



AI Is Not Rocket Science: Ideas for Achieving Liftoff in Canadian AI Adoption

By Kevin Leyton-Brown, Cinda Heeren, Joanna McGrenere, Raymond Ng,
Margo Seltzer, Leonid Sigal, and Michiel van de Panne

Prepared from remarks made at the C.D. Howe Institute and CIFAR Conference “Artificial Intelligence – Seizing Opportunities” held on May 27, 2025.

- Canada is a global leader in AI research, ranking second worldwide in top-tier researchers and first in the G7 for per capita academic AI papers. Yet adoption lags: in 2023, Canada ranked 20th in AI adoption among OECD countries.
- The federal government’s 2017 Pan-Canadian AI Strategy strengthened academic research capacity through institutes in Toronto, Montreal, and Edmonton. Other programs – such as SR&ED tax incentives, superclusters, and matching grants – have supported industry adoption, but collectively they have not overcome Canada’s persistent adoption deficit.
- A central challenge is that AI is not “rocket science.” Unlike disciplines tied to a single industry, AI is a general-purpose technology. Because AI researchers specialize in methodologies rather than specific sectors, academics and industry often struggle to connect and sustain effective partnerships.
- We describe four ideas for increasing knowledge translation from AI researchers to Canadian industry. First, a concierge service to match industry needs with academic expertise. Second, consulting in exchange for student scholarships, creating incentives for researchers, while offering practical support to firms. Third, funding “research trios” that combine an AI expert, a domain expert, and an industry partner to tailor applications to specific sectors. Finally, a dramatic expansion of AI training, from basic literacy to dedicated degrees and continuing education, to meet the fast-rising demand for expertise. Together, these measures aim to transform Canada’s academic strength in AI into broad-based industrial adoption.

1 INTRODUCTION

Canada has an enviably strong university sector, so it is perhaps unsurprising that the country is recognized as a global leader in AI research. Indeed, the roots of modern AI were laid by three Canadian academics: Geoff Hinton at the University of Toronto; Yoshua Bengio at the Université de Montreal; and Rich Sutton at the University of Alberta.¹ The federal government’s 2017 Pan-Canadian

1 Of course, every success has many mothers. Hinton and Bengio shared their Turing Award with Yann LeCun at New York University. Rich Sutton shared his with Andrew Barto at the University of Massachusetts. Many feel that German computer scientist Jürgen Schmidhuber has not received the credit he deserves; still others have made critical contributions. Overall, however, AI’s deep roots in Canada are undisputed.

AI Strategy (CIFAR 2017) was one of the world's first federal initiatives to strengthen a country's AI research capacity, funding AI Institutes in Toronto, Montreal, and Edmonton. Partly reflecting this investment, in 2024 the Organisation for Economic Co-operation and Development (OECD) ranked Canada second globally in terms of the number of top-tier AI researchers, with 10 percent of the total, and first in the G7 in per capita number of academic AI papers over the previous five years (Dobbs and Hirsch-Allen 2024).

Despite its academic strength in many sectors, Canada has long lagged the OECD in translating academic research into industrial application (Thirion 2025; Council of Canadian Academies 2018). This matters for the economy as a whole, because such knowledge translation is a key vehicle for productivity growth: between 1981 and 2019, 92.5 percent of Canadian per capita GDP growth was attributable to productivity growth. It is terrible news, then, that Canada experienced almost no productivity growth in the last decade (2015–2024), versus a 15 times higher rate in the US (Accenture 2024).

Given this context, it may not be surprising that Canada is widely recognized to be lagging in AI adoption. For example, a 2023 study based on 2021 data ranked Canada 20th in AI business adoption among OECD countries (Lockhart 2023). It is widely anticipated that AI will be one of the primary drivers of productivity growth in the decades to come,

as AI is integrated into industrial sectors across the whole economy. It is thus critical that Canada avoid repeating its historical patterns as this new industrial era takes shape.

To do better, we must understand that AI is not rocket science. After all, rocket science studies rockets. In contrast, AI studies the development and application of machine intelligence. This simplistic observation highlights a fundamental truth that should drive the way we think about AI adoption policy in Canada. Rocket science is focused on building one specific kind of thing. The field's industrial outlets are obvious (the aerospace industry, for example), and relationships between academics and industry² representatives develop and sustain over a long period of time. Rocket science is a good placeholder for a wide range of scientific and engineering disciplines, in the sense that it is focused on a body of knowledge that feeds into a single, well-defined



2 We use the term “industry” to refer to organizations that might benefit from adopting AI. While many such organizations are for-profit companies, others belong to the nonprofit and government sectors.

This and all other illustrations were generated by OpenAI's 4o model. We hope they make the article more engaging to read. Slight inconsistencies and imperfections are part of the fun.

About the Authors: Kevin Leyton-Brown is a professor of computer science at UBC; Cinda Heeren is a professor of teaching at UBC; Joanna McGrenere is a co-head and professor of computer science at UBC; Raymond Ng is a professor of computer science at UBC; Margo Seltzer is a co-head and professor of computer science at UBC; Leonid Sigal is a professor of computer science at UBC; and Michiel van de Panne is a professor of computer science at UBC.

industrial sector. Other good examples include pharmacology, forestry, radiology, and mining; the reader can surely think of others. AI is different. Its promise and its curse are that it can be applied to just about anything. The promise is that AI's potential for economic transformation is immense. The curse is that applied research always requires deep knowledge of the application area, and AI experts do not have the luxury of picking a single application area and sticking to it.

One might imagine that the solution would be to employ a small number of AI experts who advance the core technology, plus a large army of applied AI researchers who specialize in translating this technology into particular applications, whose nuances they carefully learn. Unfortunately, this model has not addressed the AI adoption lag to date and seems unlikely to improve AI technology transfer in the near future. Instead, deep knowledge of technical foundations has been and remains critical for making good choices about which AI technology to deploy and in what way. There are two key reasons.

First, progress in AI is extremely rapid: today's mainstream technologies didn't even exist a few years ago. People who are focused only on applying AI struggle to stay on top of such changes in the field's methodologies, and hence fall behind quickly.

Second, AI is an amazingly general-purpose technology that does not respect application boundaries. Even some of its most stunning breakthroughs arose not by drilling into applications but from research that initially focused on entirely different domains. For example, DeepMind's work on using reinforcement learning to play Atari games led in a few years to AlphaFold, its system for



predicting protein folding that yielded a Nobel Prize in Chemistry. And while code synthesis has been a topic of specialist research for years, the breakthrough that gave rise to “vibe coding” came when OpenAI and others built general-purpose chatbots trained on the entire internet. Overall, steering AI adoption requires strength in the core discipline. Or as a yoga guru might say, the foundation for success is core strength.

2 WHAT'S AN ASPIRING AI ADOPTER TO DO?

A hypothetical example will help us describe some of the key issues and illustrate how frustrating they make the AI adoption problem for industry. Imagine that you are the CEO of a Vancouver-based startup promoting the use of AI-enhanced, digital medical records across a mental-health patient's multiple touch points with the health system. You recognize that your pre-ChatGPT AI approach is no longer internationally competitive; even your provincial contacts are starting to ask hard questions in client meetings. So you'd like to up your game, and your first step is to reach out for some advice. How do you go about it?

After some Google searching, you'd realize that “AI + Health” isn't a huge area of research within AI.³ For example, a small fraction of papers at top

3 Dedicated AI + Health initiatives may well be established in Canada. We don't think such efforts would eliminate the issue that most AI research is general purpose and that it is hard to establish the right connections between academics and industry. But this example is nevertheless something like a best-case scenario. Consider other application areas that might have mattered to a Vancouver CEO, like AI + Critical Minerals Exploration, AI + Forestry, AI + Legal Services, etc. In those cases, the CEO's Google searches would have turned up even fewer examples of cutting-edge, applied AI research.

conferences – which turns out to be a good gauge of topic prioritization within the field – discuss any form of healthcare. While there are a smattering of researchers who do claim this specialization, they tend to have specific areas of focus (e.g., cancer diagnosis; drug development; causal analysis in longitudinal studies) that do not overlap in any obvious way with your company’s mission. Conversely, the academic papers on the AI tools used by your competitors don’t discuss medicine at all, let alone medical records. If you took the time to look through the web pages of academics in top AI departments in Vancouver or around the world (or if you used an AI tool like DeepResearch to do this for you), you’d find out that, while AI experts do a lot of applied work, they rarely describe their research expertise in terms of applications, and indeed they leave one application area behind and parachute into another on a regular basis. How then should you find the right person to help you with your problem?

Let’s say your niece is a computer science PhD student at the University of British Columbia (UBC), and you take her out to lunch to learn more. While she doesn’t work in AI herself, she knows enough to clear up some of your confusion. She tells you that progress in the academic field of AI has been accelerating over the past couple of decades, and at this point, it is moving unprecedentedly quickly; sometimes multiple generations of the same technical idea succeed each other within the same year! It is thus impossible for any single person even to stay current across the board. AI experts who want to advance the frontier specialize in specific AI modalities⁴ and indeed in specific methodological approaches to their given modality.⁵

The good news is that – far from needing some kind of inducement to climb down from their ivory towers – AI experts are eager to work on applications, because performance on real-world data is the yardstick by which research ideas are measured. So your niece tells you not to despair: your medical records domain is indeed data rich, and deploying modern AI is your whole focus. Plus, she knows from her own advisor that Canadian academics are deeply reliant on industry partnerships to hire graduate students, so she’s sure that you won’t get the cold shoulder if you reach out.

This leaves you feeling hopeful, but also more frustrated than before. It’s clear that AI could profoundly help your business; it’s what your customers are demanding. Vancouver (and Canada more broadly) is full of AI expertise, and academics are apparently eager to hear from you. But who should you reach out to? The web pages of AI professors are a sea of technical jargon. Furthermore, you’ve got a suspicion that someone who might jump at the chance to work with you might not be a leader in their own field, while the perfect person might need some persuading that your medical records startup is the right partner for *them*. You don’t understand enough about AI to make that case. What should you do?



-
- 4 Some of what we’re calling AI modalities include computer vision, natural language processing, probabilistic reasoning, robotics, multiagent systems, and theoretical foundations.
- 5 Within computer vision, researchers might focus on still images versus video; drawing conclusions about existing data versus generating new, realistic content; adding capabilities to models versus deploying existing approaches at scale; working from static data versus making models that evolve as new data comes in. We could give dozens of additional examples for computer vision and analogous lists for any other modality.



3 EXISTING MODELS FOR PROMOTING AI ADOPTION

Keeping this example in mind, let us discuss three different strategies that the Canadian government employs to leverage academic strength in a research area to boost industrial uptake, and consider the impact of each on AI adoption.

First, governments focus on strengthening the academic side of the equation, reasoning that a healthy academic ecosystem offers trickle-down benefits to industry. There is truth to this idea: many startups arise organically; serendipitous conversations between academics and industry representatives produce breakthroughs; a healthy stream of graduating students is critical for sustaining a local industrial base. However, while a strong academic sector is surely helpful to industry, the entire premise of this article is that it is not sufficient. The “trickle-down” pace is far too slow. Canada’s academic strength, but industrial underperformance in AI, shows that academic investments alone do not guarantee productivity gains. Generous Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant funding for professors would not help the CEO in our example.

Second, the government might directly subsidize industry, seeking to eliminate barriers to the adoption of new technologies. This is what Canada has done for many years through Scientific Research and Experimental Development (SR&ED) tax

incentives (which cover only a fraction of companies’ research and development investments); technology supercluster programs; special-purpose deals for large companies considering opening a new campus in a given location; and so on. As in the previous case, this strategy is not without merit. At the margin, economic incentives can surely boost investment and otherwise nudge a company towards action. But there are two significant drawbacks to the approach. First, despite decades of employing it, Canada is nevertheless experiencing abysmal productivity growth; this approach is clearly not a silver bullet. Second, it fails to capitalize on Canada’s biggest advantage in AI, which is the strength of its academic sector. Again, while our fictitious CEO might appreciate a SR&ED tax credit, this wouldn’t directly help them to figure out how to apply modern AI in their product.

What remains is for the government to seek to boost industry adoption by leveraging academic strength. For decades, this has been a major focus of Canadian industrial policy. The approach is perhaps exemplified by the NSERC Alliance program and by nonprofit “centre of excellence” programs such as Mitacs. At a high level, both programs incentivize professors to work with industrial partners as a way of funding their research. The model is typically that professors initiate these collaborations, motivated by the norm that a successful career requires a larger complement of graduate students than can be funded

on an NSERC Discovery Grant, student scholarships, and teaching assistantships alone, and by the reality that few alternative sources of funding exist. Industry partners are asked to make some combination of a cash and an in-kind contribution to the project, which is supplemented by matching funds from the government and generous intellectual property (IP) terms.


While the Mitacs/Alliance model clearly succeeds in catalyzing some knowledge transfer from academia to industry, it is not a one-size-fits-all solution. On the industry side, it does not work well for those who need only straightforward advice rather than the year-long research projects the model is set up to provide. Students are not always the ideal points of contact for industry: in addition to their comparatively lower levels of technical expertise, they can lack a broad view of their field and have less experience communicating technical ideas with non-experts. On the faculty side, it can also be a poor fit. Faculty who oversee projects often lack deep domain knowledge of the downstream industry, particularly for non-rocket-science disciplines like AI. Finally, application deadlines, the uncertainty associated with peer review, paperwork, and intellectual property negotiations add significant frictions to the process.

Faculty need to support their research programs, and Mitacs/Alliance provide the most readily available funds. This often means there is more interest in such collaborations from academics than there is industry demand. This lost opportunity is profound. While unmet demand for research funding is often framed as a problem for professors, the reality is that programs are at least partly motivated by a desire to foster technology transfer, and nearly all money flows to trainees. This means that when funding is scarce, fewer students are trained, which makes little sense: the marginal annual cost of funding all of the

additional graduate students a professor would be willing to supervise (a few tens of thousands of dollars each) is small relative to the annual cost already being paid

to employ the professor and house their research program, which goes beyond salary to include the indirect costs of administrative and technical staff support plus the facilities that the professor needs to do their job (hundreds of thousands of dollars per professor). And Canada has a huge need for skilled trainees! It is as though the BC government had built a modern hospital, purchased a state-of-the-art MRI machine, but then only funded it at a level that allowed it to see patients for five hours a day. Making matters worse, Canadian funding schemes are often regionally based (e.g., the Pan-Canadian AI Strategy makes it difficult to adequately fund AI research outside Toronto, Montreal, or Edmonton). This means that the unserved gap between student demand and student funding is unevenly distributed across the country.

Finally, as our example illustrated, matchmaking between academics and industrial partners is very hard. Due to the lopsided incentive structure just described, the task falls mostly on academics. The problem is that few of them are particularly good at it – the academic pipeline does not select for people with the ability to cold-call small-company CEOs and persuade them to part with money to co-fund research projects. This task is complicated even further when AI experts do not focus on single domains, and virtually any industry in town could be a potential partner.



4 IDEAS FOR ACHIEVING LIFTOFF IN CANADIAN AI ADOPTION

It is possible to do better. This section describes four synergistic ideas for paths forward that could help our hypothetical CEO and others like them to leverage Canadian academic strength to fuel their AI adoption. Deploying them together would yield maximum impact.

4.1 Idea 1: A Concierge Service for Matchmaking

We have seen that it is hard for industry partners to know who to contact when they want to learn more about AI. Conversely, it is at least as hard for AI experts to develop a broad enough understanding of the industry landscape to identify applications that would most benefit from their expertise. Given the potential gains to be had from increasing AI adoption across Canadian industry, nobody should be satisfied with the status quo.

We argue that this issue is best addressed by a “concierge service” that industry could contact when seeking AI expertise. While matchmaking would still be challenging for the service itself, it could meet this challenge by employing staff who are trained in eliciting the AI needs of industry partners, who understand enough about AI research to navigate the jargon, and who proactively keep track of the specific expertise of AI researchers across a given jurisdiction. This is specialized work that not everyone could perform! However, many qualified candidates do exist (e.g., PhDs in the mathematical sciences or engineering). Such staff could be funded in a variety of different ways: for example, by an AI institute; a virtual national institute focused on a given application area; a university-level centre like UBC’s Centre for Artificial Intelligence Decision-making and Action (CAIDA); a nonprofit like Mitacs; a provincial ministry for jobs and economic growth; or the new federal ministry of Artificial Intelligence and Digital Innovation.

Having set up an organization that facilitates matchmaking, it could make sense for the same office to provide additional services that speed AI adoption, but that are not core strengths of academics. Some examples include project management, programming, AI-specific skills training and recruitment, and so on. Overall, such an organization could be funded by some combination of direct government support, direct cost recovery, and an overhead model that reinvests revenue from successful projects into new initiatives.



4.2 Idea 2: Consultancy in Exchange for Student Scholarships


Many businesses that would benefit from adopting AI do not need custom research projects and do not want to wait a year or more to solve their problems. The lowest-hanging fruit for Canadian AI adoption is ensuring that industry is well informed about potentially useful, off-the-shelf AI technologies. We thus propose a mechanism under which AI experts would provide limited, free consulting to local industry. AI experts would opt in to being on a list of available consultants. A few hours of advice would be free to each company, which would then have the option of co-paying for a limited amount of additional consulting, after which it would pay full freight if both parties wanted to continue. The company would own any intellectual property arising from these conversations, which would thus focus on ideas in the public domain. If the company wanted to access university-owned IP, it could shift to a different arrangement, such as a research contract. This system

would work best given a concierge service like the one we just described. The value offered per consulting hour clearly depends on the quality of the academic–industry match, and some kind of vetting system would be needed to ensure the eligibility of industry participants.

Why would an AI expert sign up to give advice to industry? All but the best-funded Canadian faculty working in AI report that obtaining enough funding to support their graduate students is a major stressor. Attempting to establish connections with industry is hard work, and such efforts pay off only if the industry partner signs on the dotted line and matching funds are approved. There is thus space to appeal to faculty with a model in which they “earn” student scholarships for a fixed amount of consulting work. For example, faculty could be offered a one-semester scholarship for every eight hours set aside for meetings with industry, meaning that one weekly “industry office hour” would indefinitely fund two graduate students. Consulting opportunities could also be offered directly to postdoctoral fellows or senior (e.g., post-candidacy) PhD students in exchange for fellowships. In such cases, trainees should be required to pass an interview, certifying that they have both the technical and soft skills necessary to succeed in the consulting role. The concierge service could help decide which industry partners could be routed to PhD students and which need the scarcer consulting slots staffed by faculty members.

The system would offer many benefits. From the industry perspective, it would make it straightforward to get just an hour or two of advice. This might often be enough to allow the company to start taking action towards AI adoption: there is a rich ecosystem of high-performance, reliable, and open-source AI tools; often, the hard part is knowing what tool to use in what way. Beyond the value of the advice itself, consulting meetings offer a strong basis for building relationships between academics and industry representatives, in which the academic plays the role


of a useful problem solver rather than of a cold-calling salesperson. These relationships could thus help to incubate Mitacs/Alliance-style projects when



research problems of mutual interest emerge (though also see our idea below about how restructuring such projects could help further).

For academics, the system would constitute a new avenue for student funding that would reward each hour spent with a predictable amount of student support. Furthermore, it would offer scaffolded opportunities to deepen connections with industry. The system would come with no reporting requirements beyond logging the time spent on consulting. The faculty member would be free to use earned scholarships to support any student (regardless, for example, of the overlap between the student’s research and the topics of interest to companies), increasing flexibility over the Mitacs/Alliance system, in which specific students work with industry partners. Students who self-funded via consulting would learn valuable skills and would expand their professional networks, improving prospects for post-graduation employment.

Finally, the system would also offer multiple benefits from the government’s perspective. It would generate unusually high levels of industrial impact per dollar spent (consider the number of contact hours between academia and industry achieved per dollar under the funding models mentioned in Section 3). All money would furthermore go towards student training. The system would automatically allocate money where it is most useful, directing student funding to faculty who are both eager to take on students and relevant to industry, all without the overhead of a peer-review process. And it would



generate detailed impact reports as a side effect of its operations, since each hour of industry–academia contact would need to be logged to count towards student funding.


4.3 Idea 3: Grants for Research Trios

Our third proposal is an approach for expanding the Mitacs/Alliance model to make it work better for AI. Industry–academia partnerships leverage two key kinds of expertise from the academic side: methodological know-how for solving problems and knowledge about the application domain used for formulating such problems in the first place. In fields for which the set of industry partners is relatively small and relatively stable, it makes sense to ask the same academics to develop both kinds of expertise. In very general-purpose domains like AI, it holds back progress to ask AI experts to become domain experts, too. Instead, it makes sense to seek domain knowledge from other academics who already have it. We thus propose a mechanism that would fund “research trios” rather than bilateral research pairings. Each trio would contain an AI expert, an academic domain expert, and an industry partner. This approach capitalizes on the fact that there is a huge pool of academic talent outside core AI with deep disciplinary knowledge and a passion for applying AI. While such researchers are typically not in a position to deeply understand cutting-edge AI methodologies, they are ideally suited to serve as a bridge between researchers focused on AI methodologies and Canadian industrial players seeking to achieve real-world productivity gains. In our experience at UBC, the pool of non-AI domain experts with an interest

in applying AI is considerably larger than the pool of AI experts. One advantage of this model is that projects can be initiated by the larger population of domain experts, who are also more likely to have appropriate connections to industry. Beyond this, involving domain experts increases the likelihood that a project will succeed and gives industry partners more reason to trust the process while a solution is being developed. The model meets a growing need for funding researchers outside computer science for projects that involve AI, rather than concentrating AI funding within a group of specialists. At the same time, it avoids the pitfall of encouraging bandwagon-jumping “applied AI” projects that lack adequate grounding in modern AI practices. Finally, it not only transfers AI knowledge to industry, but also does the same to both the domain expert and their students.

4.4 Idea 4: Greatly Expanded AI Training

As AI permeates the economy, Canada will face an increasing need for AI expertise. Today, that training comes mostly in the form of computer science degrees. Just as computer science split off from mathematics in the 1960s, AI is emerging today as a discipline distinct from computer science. In part, this shift is taking the form of recognizing that not every AI graduate needs to learn topics that computer science rightly considers part of its core, such as software engineering, operating systems, computer architecture, user interface design, computer graphics, and so on. Conversely, the shift sees new topics as core to the discipline. Most fundamental is machine learning. Dedicated training in AI will require a deeper focus on the mathematical foundations of probability and statistics, building to advanced topics such as deep learning, reinforcement learning, machine learning theory, and so on. Various AI modalities also deserve separate study, such as computer vision, natural language processing, multiagent systems, robotics, and reasoning. Training



in ethics, optional in most computer science programs, will become essential.

Beyond dedicated training in the core discipline, we anticipate huge demand for

broad-audience AI literacy training; for AI minors to complement other disciplinary specializations; for continuing education and “micro-credential” programs; and for executive education in AI. There is also a growing need for “AI Adoption Facilitators”: bridge-builders who can help established workers in medium-to-large organizations understand how data-driven tools could offer value in solving the problems they face. Training for this role would emphasize business principles and domain expertise, but would also require firmer foundations in machine learning and data science than are currently typical in those disciplines.

5 HOW WE’RE ACTING ON THIS VISION AT UBC

UBC has been a Canadian AI research powerhouse since at least the 1990s and remains one today. Because Vancouver was not one of the cities favoured by the Pan-Canadian AI Strategy, we have taken a different path than our colleagues elsewhere in the country, developing a distinct vision for fostering AI adoption. In what remains, we discuss the AI landscape in Vancouver and some of the ways that we have already begun to build out our ideas.

5.1 Vancouver’s Strength in AI

Two objective measurements illustrate Vancouver’s national standing in AI research. First, a Canada CIFAR AI Affiliate Chair is a recognition of excellence for Canadian AI researchers working

at institutions outside the three cities that host AI institutes: effectively, all Canadian research universities other than the University of Toronto, York University, Université de Montréal, McGill University, and the University of Alberta. Over half of these chairs are held by researchers in Vancouver, and over 40 percent at UBC alone (CIFAR 2025). Second, two prominent 2024 rankings of AI research at Canadian universities placed UBC second nationally (Chaitanya 2024; Infosource 2024).

Beyond its strength in AI research, Vancouver also has one of Canada’s largest tech sectors, partly owing to its proximity to tech hubs in Seattle, Silicon Valley, and Asia. A 2025 analysis found that Vancouver ranked 10th in North America (and 3rd in Canada) in terms of its tech sector’s “depth, vitality and attractiveness to companies seeking tech talent and to tech workers seeking employment” (CBRE 2025). This study also ranked Vancouver 2nd in Canada (after Toronto) for the total number of AI specialists in the tech industry, and 3rd in North America (after the San Francisco Bay Area and Seattle) in terms of the share of its tech workforce employed in the tech industry – a driver of innovation.

Indeed, although British Columbia has the reputation of being economically driven by natural resource extraction, BC’s tech sector employs six times as many people as the natural resources sector, and by itself accounts for between 38 percent and 50 percent of BC’s exports (Tipping 2025). The world’s leading AI research conferences are furthermore regularly hosted in Vancouver; over just the past couple of years, this has included the International Conference on Machine Learning (ICML 2025), the Conference on Neural Information Processing Systems (NeurIPS 2024), the Association for the Advancement of Artificial Intelligence Conference (AAAI 2024), and the Conference on Computer Vision and Pattern Recognition (CVPR 2023). The question about how best to foster AI adoption is thus a critically important one in the local British Columbia context, just as it is nationally.



5.2 How Our Experience at UBC Shaped This Proposal

An inevitable downside to the AI institute model implemented in Toronto, Montreal,

and Edmonton is that it has created a sharp delineation between those who receive generous funding and other support for their research and those who do not. Our relative dearth of funding at UBC has spared us from the need to make such sharp distinctions, and so we have focused on supporting the broadest possible coalition of researchers who self-identify as focusing on AI and on developing connections with companies in our local ecosystem. This has led us to two major takeaways. First, local companies find it very difficult to connect with experts who have the right expertise, and many (at least initially) want quick answers rather than long-horizon research projects. Second, there is a rapidly growing group of academic researchers outside of core AI who are passionate about applying the technology in their own domains, and many see the difficulty of establishing collaborations with AI experts as a key impediment to their work.

Responding to these needs, we have established three initiatives at UBC that foster connections between AI experts and domain experts. First, UBC's analogue to an AI institute is CAIDA, founded in 2018 (CAIDA 2025a). To date, this centre has attracted 143 UBC faculty as members, spanning 27 departments from across the university (CAIDA 2025b). This group is subdivided into 34 core members (roughly, people who would surely hold an AI chair if they worked elsewhere) and 109 non-core members (in most cases, academics from disciplines outside computer science and statistics who apply AI to domains in which they have primary expertise;

elsewhere, some of these people would hold AI chairs too). Second, we have established a Faculty of Science Hiring Cluster called AI Methods for Scientific Impact (AIM-SI), an initiative which has hired six new faculty since 2022, spanning computer science, statistics, and mathematics (AIM-SI 2025). The cluster is explicitly focused on fostering partnerships between core AI researchers and experts in other scientific disciplines. Third, UBC's Data Science Institute (DSI) funds postdoctoral fellows to serve as a bridge between data science experts, other academic domain experts, and industrial partners (DSI 2025). Over the past decade, that DSI funding program has brought back external research funds almost four times as much as DSI has spent on the program, demonstrating the value of incubating research trios.

DSI has also piloted a bootcamp initiative that has informed our consultancy idea. For example, in summer 2025, DSI hosted a bootcamp on interpretable machine learning. As one component of the program, graduate students specializing in the area met with selected industrial partners who presented problems they wanted to solve and contributed pilot data. The students worked full-time for a few weeks with guidance from their supervisors. Their results were presented to the industrial partners; several of these partnerships have led to ongoing partnerships and to work proceeding towards publication.

UBC has also taken significant steps to advance AI education. Most notably, UBC is making a proposal to the Province of British Columbia to establish a Bachelor of Science in AI: Canada's first such dedicated undergraduate degree. In addition to allowing students to deep-dive into the core discipline, the degree will also offer a pathway dubbed "BSc AI+ x ." After fulfilling the core degree requirements, students will take three specially designed courses in a second field x , where x might be medicine, economics, forestry, media studies, etc. This program will help to train the sorts of bridge builders who will help various sectors of the Canadian economy to adopt AI technologies. We are also

designing a 100-level (introductory) university-wide AI literacy course. Just as it is broadly understood that educated people should understand the basics of psychology, economics, and chemistry, we have reached a time when a conceptual understanding of AI is critical for even a nonspecialist audience.


6 CONCLUSIONS

AI offers promise for addressing Canada's productivity crisis. But AI is not rocket science: it's a science applicable to almost everything. Because the field is changing so quickly and because AI methodologies are profoundly general-purpose, an application-focused strategy is insufficient. Instead, AI adoption efforts must draw on expertise from core AI researchers. Canada is fortunate to have unusual strength in this area. However, the existing Mitacs/Alliance model of academic-industrial partnerships does not fully meet the needs of disciplines that impact a very diverse set of industrial sectors. New ideas and complementary mechanisms are therefore needed to accelerate AI adoption across Canadian industry.

This paper proposed four key ideas that could synergistically bridge this gap. First, both AI researchers and companies struggle to identify counterparties with whom they can fruitfully collaborate. A centralized matchmaking service would open the door to many partnerships that would otherwise never be established. Second, not every industry inquiry is a request to establish a research project. Simple consultations about public-domain AI knowledge would be hugely valuable to industry and, if appropriately incentivized, attractive to academics.

Third, in situations where longer-term research projects do make sense, the lack of domain knowledge by AI experts is nevertheless a barrier.

Better outcomes can be achieved by building bigger teams that include academic domain experts from outside AI in "research trios." Finally, we're now experiencing only the tip of the iceberg in demand for AI expertise. As the sector grows in economic importance, meeting the resulting demand will require considerable innovation in AI education. This will span a wide spectrum: every student should be offered the opportunity to develop AI literacy; those seeking to specialize in AI should be able to pursue dedicated courses of study; students in other disciplines should be able to develop sub-specializations in AI; graduates in the workforce should be offered opportunities for both continuing education to allow them to apply AI in their existing career paths and retraining to help them to chart a new course. At UBC, we're already starting to build out the vision described here. To go further will require sustained support from government, industry, non-profits, and foundations. Canada's current strength in AI is only the beginning, and that our country is ready to enter a new chapter enabled by new models of academic-industry partnership.



REFERENCES

- AAAI. 2024. “The 38th Annual AAAI Conference on Artificial Intelligence.” <https://aaai.org/conference/aaai/aaai-24/>.
- Accenture. 2024. “Canada’s Generative AI Opportunity.” <https://www.microsoft.com/en-us/industry/microsoft-in-business/wp-content/uploads/sites/28/2024/06/Canadas-Generative-AI-Opportunity-White-Paper-FINAL-English.pdf>.
- AIM-SI. 2025. “AI Methods for Scientific Impact.” <https://caida.ubc.ca/aim-si>.
- CAIDA. 2025a. “UBC ICICS Centre for Artificial Intelligence Decision-making and Action.” <https://caida.ubc.ca/>.
- _____. 2025b. “Faculty Members of CAIDA.” <https://caida.ubc.ca/faculty-members-caida>.
- CBRE. 2025. “Scoring Tech Talent 2025.” <https://www.cbre.com/insights/books/scoring-tech-talent-2025>.
- Chaitanya, Sai. 2024. “Analytics Insight: Leading Universities in Canada for AI and Data Science.” <https://www.analyticsinsight.net/artificial-intelligence/leading-universities-in-canada-for-ai-and-data-science>.
- CIFAR. 2017. “Pan-Canadian AI Strategy.” <https://cifar.ca/ai/>.
- _____. 2025. “Canada CIFAR AI Chairs.” <https://cifar.ca/ai/canada-cifar-ai-chairs/>.
- Council of Canadian Academies. 2018. *Competing in a Global Innovation Economy: The Current State of R&D in Canada*. Ottawa: Expert Panel on the State of Science and Technology and Industrial Research and Development in Canada, Council of Canadian Academies. https://cca-reports.ca/wp-content/uploads/2018/09/Competing_in_a_Global_Innovation_Economy_FullReport_EN.pdf.
- CVPR. 2023. “The IEEE/CVF Conference on Computer Vision and Pattern Recognition 2023.” <https://cvpr.thecvf.com/Conferences/2023>.
- Dobbs, Graham, and Jake Hirsch-Allen. 2024. “Canada’s Plans to Bridge the AI Compute Gap and How It Can Make Industry Policy Inclusive and Sustainable.” OECD AI Policy Observatory Blog. April 16. <https://oecd.ai/en/wonk/canadas-ai-compute-gap>.
- DSI. 2025. “Data Science Institute.” <https://dsi.ubc.ca/>.
- ICML. 2025. “The Forty-Second International Conference on Machine Learning.” <https://icml.cc/Conferences/2025>.
- Research Infosource. 2024. “Spotlight on University Artificial Intelligence Publication Performance 2018–2022.” <https://researchinfosource.com/cil/2024/focus-on-artificial-intelligence-research/spotlight>.
- Lockhart, Angus. 2023. *Automation Nation? AI Adoption in Canadian Businesses*. Toronto: The Dais, Toronto Metropolitan University. <https://dais.ca/reports/automation-nation-ai-adoption-in-canadian-businesses/>.
- NeurIPS. 2024. “The Thirty-Eighth Annual Conference on Neural Information Processing Systems.” <https://neurips.cc/Conferences/2024>.
- Thirion, Gilles. 2025. “Raising Business Sector Productivity.” In *OECD Economic Surveys: Canada 2025*. Paris: OECD Publishing. https://www.oecd.org/en/publications/2025/05/oecd-economic-surveys-canada-2025_ee18a269/full-report/raising-business-sector-productivity_443bcd88.html.