through scalable power management and cleantech innovation.

STRATEGIC RATIONALE

Core AI capabilities (OOD generalization, causal and physics-informed ML, explainability, multimodality) are closely linked to environment and energy applications. The coalition embodies Mila's founding principle that AI must be aligned with societal benefit, through interdisciplinary participation in the AI innovation process with problem owners from across use cases of importance to society.

Leadership & Contributors

8 CORE PROFS

SPONSOR

David Rolnick

PRIMARY CONTRIBUTORS

Pierre-Luc Bacon
Alex Hernandez-Garcia
Tegan Maharaj
Derek Nowrouzezahrai

SECONDARY CONTRIBUTORS

Glen Berseth
Hugo Larochelle
Dhanya Sridhar

Mila

Information Integrity Coalition

2.0 · STRATEGIC COALITIONS

PROPOSAL SUMMARY

What shapes what we believe and consider true? As AI becomes the primary interface to information, the coalition studies the full epistemic pipeline (data provenance, truthful and calibrated models, resilience to manipulation and AI swarms, impacts on human belief) and designs AI-facilitated systems for a healthy information ecosystem.

STRATEGIC RATIONALE

Generative models are becoming the main gateway to information, yet they hallucinate, drift and can be poisoned or steered, while making persuasive misinformation cheap at scale. The WEF ranks misinformation as the most severe global threat, and Canadian intelligence reports it already targets elections. Mila can make Montreal a hub for trustworthy information systems.

Leadership & Contributors

7 CORE PROFS

SPONSOR

Reihaneh Rabbany

PRIMARY CONTRIBUTORS

David Adelani

Golnoosh Farnadi

SECONDARY CONTRIBUTORS

Yoshua Bengio

Laurent Charlin

Alex Hernandez-Garcia

Guy Wolf

Mila

Embodied AI and Robotics Coalition

2.0 · STRATEGIC COALITIONS

PROPOSAL SUMMARY

Foundation models are being deployed on robots without the robustness, speed or grounding that automation requires. The coalition asks how to deliver physically faithful low-level control, efficiently combine internet-scale and physical data, transfer knowledge from robotics back to VLMs (e.g., spatial reasoning), and ensure resilient, reproducible and trusted physical AI.

STRATEGIC RATIONALE

73% of core profs identified embodied AI and robotics as a serious limitation of state-of-the-art AI, the top candidate in the forecasting survey. Physical AI is widely seen as the next frontier, attracting massive investment with major economic and workforce impacts. Mila can lead on responsible deployment and keep Canada and Québec at the forefront.

Leadership & Contributors

7 CORE PROFS

SPONSORS

Glen Berseth

Aishwarya Agrawal

PRIMARY CONTRIBUTORS

Liam Paull

AJung Moon

SECONDARY CONTRIBUTORS

Sarath Chandar

Derek Nowrouzezahrai

Amir-massoud Farahmand

Mila

Multi-Agent Systems Coalition

2.0 · STRATEGIC COALITIONS

PROPOSAL SUMMARY

AI is shifting from isolated models to ecosystems of interacting agents. The coalition will build the science of robust multi-agent systems around one unifying concept, asymmetry (in information, capability, compute and incentives), studying emergent specialization, decorrelated failure, multi-agent-native benchmarks and market-based coordination with calibrated self-reporting.

STRATEGIC RATIONALE

Agentic systems are already entering clinical workflows and emergency response, and over 20% of organizations plan to deploy agents (Stanford AI Index 2026), yet benchmarks remain single-agent. If small asymmetric agents can match large models at a fraction of the compute, capable AI becomes accessible to hospitals, smaller institutions and data-sovereign settings.

Leadership & Contributors

6 CORE PROFS

SPONSORS

Gauthier Gidel

Aaron Courville

PRIMARY CONTRIBUTORS

Doina Precup

Aditya Mahajan

Dhanya Sridhar

SECONDARY CONTRIBUTORS

Simon Lacoste-Julien

Mila

NeuroAI Coalition

2.0 · STRATEGIC COALITIONS

PROPOSAL SUMMARY

AI and neuroscience share a core mission: understanding the nature of intelligence. The coalition investigates intelligent systems through biological and synthetic lenses to identify universal principles that hold across substrates, and to map the trade-offs between silicon and biological architectures that will shape next-generation AI.

STRATEGIC RATIONALE

NeuroAI serves three strategic fronts: more resource-efficient architectures and learning rules as scaling meets compute and energy ceilings; a deeper understanding of how human minds and AI work together; and AI for neurology, one of the fastest-growing areas of healthcare research, from brain imaging to brain-computer interfaces.

Leadership & Contributors

5 CORE PROFS

SPONSOR

Guillaume Lajoie

PRIMARY CONTRIBUTORS

Blake Richards
Danilo Bzdok
Tal Arbel

SECONDARY CONTRIBUTORS

Irina Rish

Mila

Continual Learning Coalition

2.0 · STRATEGIC COALITIONS

PROPOSAL SUMMARY

Today's AI systems are static: trained offline and deployed without learning from experience. The coalition asks how to design systems that learn continually without catastrophic forgetting or loss of plasticity, across four axes: optimization and algorithms, architectures and dynamic memory, continual decision-making, and continual alignment and safety.

STRATEGIC RATIONALE

71.9% of core profs agreed continual learning will be a serious limitation of AI systems by 2028, tied for the #1 consensus priority in the survey. Replacing full retraining with incremental updates cuts compute costs and emissions, democratizes AI development beyond a few industrial giants, and strengthens Canada's leadership in the next ML revolution.

Leadership & Contributors

5 CORE PROFS

SPONSOR

Sarath Chandar

PRIMARY CONTRIBUTORS

Amir-massoud Farahmand

Irina Rish

SECONDARY CONTRIBUTORS

Doina Precup

Guillaume Rabusseau