Open Source as a Pathway to AI Adoption in Financial Services

Opportunities and Risks

Andres Rojas

April 14, 2026

VECTOR INSTITUTE | INSTITUT VECTEUR

PART 1

Traditional closed models won't work for AI

The speed of AI development makes going it alone increasingly costly.

But openness also unlocks something closed development cannot:

Trust at scale.

AI algorithms are not IP.
Only trained models are.

NOT YOUR IP

— Transformer architectures

— Attention mechanisms

— Diffusion models

— RLHF & alignment methods

YOUR MOAT

When you train on your data,
you generate model weights
that encode your domain,
your clients, your judgment.

That's the software.
That's the IP.

VECTOR INSTITUTE | INSTITUT VECTEUR

No institution can monitor this field alone.

CONFERENCE	DOMAIN	2024	2025
NeurIPS	General ML	4,497	**5,290**
ICML	General ML	2,609	**3,260**
ICLR	Deep Learning	2,260	**~3,700**
AAAI	General AI	~2,500	**3,032**
CVPR	Computer Vision	2,719	**2,878**
ACL	NLP / LLMs	940	**~1,600**
EMNLP	NLP / LLMs	1,271	**~2,000**
NAACL	NLP / LLMs	565	**719**
IJCAI	General AI	791	**1,023**
KDD	Applied ML / Data	562	**~600**
Top 10 conferences		**~18,700**	**~24,100**

+28% in one year

AI has a trust problem that closed development makes worse.

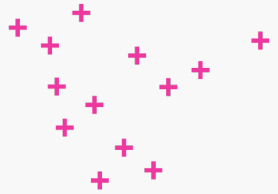

- AI models' decisions are hard to explain and can encode bias
- They can behave differently at scale than in testing, and drift over time without obvious symptoms
- **In a regulated industry, "trust us" is not a governance strategy**
- Regulators are moving toward explainability and auditability
- In 2026, the EU AI Act is live; OSFI, the FCA and others are actively developing AI-specific guidance

The same technique.
Two entirely different domains.

🔥 Wildfire Propagation

Anticipate wildfire spread.
Protect critical infrastructure
with bank exposure.

Each node's forecast influenced by its neighbours.

GRAPH-BASED SPATIAL FORECASTING

🏧 ATM Cash Demand

How much cash to hold
at each machine, given what
neighbouring ATMs are doing.

Each node's forecast influenced by its neighbours.

Neither team invented the technique. Both created real value by applying it. You are not competing on the engine.

Same institution. Same technology. Different risk. Different answers.

Portfolio Allocation Model

Core to strategy. High adversarial surface.
Should almost certainly stay closed.

→ CLOSE

Chatbot Bias Filter

Low strategic sensitivity.
Transparency required for credibility.

→ OPEN

The question is not open vs. closed — it is whether you have the infrastructure to tell the difference.

PART 2

There are proven strategies for working collaboratively with AI

The structural blockers are shared.
Solving them in isolation seems irrational.

VECTOR INSTITUTE | INSTITUT VECTEUR

Most institutions spend energy on all three layers. They should only spend it on the top one.

YOUR COMPETITIVE EDGE

Proprietary data · Domain expertise · Client relationships · Institutional judgment

CO-DEVELOPMENT LAYER

Model validation frameworks · Explainability toolkits · Governance templates · Reference implementations

ALGORITHM + CLEAN REFERENCE IMPLEMENTATION

Not just the paper — working code that makes a technique accessible, not just known. Shared openly.

GUARD ↑

SHARE ↓

VECTOR INSTITUTE | INSTITUT VECTEUR

Vector's co-development model: from published technique to working proof of concept.

01
Identify
Monitor thousands of papers. Surface relevant techniques early.

02
Implement
Build vanilla code + sample data + compute. Partners start running, not reading.

03
Discover
Community finds use cases. Open-ended — surfaces what Vector wouldn't anticipate.

04
Build PoC
Each team builds in a shared environment with guidance and lateral knowledge flow.

05
Share & Deploy
Daily debriefs. Showcase day. Then each org builds production in-house.

Collaborative →→→ Proprietary

VECTOR INSTITUTE | INSTITUT VECTEUR

Partners are not starting from zero.

6 months → 6 weeks

Reducing the exploration cycle

Maximizing the opportunity

200+
organizations

550+
use cases

250K
hours co-developed

200+
reference implementations

What Vector absorbs on behalf of all participants:

- Early identification of relevant techniques from thousands of papers
- Workshops and tutorials translating theory into practice
- Working starter code + sample datasets + configured compute
- Project management, technical guidance, and knowledge facilitation

PART 3

Being open is not binary
Prepare accordingly

The risk profile of openness depends on the model, the use case, and the institution's governance maturity.

Degrees of Open Source AI

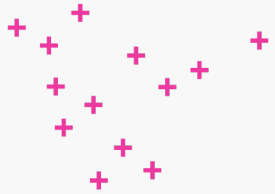

Open Publication
Ideas, papers, conferences. Already universal.

No risk

Open Benchmarks
Shared evaluation standards. Low risk, high value for model validation.

Minimal

Open Code
Working implementations without trained weights. Risk profile similar to open-source software.

Familiar

Open Datasets
Valuable but constrained by privacy regulation and competitive sensitivity.

Context-dependent

Open Weights - Foundation
Trained model released publicly. Provenance and control questions arise.

Elevated

Open Weights - Fine-tunable
Base is shared; your specialisation is yours. Provenance discipline essential.

Managed carefully

VECTOR INSTITUTE | INSTITUT VECTEUR

Openness should be the outcome of a structured risk-benefit analysis.

Portfolio Allocation Model

Core to strategy. Openness exposes the strategy itself.

CLOSED

AML Transaction Monitoring

Auditability favours openness. Adversarial exposure favours caution. No universal answer.

CAREFUL ANALYSIS

Chatbot Bias Detection

Transparency is essentially required for credibility. Community scrutiny improves robustness.

OPEN

The same institution. Three different answers. All correct.

VECTOR INSTITUTE | INSTITUT VECTEUR

The real vulnerability is not open vs. closed. It's the absence of governance infrastructure.

01 **Model classification framework**

Consistent criteria for assessing strategic sensitivity, adversarial exposure, and regulatory profile

02 **Openness policy by tier**

Not one policy for all models — a set mapped to classification, with licencing and review requirements

03 **Provenance & supply chain controls**

Knowing exactly what went into every model you deploy — whether built internally, co-developed, or on open weights

04 **Contribution governance**

Policies for what teams can share externally, under what licences, with what internal review

05 **Ongoing review**

The risk-benefit calculus changes as models become more central and threats evolve — classification must be revisited

For the right models, transparency is not a risk to manage — it's a strategy.

⚖️

Regulatory Trust

Explainable, auditable models are faster to approve. In 2026, some regulators are moving from preference to requirement.

🎓

Talent

Researchers want to publish, contribute to open problems, and build on community knowledge. Closed systems are a recruitment and retention disadvantage.

⚡

Speed & Robustness

A community stress-tests your model faster, more cheaply, and more adversarially than any internal QA process.

Transparency isn't something you offer despite the risks. For the right models, it is how you manage risks you already have.

Canada is not just a participant in this ecosystem.

- Canada trained the researchers who launched modern AI
- Vector Institute has spent 6 years proving that structured co-development works at scale
- The opportunity: define the governance standards and trust infrastructure the world will need

That work starts here

Open models don't ask for blind trust.

They can be verified.

- Build on open foundations: the algorithms are already shared
- Invest your edge in data, domain expertise, and good governance
- Join communities that accelerate trust