
Microsoft
aiPAB
AI Customer Insights

About aiPAB

The Microsoft Canada Artificial Intelligence Partner Advisory Board (aiPAB) is an industry group with the objective of accelerating the adoption of AI and other emerging technologies within Canada. The Board is comprised of leaders from Microsoft Canada, Certified Microsoft Partners, and Advisors from academia, industry, and non-profit organizations.

Board members are all senior-level experts in data and AI who meet quarterly to discuss pressing issues in the Canadian technology landscape. In addition to regular Board meetings, the Board frequently forms working subcommittees to advance initiatives identified during discussions and produce public-facing assets such as this report.

Acknowledgements

This research was conducted by the Human Futures Studio, a Toronto-based strategic insight and foresight consultancy, under the guidance of Microsoft Canada, the aiPAB, and a working subcommittee of the board.

Executive summary

There is a contradiction surrounding artificial intelligence (AI) in Canada: we are world leaders in AI research but laggards in industry adoption. This report details the findings of a qualitative study attempting to understand this contradiction at a deeper and more productive level.

We spoke to 18 senior leaders in data and AI from industries and geographies across Canada, asking them about their AI successes and failures, team structures, valued opportunities, and barriers to adoption. What we found was a set of insights that painted a maturing Canadian landscape with fragmented challenges to overcome, most of which stemmed not from the technical issues, but from an organization's ability to effectively integrate AI teams and projects.

Key barriers identified to successful AI projects included poor collaboration with non-technical teams, knowledge gaps leading to mistrust, inflated expectations of value, outdated forms of technology management, non-standardized team structures, absent regulation, and even speaking the same language.

Through a holistic view of the interviews and insights of this project, we propose seven key takeaways for organizations to consider in order to increase the likelihood of success for making your organization an AI-driven one.

1. Define and use specific language when talking about AI
2. Collaborate with different parts of the business to identify valuable problems
3. Train your colleagues on AI basics to increase trust and comfort levels
4. Manage and lead AI projects differently from traditional IT projects
5. Prioritize change management as the most critical step in any AI project
6. Build your AI practice slowly, simply, and intentionally
7. Be transparent with your plans to both your organization and the public

Table of contents

Project context

During a regular board meeting, members of the aiPAB recognized a gap in the Canadian knowledge base regarding AI and is application across the country. While numerous quantitative studies were available that explored topics like organizational maturity levels, high-potential use cases, key barriers to adoption, and perceived opportunities, the findings of these studies, while valuable, were limited in their depth of understanding. The board sought to understand Canadian AI adoption from a deeper perspective so that we might better serve organizations' needs and proliferate this transformative technology across the Canadian landscape. This set the stage for a qualitative research study that would ultimately dig deeper and understand the *why* behind the aforementioned topics.

What we did

For this project, we conducted 18 one-hour interviews with Canadian AI business decision makers from a wide diversity of industries, geographies, and sizes. Within these interviews, we used an open-ended research approach to explore topics such as the individual's and organization's experiences with AI, value derived from AI, high-potential uses cases and opportunities, barriers to AI adoption, risks associated, team makeups and structures, knowledge sources, and working with vendors. These interviews established an evolving narrative for deriving insight, whereby observations of earlier interviews would be further tested and explored in later interviews. The synthesis of these interviews produced the insights contained in this document.

How we did it

At HFS we ground our work in anthropological methods, which include an semi-structured interview format that gives the participant the most space possible to craft their answers as they choose. The purpose is to let them put their responses in their own words—because the way they structure their answers, the words they select, and what they choose to discuss and omit are as important as what they say.

Usually, not every topic needs to be directly raised. In interviews of this type participants often raise and discuss most of the topics of interest with only light guidance, given the context under which the interview is occurring. Throughout the interview, the researcher will guide the conversation so that if the respondent does not offer an answer that directly addresses one of the topics or sub-areas, the conversation will shift to address them more directly. Each topic will be addressed at least once, although usually the researcher will be able to guide the discussion so that they are covered from more than one direction.

The purpose of this style of interview is to avoid questions that elicit 'yes' or 'no' answers and then to allow the participant time to build an argument. Instead, the researcher will ask open-ended questions like:

"Tell me about your experiences with data and machine learning."

"What value are you looking to derive from new technologies?"

Who we spoke to

We spoke with 18 AI leaders from organizations across Canada. Participant names and organizations have not been shared for this research to respect the privacy of the interviewees and allow them to speak freely about challenges & opportunities within their work. Participants were senior leaders who oversee data and AI projects or practices within their organization. Industries represented include financial services (5), healthcare (4), telecommunications (2), retail (2), food & beverage (2), transportation (1), education (1), and real estate (1). Respondents came from a mix of private (9) and public (9) organizations. Organizations involved included small businesses (4), mid-market (4) and enterprise (10). Several geographies across the country were represented including British Columbia, Alberta, Ontario, and Quebec. Typical titles included CIO, CTO, VP, Director, and Manager, all overseeing some combination of data, AI, analytics, and digital strategy. Some came from very technical backgrounds while some were business-educated managers. Some had traditional IT/technology backgrounds while others were trained directly on machine learning and data science. Some oversaw large development teams while others largely outsourced their data and AI work to vendors. What they all shared was both the managing of various data & AI projects and the need to build some sort of practice around this work.

Artificial intelligence (AI) is not a settled idea. Despite having been a dream for over a century and a technical challenge for more than 60 years, AI is still not fully defined. It is no wonder that it is hard for people to conceptualize. And in this small point lies the crux of the problems facing AI in Canada. Something that is hard to conceptualize is hard to define, describe, implement, and use.

While the idea of AI is a familiar one, having been part of books, films, and stories long through human history, the concept of a practical, applicable AI is very new. Novelty is one of its key problems, because people do not yet have a solid understanding of how, when, and why to use AI.

AI also currently lacks a solid social role. While this role will develop over time, its incoherence is felt as an immaturity. We are all still in the process of experimenting with AI's application. Once we have a solid idea of how to use AI, at least at first, we will then get on with experimenting with its purpose. With this established, we will then know what AI is for and it will be easier to implement AI applications in the future. The insights in this report point back to the lack of this social role and the confusion it brings.

To push for the proliferation, application, and adoption of AI tools and technologies in Canada, we will have to work to provide AI with cohesive meaning and value by grounding it in an unambiguous place in our working and daily lives. Doing this will not only enable positive change, it will also go a long way to overcome the barriers to adoption that arise from the confusion over what AI is for and why anyone should care.

All of this is necessary to pull applied AI out of fiction and the curiosity cabinet where it has been placed. Once done, AI will fill many roles and jobs with the weight of the logical inevitably behind it.

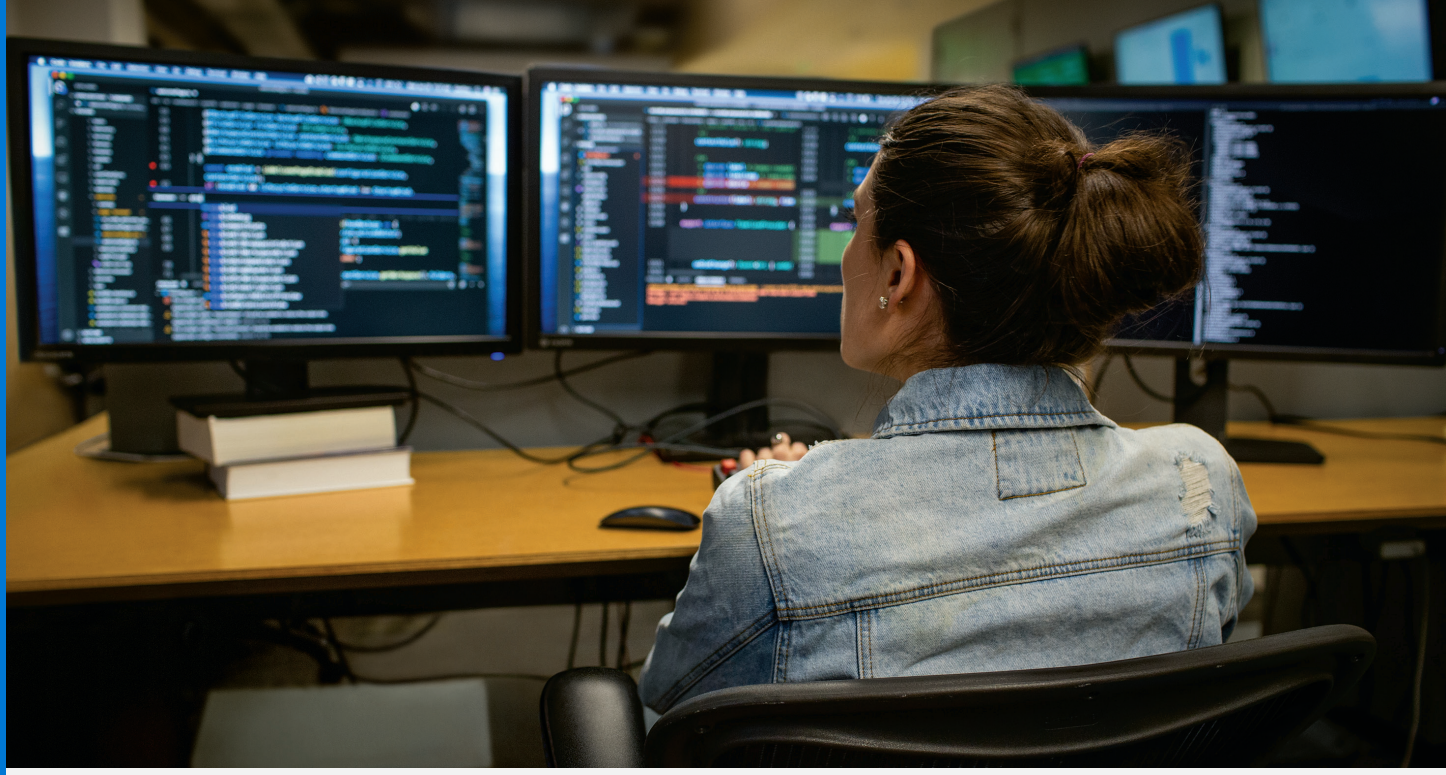

aiPAB

Understanding AI's value through its meaning in our lives

Value is best understood not as a property of a commodity, or as a property of exchange in motion; these are very limited notions of value. The value of AI begins with it being meaningful. It must have a form, structure, purpose, and place in people's estimation. AI is best understood in action as a tool for solving problems, rather than as a monolithic technical object.

For these things to happen, AI must first have a settled definition. This first set of insights highlights the problem of AI's meaning gap, and how this subsequently implicates its lack of clear value. This makes it difficult to integrate AI into any organization and, because it lacks obvious parallels to past technologies and ways of working, even to gauge the value of AI against established norms.

1.1 | Artificial intelligence means everything and nothing

The hype around AI has drowned meaningful discussions in a sea of hollow buzz words.

Though each interview was primed with an email discussion describing the intent to explore questions around the AI space, nearly every participant in this project began the conversation by clarifying what we meant by "AI", especially the more technical individuals. Though discussions were intentionally kept general, participants naturally gravitated towards the types and definitions of AI that they were most comfortable with, leaving our interviews as varied as the individuals we met with. This highlighted the diversity of meaning of AI to different audiences and one of the largest challenges for those working in the space: speaking the same language.

Implications

Before any meaningful discussion about AI can begin, the basic definitions of what AI is and where it will be applied will be critical. Any meaningful conversation about AI with a technical individual requires greater nuance in terms of developmental approach or use case application. Specificity in these instances, will be crucial to building a common understanding, even if it requires patience and shared learning to find this common ground. Experienced organizations barely talk about AI at all; they instead focus on the challenge/ opportunity at hand and the value that their project hopes to achieve. AI is simply the tool to do part of the job.

> "When you ask about AI, what **exactly are we even talking about?**"
>
> **-Private Enterprise Retail**

1.2 | The value and risk of buzzwords

Why use 'artificial intelligence' at all? It excites and it sells.

Though the previous insight highlights a risk around the term AI, there still lay value in its use. However, that value, like that of all buzzwords, lay in its use with less knowledgeable individuals. As a term for lay individuals, AI sparks the imagination of a world full of possibilities. The common meaning comes mostly from Hollywood and futurism, which means the term AI contains just enough specificity to give meaning to a conversation and just enough diversity to let that conversation go anywhere. The value in this term is to excite the possibility of a magic, catch-all technology that can solve all of an individual's greatest challenges, however, the risk of the term is that conversations never become tangible enough to make any real progress.

Implications

While the term AI can have tremendous value in the opening gambit of any exploratory conversation, it is important to very quickly guide an individual towards more specific terminology, applications, and outcomes to avoid the risk of spending too much time speaking at a uselessly high level. For less experienced individuals, this guidance often must come in the form of education, done not from an authoritative position, but instead established collaboratively, based on the knowledge and experiences of the individual you're speaking to.

> "I mean, it certainly **gets everyone excited...** and that can be a good thing and a bad thing."
>
> **-Private Enterprise Financial**

1.3 | AI loves problem-rich environments

AI models are tools, but like all tools, their value lay only in problem-specific applications.

The organizations who seemed to have the greatest success with their AI projects were those who started their funnel with numerous potential problems across a variety of business functions where AI could be applied. This is because, like all good innovation portfolios, AI efforts should be a funnel of projects. When striving for AI value, efforts should first start with idea generation around the art of the possible. These ideas should then be filtered through questions of data availability. Those with suitable data sources should then be considered for their potential value to the organization. This inevitable filtering of new AI opportunities will ultimately reduce to a handful of plausible projects, meaning that the top of the funnel requires a lot of potential problems or opportunities.

Implications

Initial conversations around the application of AI should be far more design thinking than machine learning. Before a single bit of data is cleaned or a single line of code written, efforts should be undertaken to identify the feasibility, data availability, and potential value of any project. In this approach, far greater value and likelihood of success is possible when building POCs or MVPs.

> "We have the data, yes, but we also have **no end of challenges** for my team to focus on."
>
> **-Public Enterprise Food and Bev**

1.4 | Benchmarking value of breakthroughs

Gauging the value of AI projects is made difficult with no prior frame of comparison.

The value of AI projects was difficult to quantify for some organizations because in many cases, the types of projects they ran were wholly new and made drastic improvements to metrics that the organization had previously not even considered, let alone tracked. Though participants often talked about such breakthrough projects with great enthusiasm, some of their leaders struggled to share that enthusiasm because they lacked a benchmark from which to quantitatively compare the value of the AI project.

Implications

Before a project begins, it's targeted value should be very clearly articulated beneath an organization's strategy. Even if the potential value of a project is different from an organization's traditional operations, there should be an agreement as to how the project's success will be gauged. As important as knowing the metric by which success will be measured, a benchmark and target improvement should also be established in order to quantify the value of AI and the ROI for the organization, even if it is speculative.

> "It's hard to have metrics to compare things to when **no one's ever done anything like this before.**"
>
> **-Public Enterprise Healthcare**

1.5 | Strategic limitations of platform tools

Many customers are not simply using cloud platforms, they are planning around them.

Beyond the obvious implications of the technological limitations of various cloud platforms, organizations that are deeply embedded in one ecosystem (Azure, AWS, GCP, etc.) are also having long-term, strategic discussions about how their needs align with the roadmap of these platforms. This means that their knowledge of a platform's capability could directly inform their digital strategy, project selection, hiring, and vendor selection. Lacking awareness of today's cloud capabilities could inform outdated project builds and tactics, however, awareness of tomorrow's capabilities could influence an organization's planning and strategy.

Implications

Your conversations should not simply focus on immediate needs and on current-state technology, but instead consider the Azure (and associated capability) roadmap to highlight alignment with future customer needs. For many, their AI practices and road maps are multi-year endeavours whose planning will factor in the capabilities and limitations of the tools at their disposal.

> "Azure's capabilities and future roadmap play an important role in determining **what we can do now and in the future.**"
>
> **-Public Small Real Estate**

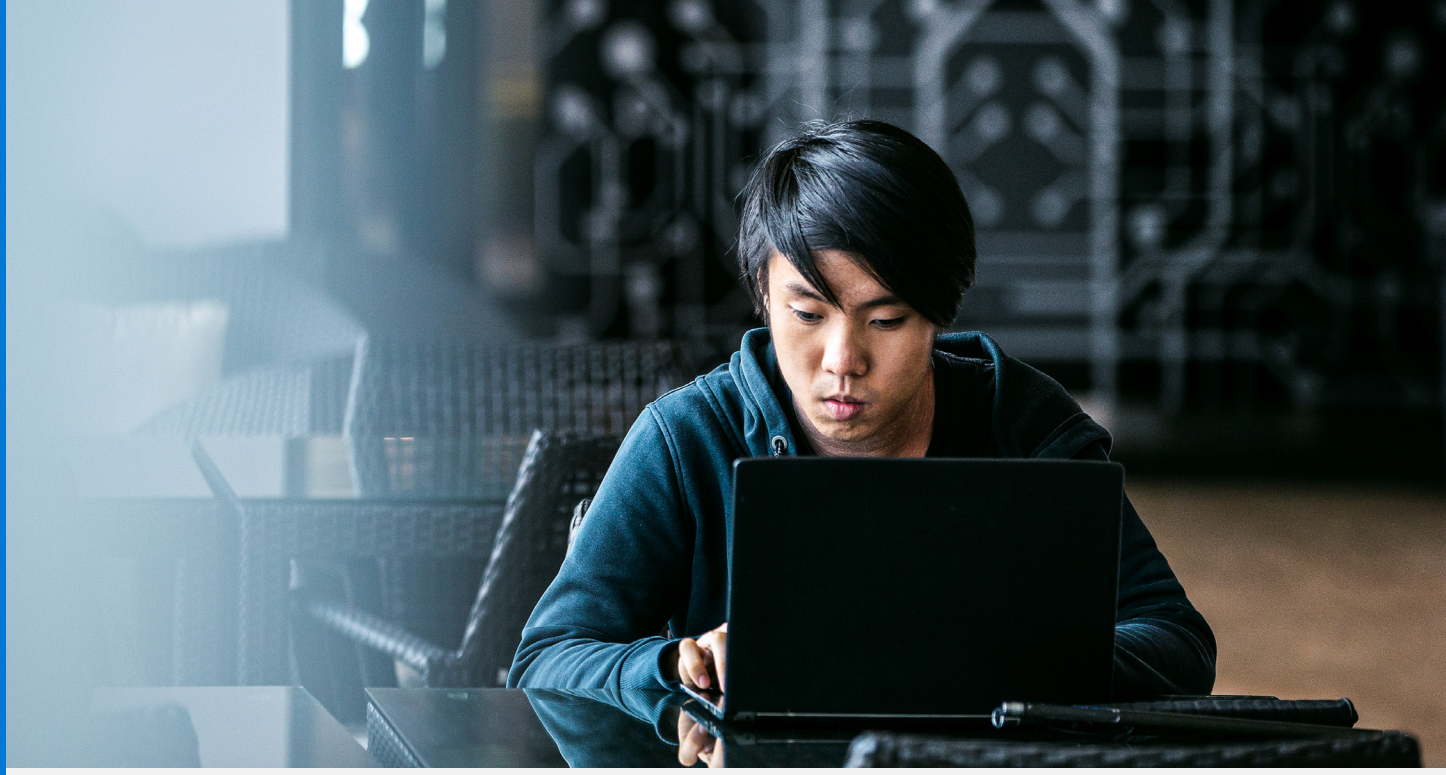

aiPAB

What is AI anyway, and whose tool is it?

A lack of knowledge and experience, even among experts, is presenting a considerable challenge for the manifestation of AI in Canada. The baggage brought by practitioners of older forms of technical problem solving and IT thinking constrain, and even contradict the transformative power of AI and its application.

Simply put, most people do not know enough about AI to apply it properly. This gap prevents the democratization of AI as a tool for all and limits its application to the narrow use cases, ideas, and processes of traditional technologists. However, those currently responsible for overseeing this application may be hindering its adoption and recognition of value. For AI to mature and expand in application, new people, perspectives, and problems are required to bring different ways of thinking for a different way of working.

2.1 | Greatest barrier to adoption is knowledge

Many individuals and teams lack basic knowledge of what to ask for or how to ask for it.

Surprisingly, when asked about why they weren't using AI more within their organization, few individuals mentioned costs. This was either because the costs paled in comparison to other challenges or because many of them didn't even know the costs associated. And one can't blame them. AI is a complex space with a wide range of different expert roles, governance structures, regulations, approaches, use cases, and other considerations. You will rarely find an individual who knows this space fully. Usually good teams are filled with a panel of diverse experts in machine learning, business analysis, data governance, ethics, and more, with a leader who is just dangerous enough in these topics to lead them. More commonly, you'll find a small, rag-tag but enthusiastic team with serious knowledge gaps.

Implications

If you sell AI services, your first job should be that of educator. No matter how much someone knows in the data & AI space, they can always learn more. The challenge is to be able to identify where an organization or individual is at in their AI maturity and how best to share new information with them to help guide them up the scale.

> "We're at the **'Mad-Men' stage of AI** right now; things are being bought that people don't even understand."
>
> **-Private Mid-Market Financial**

2.2 | Nobody trusts a black box

Skepticism surrounds AI when its results can't be explained or directly correlated to known data.

Alongside the knowledge gap, a lack of trust due to this knowledge gap can be a huge barrier to employees adopting a new AI tool once it is deployed. If a tool produces obvious results that align with employees' beliefs, they view the tool as being a waste of time and money. If a tool produces results that challenge those beliefs, they refuse to believe in the accuracy of that tool, particularly because they cannot understand how results were obtained. This leaves a very narrow space in which AI can affect change, however, it must do so in a way that respects the beliefs, rationale, and comfort levels of the employees who use it. As AI literacy increases, this effect begins to widen, however, this is a process that takes time and incrementally challenging, positive experiences.

Implications

In a perfect world, all AI would be explainable AI. The value of comprehension of an outcome cannot be overstated, however, that comprehension is a product of both the explainability of the models used as well as the knowledge of the audience. This means that your efforts must be twofold: create models that can show enough correlation to known data to appease its users' curiosity and take the time to educate would-be users on why AI is different from traditional sources of business intelligence.

> "No one is going to use one of these tools instead of their old manual process unless **they get how it works.**"
>
> **-Public Small Financial**

2.3 | Not everyone is allowed to play

Risks associated with centralized data structures necessitate a barrier to entry for who gets to work with them.

Particularly for more mature organizations that have centralized their data, processes, and tools, users of these systems have the potential to, at best, waste valuable time playing with models that have no business value, and at worst, cause irreparable damage to datasets that can set the organization back. While some of these challenges can be prevented through appropriate permissioning, ultimately the solution to this problem is that of good judgement. Technically minded people need to develop good business judgement in order to know how to use their time effectively. Business minded people need to develop good technical judgement to know how to respect the complexity of such systems and play within their own capabilities.

Implications

Some participants spoke to their ideal goal of having badges or certifications required before anyone can access certain datasets or tools. Ultimately, this is a question of education surrounding people who play with data and AI, however, with technologies, tools, and approaches constantly evolving, the education becomes less about specific skills and knowledge, and more about instilling a sense of good judgement and responsibility around use.

> "I feel like we need a sign around all of our data and tools that reads **'here be dragons'.**"
>
> **-Private Mid-Market Financial**

2.4 | IT can be a danger to AI

IT experts from the last era are only as useful as their ability to keep with the times.

CIOs, CTOs, and other IT professionals often inherent responsibility for data and AI projects. On paper, this makes sense; your technology experts should theoretically own responsibility for this new technology. However, in practice, the management of data and AI services is different enough from traditional IT to make these leaders not simply unsuitable, but at times an active liability. Traditional forms of technology management can unintentionally erode the value of AI projects. Even agile methodologies have the potential to be the kiss of death for AI since building a machine learning model often can't be broken down into bite-sized steps and failure can emerge from anywhere. However, some in the tech community push their teams forward with old ways, fail unexpectedly, and are confused when things don't improve.

Implications

If you are a CIO or other senior IT professional, recognize that your years of experience may be actively hurting your AI efforts. A new technology requires new knowledge, new approaches, and a new way of thinking. Upskilling may be required not simply for your dev efforts, but your leadership efforts as well.

If you are working with a CIO or other senior IT professional, be vigilant in watching for old school approaches to new technologies and politely (we are Canadian, after all) encourage management systems that give room for AI projects to succeed.

> "So many of these **old guys are pretending** to know this stuff, but are still operating like it's the 2000s."
>
> **-Private Enterprise Financial**

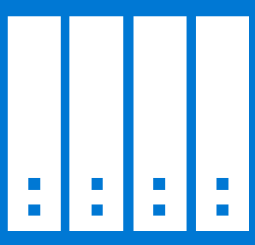

2.5 | AI risk is reputational risk

The risk of failure with AI projects is far greater than simply lost time or money.

When discussing risks, many participants cited public and often embarrassing failures of themselves or other companies and the ensuing brand/reputation damage. Given the relative immaturity of most organizations with respect to AI, failure is somewhat expected. However, individuals described two types of failure: the quiet, internal failures where time and money were lost and the loud, public failures where reputation was lost. Both of these failures are viewed as costly, however, only one of these is viewed as an acceptable risk.

Implications

When working on AI projects, go ahead and fail. Fail early, fail often, fail big, fail small, but do not fail loudly. Given the unproven nature of AI, failure is to be expected, however, by the time a project reaches the public or full deployment, it should have gone through enough testing and beta use to work reliably… or at minimum, be designed to fail gracefully.

> "If we can't show how we got a result… that is a risk not just to our offering **but to our reputation.**"
>
> **-Public Small Real Estate**

aiPAB

AI is nothing more than a tool in the end

AI finds its value in application. No matter how ambitious the project or how incredible the outcome, the reality of such projects is that AI is only a small component of a much larger ecosystem of people, processes, and other technologies. Just like a hammer is a critical tool to build a house, AI may be critical to many of such projects, however, it cannot be the point of the project. This is the failing of many technologists: to fetishize their toys to the point of thinking that they are more important than the people, organization, or even objectives.

AI is a tool, in this is it not much different from a hammer. It is a powerful tool, however, still just a tool. AI cannot do anything a human could do, it simply does things faster, more accurately and more precise. And like all tools, its value is recognized through the hands and mind of the user.

3.1 | Project chicken & governance egg

Which came first? It's a trick question.

In a perfect world, data governance and structures would be well-established long before serious AI projects are underway. The reality is that few senior leaders will see the value in spending on structural/backbone efforts, instead looking for more tangible value out of quick-win POCs/MVPs. Moreover, while basic data structures can often be implemented using pre-existing standards, more complex governance is often specific to the project types and use cases being explored by the organization. Although good governance is a critical foundation to implementing any sizeable AI project, the two efforts will ironically be at odds from a budgeting and resourcing perspective.

Implications

Business leaders want low-hanging, shiny fruit. Technologists want platforms to build upon. Wise projects are designed with the unsexy but necessary platform builds within the guise of a shiny, tangible AI project. This requires visionary leadership and planning since governance and structural components will often have to be pieced together over time across multiple projects.

> "I have to **hide the foundational stuff** that I need into the tangible projects that they want."
>
> **-Public Small Financial**

3.2 | Ask, don't tell

Even the greatest AI project has no value unless it solves a business need that compels people to use it.

The massive hype around AI has caused many people surrounding AI efforts to be a strange combination of skeptical and overly enthusiastic. This has led to equal parts of confusion about what's possible, fear about how jobs will change or be lost, and unrealistic expectations about what AI projects will achieve. The makings of this problem are the fault of AI teams *telling* the rest of the organization about the incredible things that AI can do without spending the time to *ask* what people actually want to change about their work. This can lead to projects that, no matter how technically advanced and successful, may never recognize their value because they do not address a genuine business need from some part of the organization.

Implications

Beyond basic AI education to raise the data literacy of an organization, the most important thing that we can do to proliferate AI is to ask questions. AI is simply a tool (albeit powerful) to solve business problems, however, this tool is useless if applied to areas with no value. Less mature AI teams should spend their time collaborating with different parts of the organization to identify how this tool can solve some of the greatest pain points, as opposed to screaming from the rooftops about how incredible AI is. More mature teams don't even need to ask; they let the value of their work speak for itself and put the onus of identifying business value on the teams that bring them new opportunities.

> "We need to stop telling people that we're going to make their jobs better or obsolete and **start asking them what they want.**"
>
> **-Public Mid-Market Transportation**

3.3 | Nobody likes change. AI is a huge change.

The more transformational the AI project the greater the change management needed.

The nature of a disruptive technology such as AI is that it will disrupt not just an industry, but also an organization. In order to recognize AI's value, this is necessary. However, this disruption is at odds with the realities of most organizations and employees; we don't like change. Several participants identified that while they may have numerous AI projects going on at any given time, they typically could only have one project being rolled out to an entire organization at a time. This is because the true bottleneck at this point is not the technology, but instead an organization's willingness and capacity for change. Not respecting this capacity for change is a surefire path to failure, even for the most successful AI models. This is, of course, assuming your workforce will embrace change at all… even despite your best efforts.

Implications

First, identify that an organization has a limited appetite for transformational change and prioritize your projects and efforts accordingly. Second, realize that the technological component of these transformations will always take a back seat to the human considerations and change management efforts in any successful project. Third, even done properly, anticipate resistance to change and expect the unexpected.

> "The organization can realistically only handle **one big project at a time.**"
>
> **-Public Mid-Market Transportation**

3.4 | Start with easy

Knowledge gaps, skepticism, and fear mean that your best starting point for AI is a simple and incremental project.

While the promise of AI is to disrupt organizations and industries alike, these kinds of innovative projects can frankly be exhausting. Particularly when a less mature organization lacks the knowledge and/or appetite for AI-driven change, the best course of action to get people interested in building a real AI practice is to start with a simple, straightforward, easy, and even dull project. Though your own data & AI teams may moan at the monotony of these projects, they are critical to building trust and highlighting the value of investing in AI at an organizational level. Participants who tried to run before they could walk acknowledged that they may have set their organization's AI efforts back months or years through overly ambitious initial projects that failed to live up to expectations or failed completely because the org wasn't ready.

Implications

Numerous participants acknowledged that their AI journeys started with off-the-shelf (OTS) AI add-ons from established service providers for things like cybersecurity and IT monitoring. OTS was described by one participant "gateway drug" to AI… but a necessary one that gave senior leadership the proof to invest further. These sorts of simple, straightforward, quick-win projects are a key step in the journey of evolving an organization's comfort and mindset around AI.

> "Everyone is interested in innovation, **but they're not interested in putting in the work.**"
>
> **-Public Mid-Market Financial**

3.5 | Expect failure and you'll never be disappointed

Vigilant monitoring, graceful failure, and other safety nets for the age of AI.

Outside of the most common AI use cases, most participants acknowledged that AI is still an emerging technology that is prone to many types of failure. Without even considering secondary risks like privacy, bias, and user experience, technical failures and unexpected results were still common in many described projects. Typically, insufficient or poorly labelled data was the source of these problems, however, some also described model drift and basic code bugs causing errors, even post-integration. These may not be the infant years of AI, but they're probably the awkward teen years, with some maturing left to do and failure modes always only one minor mishap away.

Implications

The most savvy organizations didn't just expect failure, they planned for it. They recognized that the black box nature of some AI projects necessitated a safety net and created vigilant monitoring protocols that in some cases generated as much or more data than the actual use cases themselves. Though extra human effort is required, the ongoing monitoring of AI applications serves the purpose of both catching failures before they hit outside users, as well as monitoring models for improvements and optimization.

> "Anytime a user reports a problem before you know about it... **you've failed.**"
>
> **-Public Enterprise Healthcare**

aiPAB

The next step for AI is building a culture around its practice

As is always the case with technology, once it is placed in the hands of people, the magic can begin. People will begin to develop work cultures, practices, and methods for making the most of the tools they have available. Currently, AI tools are largely limited to those 'in the know': data scientists, ML engineers, and developers. This is unfortunate because the true value of the technology will only be realized once it is in the hands of the many, not the few.

What is needed is greater dialogue between those in the know and those on the outside. This will serve two critical purposes. First, the outsiders will develop a better appreciation for the true value that AI can bring to their work. Second, those in the know will gain a better understanding for the needs of the organization around them and how AI can help. Only through this dialogue can an organization's AI culture mature to the point of delivering its true value.

4.1 | AI practices start as projects

The reality of building a piecewise team occurs around tangible initiatives.

Very few organizations have the resources, patience, and risk tolerance to build a complete and well structured AI team. More often, teams are slowly kluged together around specific exploratory projects and prototypes. However, there is no standard team structure or growth path because there is seldom a standard AI project roadmap. Some teams start with a data scientist to look through and clean existing data for opportunities. Some teams start with ML engineers to attempt to build POC models. Some teams start with non-technical people to oversee the building of an MVP by an external vendor. No two paths to practice maturity are alike.

Implications

When thinking about the needs of any organization and AI team, it is a must to get to know the individuals, their capabilities, and their gaps. In addition to understanding an organization's strategic needs for AI, it is crucial to also understand a team's need for specific AI skills, as well as how to design a project that can integrate into the organization and be successfully managed by the team.

If working with a less mature organization on an early AI project, realize that the people, processes, and tools established will not only implicate that project, but the future of the organization's AI practice.

> "I wasn't **given a budget to build a team…** all of this was built around projects."
>
> **-Private Mid-Market Healthcare**

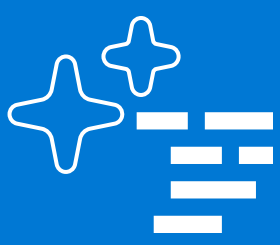

4.2 | No two AI teams are the same

Diversity of use cases, approaches, and organizations implicates a lack of any standardized team models.

Because AI is an emerging technology for many organizations and because an organization's team is often purpose-built over time and the needs of a project, no two AI teams will look the same. Some participants oversaw a team of pure data scientists while others managed a mix of developers, analysts, ethicists, and designers. Given the relative immaturity of AI in organizations, this even means that different roles will lack common training and standardization; for example, data scientists described by our participants often came from a variety of educational and career backgrounds. This isn't a bad thing, however, it does mean that cookie-cutter approaches to managing or working with teams will inevitably fail. Hardline waterfall could be as bad as hardline agile or kanban… it depends on the team.

Implications

Leading AI teams requires a dynamic leadership style and approach. And with every new hire to a team, there is no guarantee that the same approach will work for that individual or the evolving team structure.

The same is true for vendors; what worked with one client may not work for another. Instead, you should look to the individuals who you are working with, their capabilities and needs, and how to best augment or support teams to achieve success.

> "Initially, **our team had no data or AI people…** for a while we just managed vendors."
>
> **-Private Enterprise Retail**

4.3 | Criticality of the AI translator

The most important disconnect in AI occurs between leadership's strategy and developers' implementation.

Though change management may bottleneck the backend of AI projects, the other major bottleneck and a large source of failure for AI projects is the lack of a sort of AI translator role: an interpreter of strategic vision, tangiblizer of AI value, and architect of AI models and approaches. This person could be a strategically-minded developer who evolves into a higher calling or a business leader with a penchant and curiosity for learning about tech. Their origin is far less important than their function. The translator acts as a go-between for the technical and strategic, constantly swapping hats to ensure that both sides are well considered and, as change inevitably happens over the course of a project, that people on both sides of the wall are aligned on the intent, objectives, and limitations of the project.

Implications

If you can't identify this person in an organization, there is a chance it could be you, or no one. Many people informally and unwittingly play this role, however, there is value in making the role explicit. By doing so, both sides have someone to hold accountable and be accountable to. This role is most important upfront, when filtering a long list of potential projects down by jointly weighing the feasibility and viability of a project. However, the role retains value over time as a project evolves to ensure alignment with all stakeholders.

> "Half of my time used to be spent as the **interpreter between leadership and my developers.**"
>
> **-Private Enterprise Retail**

4.4 | It takes a village to integrate AI

Rolling off key team members too soon can drop valuable POVs in a critical moment.

Though early AI projects can often exist in siloes with a small group of technical individuals during the development phase, integration of the project back into the organization or onto a public-facing service requires an extensive team: data scientists, developers, system architects, cloud experts, UX designers, business analysts, and more. Though not all of these roles will be active at all times during integration/deployment, each of them provides a critical lens through which they view the project and will ask important questions to help mitigate risks and ensure a successful rollout. Remove even one role too soon and you risk exposing yourself to be blindsided by the lack of expertise required by this complex ecosystem of people and technology.

Implications

The implementation of large and complex AI projects requires a sort of roundtable of diverse experts. Though some will be more active participants in the deployment of these projects, others may simply need to be informed of implementation decisions so that they can bring their expertise to review plans and periodically ask a single question that may prevent catastrophe down the line.

> "No, we don't roll people off projects that go live anymore, **they just switch to being advisors.**"
>
> **-Private Enterprise Telecom**

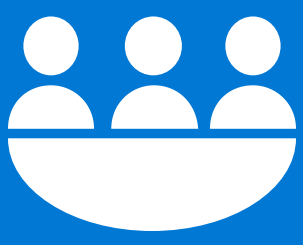

4.5 | Vendor management is an AI job

To get true value out of a vendor, manage them with someone who can speak the language.

Several participants described working with vendors where the day-to-day client-side manager was business educated. Though these projects seemed to run well in the moment, in reality, they were on borrowed time. The inexperience of the client manager meant that both critical information from the vendor was not landing correctly and that client-side organizational changes that seemed irrelevant (but weren't) didn't get communicated back to the vendor. Ultimately, the right decisions were not happening around the project at the right time and projects would slowly spiral out of control until a critical failure would unearth a backlog of misalignments and unresolved issues.

Implications

People managing vendors need to be able to speak the same language. Though vendors can help to educate and guide these individuals, the information asymmetry of both sides (client about the technology, vendor about the organization) will mean that gaps are bound to occur. A healthy project remains so through constant questions and updates using the same language. Any missed information has the potential to become the seed for future disaster despite the best intentions of both sides.

> "The technology wasn't mature, the project was complex, and **a PM wasn't enough to manage them.**"
>
> **-Private Mid-Market Food & Bev**

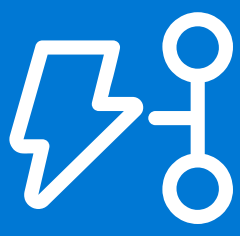

aiPAB

The modifications of AI use and practice are yet to come

Regulation will create the necessary brakes on the system. But because they are developed as a reaction, they have not yet reached a point of maturity to equal the expansion of AI application. Canadian organizations know that more policy is coming in the AI space, and many of them identify this as a source of risk. What fewer organizations realize is that they will have a hand in crafting this policy, either passively or actively.

Each time a new model, service, or product is deployed, it carries the passive potential to trigger new regulation when some implicit, unspoken rule is broken around engaging with the Canadian populous. However, this reactive regulation will almost necessarily be an overreaction, slowing the adoption of AI for everyone in Canada due to the overzealous actions of a few. Wise organizations, on the other hand, understand the knowledge and resource limitations of government and are actively trying to work with them to guide responsible yet fair policy into place so that Canada can leverage the true value of AI.

5.1 | The biggest problem with regulation is its absence

Everyone is operating under assumed regulation to come. . . but with no guarantees.

AI regulation can't be a problem yet because in most contexts, industries, and geographies, it doesn't exist. However, most participants from this research are well aware of that vacuum and identify it as a massive risk for the future of AI. Currently, Canadian organizations look to regulation from other markets, industry best practices, or best guesses when trying to speculate on how future regulation could impact them. This leaves many feeling insecure, investing heavily into projects built upon assumed standards and pending regulation that could never come to be.

Implications

One strategy is to work with regulators to fill this gap (see 5.2). Another is to simply be a good steward for society when deploying AI. Responsible and ethical AI discussions aren't always the most exciting, however, they may be critical to fueling much needed public trust around AI and for proving to government that your organization can be trusted with such powerful tools instead of becoming the target of reactionary policy like we've seen with facial recognition applications.

> "As it is, we're operating under **big assumptions,** just waiting for some new law to change everything."
>
> **-Private Enterprise Retail**

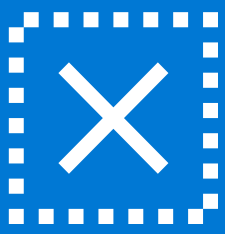

5.2 | Leading regulators in more ways than one

Technology outpaces regulation. More and more, technologists are looked to for help.

No one is surprised by government's inability to keep up with the pace of innovation in the technology space, particularly around complex and ethically charged topics such as AI. However, numerous participants - particularly those in crown corporations, public sector organisations, or those perceived by government as lacking bias - talked about ministers and other policymakers looking to them for guidance on policy development. Even private financial institutions spoke about working collaboratively with OSFI on establishing new standards and structures. Lawmakers in Canada identify that they lack expertise and are seemingly happy to work with organizations willing to support the process, so long as it is done in a transparent, open, and fair manner.

Implications

While American brash decision making may fuel tremendous technological innovations, Canada's business conservatism may come in handy here. Rushing too far ahead too fast can cause reactionary policy that can stifle innovation for years to come. Instead, wise organizations get close with their public counterparts and learn to approach regulation less as a set of rules to follow and more as a collaborative set of guidelines to be evolved to serve the joint best interests of the organization and the Canadian populous.

> "We're publishing a framework in the hope that government will **consider it when making legislation.**"
>
> **-Public Enterprise Healthcare**

5.3 | Our risky southern brother

American innovation comes at the cost of American controversy.

Particularly in discussions with organizations headquartered in the US, some participants acknowledged the cultural and operational differences between Canadian and American counterparts. While they spoke enviously and admirably of the rapid progress and huge leaps forward made in the US, they also spoke cautiously about the brazen and reckless advances of those down south.

Implications

If history holds, the US will always take more and larger business risks than we do north of the border, and AI should be anticipated to be no different. Instead of looking on with jealousy, we should be monitoring meticulously and taking notes. Given our cultural, economic, and regulatory similarities, much of what succeeds and what fails in the US should also hold true in Canada. Consider American innovation to be a free testbed for us to watch, learn from, adapt to, and improve upon.

> "It's not just the money or lax regulation that sets them apart, **there's an entire culture of 'let's see' down there.**"
>
> **-Private Enterprise Telecom**

aiPAB

A *human futures* vision of AI in Canada

The success and wider application of AI in Canada is not predominantly a technical challenge. It is not even really a business challenge. It is a human challenge. Because AI is still an immature technology, it is not surprising that people are not ready to see its immediate utility.

However, once the remedial work outlined in this document is underway, the application of AI will become easier. Once a clear social role and meaning has begun to form, the business cases for AI application will proliferate and grow. Once the value of AI has been demonstrated, a period of experimentation will follow where novel, and even more useful, applications will arise.

While starting with the human side of AI might seem counterintuitive to those in the technology sector, it is actually the only way that AI can fufill its promise. We must overcome the human barriers to adoption and apply AI within the logic provided by its meaning, value and structure. With these elements in place, the rest settles into the realm of incremental, thoughtful innovation.

The future of AI in Canada is as much a human story as it is a technical one.

The insights shared in this document provide a valuable and critical look at the Canadian AI landscape. Each alone could provide your organization with an impetus for change in some aspect of your business or operations. However, considered together, these insights may leave the reader with a challenging question: *what should I do*?

To answer this question, we synthesized our research findings into a list of seven key takeaways. While these recommendations will not guarantee the success of your AI projects and practices, they are a critical foundation to recognizing the value of AI in your organization.

1. **Define and use specific language when talking about AI**
   AI has come to mean everything and nothing. To ensure productive conversations around AI planning and projects, define key terms upfront and use them in regular conversations. And when in doubt, always clarify what people are talking about.

2. **Collaborate with different parts of the business to identify valuable opportunities**
   AI's tremendous potential can only be realized when applied to real business problems. Collaborate with different teams in your organization and ask them about their day-to-day challenges to identify high-value opportunities.

3. **Train your colleagues on AI basics to increase trust and comfort levels**
   AI's black box reputation is a barrier to internal adoption. Try to create explainable models and raise the AI knowledge levels of your colleagues to make them more willing to use AI because they understand how it obtains its results.

4. **Manage and lead AI projects differently from traditional IT projects**
   AI requires different leadership and management than traditional IT. Trust your expert talent, retrain your legacy talent, and be prepared to lead AI projects dynamically; manage work with the style and approach best suited to the project.

5. **Prioritize change management as the most critical step in any AI project**
   The best AI models fail if no one uses them. For every dollar invested into technical development, a dollar or more may be needed for change management. Start early and involve key stakeholders from opportunity identification to deployment.

6. **Build your AI practice slowly, simply, and intentionally**
   Running before you can walk is a good way to trip. Early AI projects establish the tone, team, and tactics for an AI practice. Start with easy projects to build confidence and highlight value. Be intentional about early projects that establish the practice.

7. **Be transparent with your plans to both your organization and the public**
   AI can be disruptive; disruption can be scary. Transparency allows stakeholders in and around your organization to make informed decisions regarding if and how they will use your AI. Hiding plans leaves you on borrowed time for a reckoning.

Want to learn more?

If you would like to know more about Microsoft's AI services, the Microsoft Canada AI Partner Advisory Board, or the insights shared within this report, you can contact the aiPAB through the email address below.

Email: **aiPABca@Microsoft.com**.

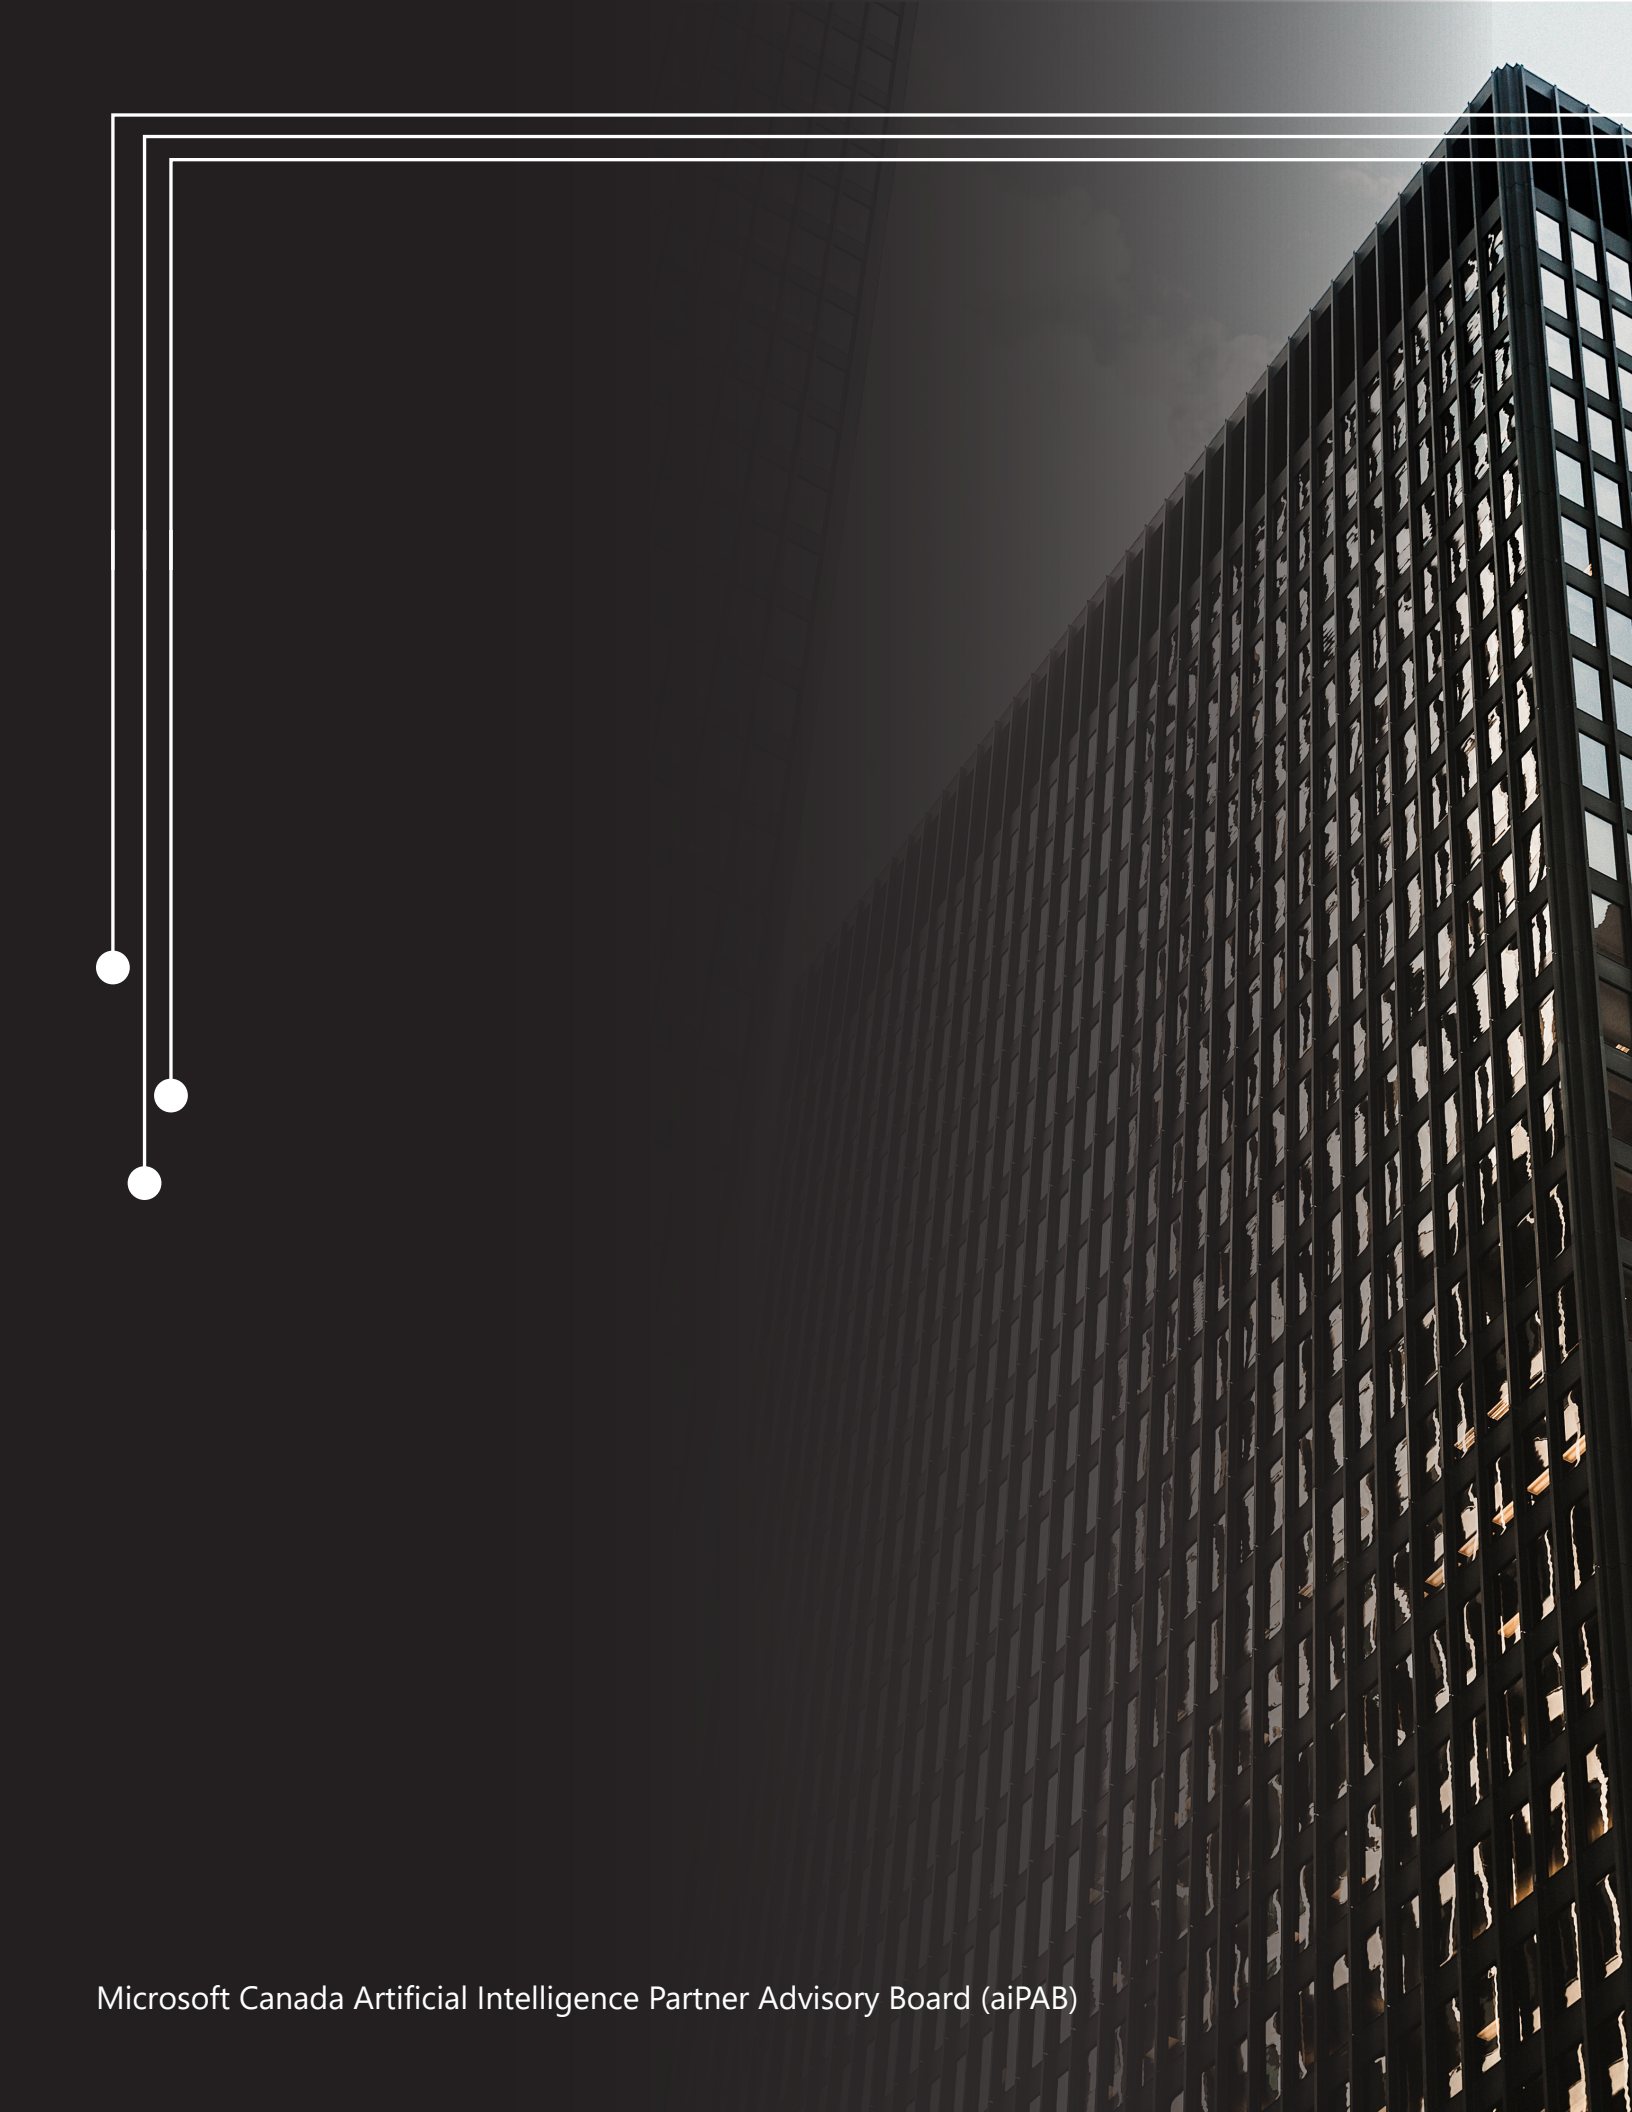

Microsoft Canada Artificial Intelligence Partner Advisory Board (aiPAB)