LMIC Insights

Conversations With Experts: Six Things to Think About When Using AI and Emerging Technologies in Labour Market Information

May 2026

The Labour Market Information Council (LMIC) is a pan-Canadian non-profit that produces accessible, evidence-based insights on Canada's labour market. Through research, collaboration, and data innovation, LMIC supports governments, employers, workers, and educators in making informed decisions. Our work helps bridge information gaps, improve labour market outcomes, and strengthen Canada's workforce development ecosystem.

To learn more about our research and initiatives, visit **lmic-cimt.ca**

Authors

Lorena Camargo, Principal Researcher, LMIC
Ken Chatoor, Director of Research and Strategic Foresight, LMIC

For more information about this report, please contact research@lmic-cimt.ca.

How to cite this report

Camargo, L., & Chatoor, K. (2026). *Conversations with experts: Six things to think about when using AI and emerging technologies in labour market information (LMIC Insights).* Labour Market Information Council (LMIC).

Table of Contents

Executive Summary

Artificial intelligence (AI) and emerging technologies are changing how labour market information (LMI) is produced, analyzed, and shared in Canada. These new tools draw on online job postings, administrative datasets, surveys, and other large-scale data sources to provide timely and detailed insights into labour demand, skills, and occupational change.

However, our conversations with Canadian vendors and developers of AI-driven LMI tools suggest that the value of these technologies depends not only on their novelty, but on how they are designed, governed, and used. Data quality, transparency, explainability, bias management, and human oversight consistently emerged as central to trust and reliability—the cornerstones of our work and our field, even with innovations on the horizon.

This LMIC Insights piece draws on eight semi-structured interviews with Canadian vendors active in the LMI space. The information we gleaned from them is complemented by a review of emerging academic and policy literature on AI governance and responsible AI.

We have distilled these perspectives into six practical considerations for organizations and policy-makers looking to adopt AI-driven tools as well as for developers seeking to build and maintain trustworthy, fit-for-purpose systems. These considerations sit against a backdrop where Canadian AI policies speak in general terms and there is still no unified governance framework for AI use in LMI.

These considerations are not exhaustive, but they highlight shared practices and concerns that can support more responsible use of AI in the LMI space. A central takeaway from our conversations is that the most important determinants of trustworthy AI in LMI are not technical features alone, but governance choices, such as how data are sourced, how models are validated, and how human expertise remains involved.

This LMIC Insights piece is part of a broader series on AI and emerging technologies in LMI that will also include a practical guide and an international perspectives piece.

Table 1: Six key considerations for AI adopters in the LMI space

Consideration	What it means in practice
Recognize that data sources matter	Ask where the data come from, how often they are updated, how they are validated, and what is missing. Look for vendors that are open about data gaps and limitations.
Keep humans in the loop	Favour tools that include ongoing human oversight in training, validation, and maintenance. Check whether domain experts or target communities review outputs.
Understand the range and risk of AI use cases	Distinguish between lower-risk uses (like taxonomy updates or exploratory analysis) and higher-risk uses (like job-matching or career-advising), where explainability and accountability matter more. Start with lower-risk applications to build capacity.
Build user awareness	Invest in helping staff and partners understand what a tool can and cannot do. Ask vendors for clear documentation, plain-language explanations, and support materials that build AI literacy.
Prioritize transparency and explainability in higher-stakes contexts	In contexts that affect funding decisions or people's careers, prioritize tools that can explain how outputs are generated. Look for vendors that disclose data sources, model logic, and key assumptions.
Acknowledge and mitigate bias	Recognize that bias cannot be fully eliminated, and ask vendors to be open about where their data may over- or under-represent certain industries or communities. Look for regular benchmarking and inclusive validation processes that identify and address these gaps.

Introduction

AI and emerging technologies are reshaping how LMI is collected, analyzed, and shared. From parsing millions of online job postings daily to supporting the identification of emerging skills and occupations faster and more efficiently than prior methods, these tools have the potential to significantly enhance the timeliness and depth of LMI in Canada and beyond (Vankevich & Kalinouskaya, 2021).

Our conversations with Canadian vendors show that the benefits of AI are not automatic. How these systems are developed, trained, validated, and implemented matters just as much as what they are technically capable of doing. Choices around data sources, model design, human involvement, and transparency shape whether people view AI-driven LMI tools as credible, useful, and trustworthy (Glikson and Woolley, 2020). In this field, trust is the gold standard. Without it, users will not feel confident in what they read or receive.

Drawing on these conversations, this piece identifies six considerations that vendors repeatedly raised when discussing the reliability and responsible use of AI in LMI. Together, they point to what organizations and policy-makers should look for when adopting AI-driven tools and to what developers and vendors should demonstrate to build and sustain trust in their systems.

LMI refers to any data or insights that help people understand current and future labour market conditions so they can make informed career, education, or hiring decisions.

LMI can include

Quantitative Data

such as employment rates, wages and salaries, business and industry trends, job vacancy data

Qualitative Insights

information from interviews, surveys, or lived experiences that shed light on how people interact with the labour market

Forecasts and Trends

projections about future skills needs, occupational demands, or sectoral shifts

Informal Sources

advice from mentors, peers, or online communities that help individuals make employment decisions

At its best, LMI is timely, relevant, accurate, and accessible—enabling individuals and organizations to act with confidence in a changing labour market.

Source: LMIC

Background & Context

The rise of AI in LMI

Across Canada's LMI ecosystem, AI and emerging technologies are changing how data are gathered, processed, and interpreted. These tools enable near real-time analysis of large volumes of information, including job postings, resumés, and other data sources relevant to the labour market, such as surveys (Adkins-Hackett and Trehan, 2024). LMI products can be produced more quickly with the use of AI tools—and, in some cases, with greater detail—than with traditional approaches (Tzimas et al., 2024).

At the same time, these innovations raise important questions. Concerns around data quality, explainability, and bias came up frequently in our conversations with vendors. While AI can highlight patterns at scale, its outputs are shaped by the data on which it is trained and the assumptions embedded in model design. Without careful validation and oversight, faster, more granular, or local insights do not necessarily translate into better or more reliable ones.

The policy and institutional landscape in Canada

Canada's AI space is influenced by a growing set of policy and research frameworks that emphasize responsible, transparent, and human-centred innovation. These include the *Directive on Automated Decision-Making, the Pan-Canadian AI Strategy*, and federal investments in responsible AI research through organizations such as Employment and Social Development Canada and the Social Sciences and Humanities Research Council (SSHRC, 2025).

With the 2025 election, Canada saw the *Artificial Intelligence and Data Act (AIDA)*, introduced in 2022, go defunct, but it also saw the appointment of Canada's first Minister of Artificial Intelligence and Digital Innovation, Evan Solomon (Castaldo & Norman, 2025). A 26-member task force was launched to actively engage Canadians and develop policies and guidelines in the Canadian AI space (CPA Ontario, 2025). This signals the importance that the Government of Canada places on AI as part of its efforts to establish global leadership in the area.

In fact, the Canadian government has been working for close to a decade to support the adoption of AI across the country through the *Pan-Canadian Artificial Intelligence Strategy* (Innovation, Science and Economic Development Canada [ISED], 2025). The Canadian Institute for Advanced Research (CIFAR) is leading phase 1 of the strategy, which was launched in 2017 and aims to strengthen Canada's talent base and competitiveness in AI research on the global stage (CIFAR, n.d.). CIFAR works closely with AI institutes across Canada—including Amii (Edmonton), Mila (Montréal), and the Vector Institute (Toronto and Waterloo)—to attract, retain, and develop top AI talent in Canada (ISED, 2025).

Canada continued to invest in AI infrastructure when it directed 2025 Budget funds to enable a large-scale sovereign AI data centre and develop the *Canadian Sovereign AI Compute Strategy*. This strategy seeks to ensure that people in Canada who want to work, research, and develop in this space have access to the compute capacity required to do so (ISED, 2025).

Vendors noted that these frameworks and actions shape expectations around transparency and governance, even when tools are developed and deployed in the private sector and for customers in the public and private sectors. At the same time, there is still work to be done on how these policies apply specifically to AI-driven LMI tools. Most speak in general terms, and specific regulations for AI use in the LMI space remain absent. Despite multiple initiatives, there is no single, unified governance framework for AI use in LMI in Canada.

The conversations we had with Canadian vendors and developers of AI-driven LMI tools raised raised broader questions about how policy frameworks translate into practice: how they influence what vendors actually do and what they mean for trust and governance in AI-driven LMI systems. While this piece does not resolve these issues, the questions nonetheless form an important backdrop for understanding the considerations that follow and the work that LMIC will continue to focus on in this AI and LMI series.

Methodology

This LMIC Insights piece draws on a detailed literature review, a scan of Canadian policies, and eight semi-structured interviews with Canadian vendors and developers of AI-driven LMI tools. We selected participants from across the Canadian marketplace, including organizations engaged through our ongoing relationships in the LMI landscape as well as vendors providing services and tools to private-sector clients and LMI disseminators (such as sector councils and government departments).

The conversations focused on how AI is integrated into data collection, analysis, and delivery processes. Key areas of discussion included:

- data sourcing, training, management, and quality assurance
- human involvement in model training and validation
- transparency, explainability, and bias mitigation when innovating in the AI space in LMI
- governance, trust, and user education

The purpose of these conversations was to identify shared practices, challenges, and trust-building signals that could inform both the adopters and developers of AI-driven LMI tools. This work is meant to serve as a knowledge base to help interested AI users and developers identify common practices, know what to look for, and understand how to inform themselves when working with these types of tools.

Our findings reflect insights from a limited sample of vendors that are active in the labour market space. The findings are not intended to be exhaustive, but to highlight emerging themes. This work represents a starting point for a broader—ideally, ongoing—conversation.

Findings

Six things to think about when using AI in LMI

Each of the six themes discussed in this piece is drawn from points that vendors emphasized as important when considering, developing, and adopting AI in the LMI space. For each theme, we present two perspectives:

- What adopters should look for when considering an AI-driven LMI tool
- What developers should demonstrate to support credibility and nurture continuous confidence among users

1. Recognize that data sources matter

Across conversations, vendors consistently emphasized that data quality is central to model reliability and overall product outcomes. As one interviewee put it plainly, "garbage in, garbage out." The sources used to train and develop AI systems shape the accuracy, relevance, and limitations of the resulting outputs.

Vendors highlighted the importance of using verified, reputable data sources—including official Canadian datasets, such as the Job Bank, Labour Force Survey (LFS), and Job Vacancy and Wage Survey (JVWS)—as well as reputable online job-posting data for training and model development. They described regular benchmarking and validation against official statistics as critical practices.

At the same time, there was broad acknowledgement that there is no such thing as perfect information. Transparency about data gaps and limitations was seen as more valuable than claims of completeness. For providers, this means being open when there is not enough information to speak reliably about a given market, and informing users when insights may overrepresent certain industries or may not be as accurate as others. For example, if data on rural communities are missing, that should be explicitly stated. Being forthcoming about what is and is not in your data is paramount to building trust so that users do not infer things that are just not possible from these tools.

We've built in human validation, even though our systems can detect emerging skills. We try to reduce how much time that takes, but it's essential—because it's 'garbage in, garbage out.'

For adopters:

Ask clear questions about where data come from, how often they are updated, and what validation processes are in place. Tools that integrate official Canadian datasets or reputable job-posting sources into the training and development of their AI systems provide a good initial signal. Importantly, adopters should expect and accept that all datasets have limitations and prioritize vendors that are open about them.

For developers:

Maintain clear data documentation and quality control processes. This includes being open about data-collection processes, sources, data coverage of assessed markets, potential limitations, and missing data. Regular benchmarking against official statistics and the inclusion of metadata (data that describe other data) (Badman & Kosinski, n.d.) can help users understand data coverage and the strengths and weaknesses of certain tools and outputs.

2. Keep humans in the loop

Human involvement emerged as a recurring theme across all stages of AI development and deployment. Vendors stressed that human oversight plays a key role in training, validating, and updating models over time. Many told us they do not envision a world where humans do not work hand in hand with these technologies.

Subject-matter experts, including those with industry and community-specific knowledge, are critical for contextualizing results and assessing whether outputs align with lived realities. This reflects the importance of combining large-scale data with local and domain-specific expertise. Human input remains essential not only for developing these technologies, but also for interpreting and strengthening their outputs. In some cases, vendors also emphasized the value of using simpler or more interpretable models as well as benchmarking AI outputs against expert human analyses.

For adopters:

Look for tools that are transparent about their processes and include human oversight at key points, such as during training, validation, and maintenance; this can be a strong signal of reliability. References to outputs being reviewed by domain experts or target communities can provide additional assurance about reliability and relevance.

For developers:

Build systematic human review into model training and quality assurance processes. Engaging subject-matter experts and affected communities can improve relevance and support interpretability. Several vendors emphasized the importance of not overcomplicating models when simpler approaches are sufficient.

> **There is a clear intersection between big data and local or domain-specific knowledge. We are not going to move away from that any time soon.**

3. Understand the range and risk of AI use cases

AI is currently being used in LMI for a wide range of tasks, from updating occupational and skills taxonomies to predicting broader economic trends. For example, large language models (LLMs) are being used to update pre-existing taxonomies, make labelling more efficient, and expedite the translation of large volumes of text.

While these models offer new possibilities, vendors pointed to risks associated with "black box" systems—that is, systems whose internal decision-making processes are not visible (Kosinski, 2024). When even developers cannot fully explain how outputs are generated, transparency becomes a critical challenge, particularly in contexts where results may influence career, education, or training decisions. In these cases, a lack of explainability can undermine trust.

Organizations may need to assess available tools based on the problem they are trying to solve, including whether to prioritize responsiveness or interpretability. This includes evaluating trade-offs between performance, transparency, and data sensitivity.

Low-risk applications—such as updating existing taxonomies or exploratory analysis—were commonly described as good entry points for experimenting with LLMs in LMI tools. Other low-risk uses include translation and labelling tasks, where humans can review and validate outputs.

Higher-risk applications include forecasting models used to project growth across occupations or regions. These uses are not inherently inappropriate, but they require greater attention to model design, validation, and the ability to clearly explain results.

> **The best approach is often to test multiple tools and determine what works best—while being mindful of data sensitivity. Even with enterprise tools, there are cases where data cannot be shared due to security or governance requirements.**

For adopters:

Differentiate between low-risk and high-risk uses of AI. Tools that support tasks like taxonomy maintenance or exploratory analytics carry different implications than those used for job-matching or career-advising (higher risk). Starting with lower-risk use cases—like internal exploratory work—can help organizations build internal capacity and learn about AI capabilities in a variety of contexts before moving into higher-stakes applications.

For developers:

Be explicit about what your tools are designed to do—and what they are not. Clear documentation of the limitations and appropriate contexts for use can support responsible adoption. In higher-impact domains, prioritize explainability. Consider piloting LLMs in controlled, low-risk environments before deploying them more broadly.

4. Build user awareness

Several vendors noted that the biggest barriers to effective AI use and broader adoption are not always technical. Instead, gaps in understanding about what AI can and cannot do—and uncertainty about what AI tools should and should not be used for—often shape how tools are perceived and used. As one interviewee mentioned, "I think the roadblock right now is really just people knowing it and using it (AI tools)."

Some vendors told us they see themselves not just as technology providers, but educators. They understand clear communication about limitations, assumptions, and appropriate use cases as part of building long-term trust with users and providing a necessary stepping stone toward addressing adoption challenges.

For adopters:

Invest in staff education and ask vendors for documentation and methodological transparency to help ensure tools are used appropriately. Do not hesitate to ask questions of companies that provide AI-driven tools and services: ask about the limitations of the models and tools, what is considered responsible use, and the intended applications. The answers can help ensure that everyone has a common understanding of how the tools should be used.

For developers:

Provide clear, accessible guidance on model logic and limitations. In our interviews, honest communication about caveats or shortcomings was consistently framed as a trust-building practice. Several vendors also highlighted the value of incorporating AI literacy materials into client onboarding and ongoing support so that users can understand both the strengths and limitations of the tools.

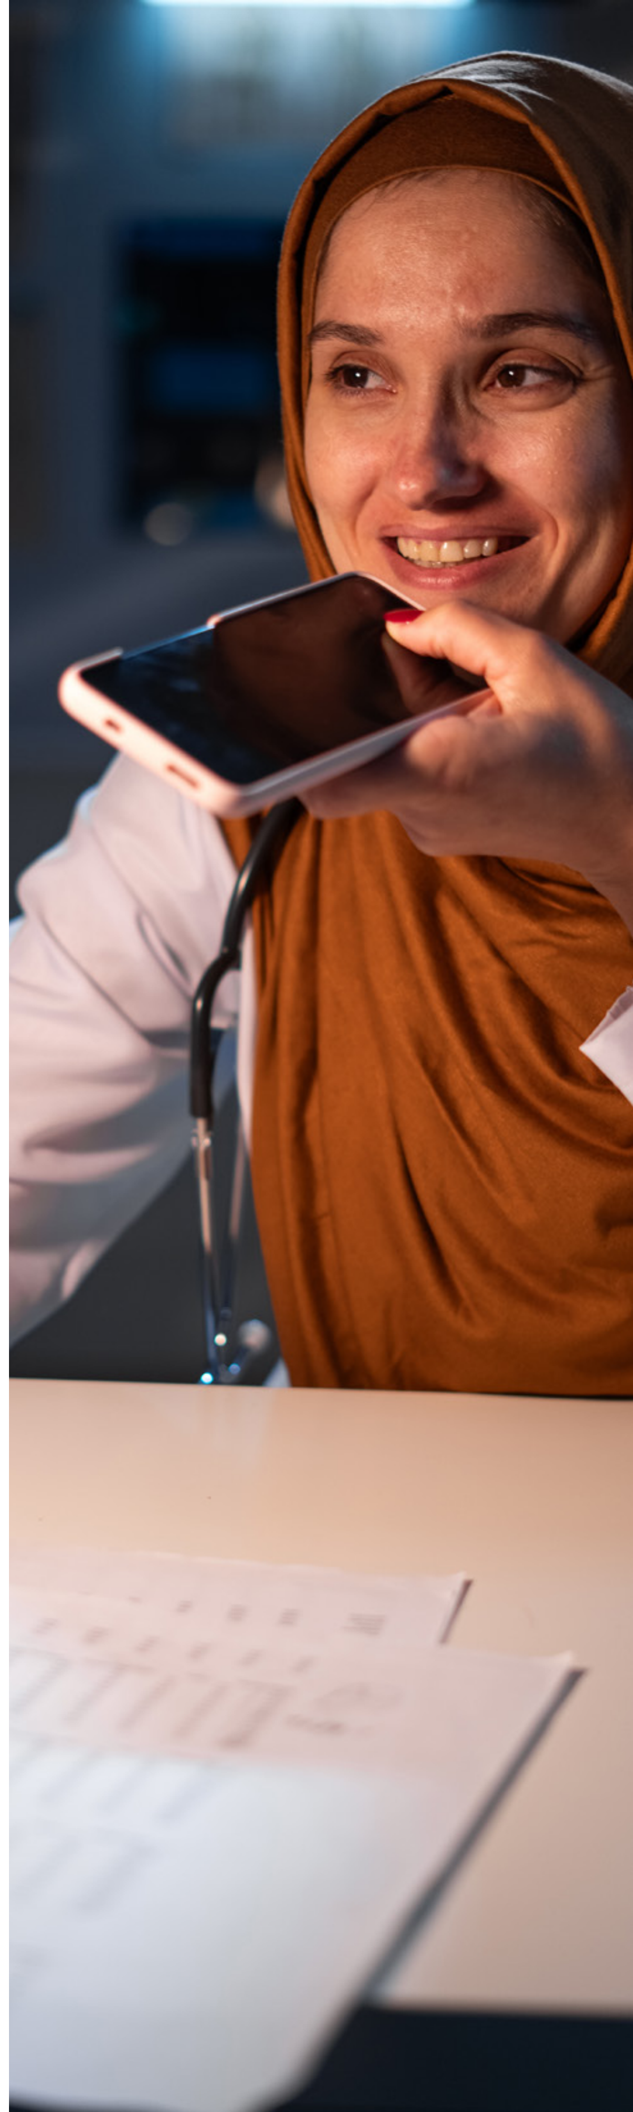

5. Prioritize transparency and explainability in higher-stakes contexts

Our interviewees widely viewed transparency and explainability as essential, particularly as AI tools are applied in higher-stakes contexts. Vendors emphasized that systems need to balance accuracy, interpretability, and speed, noting that the appropriate balance depends on the level of risk associated with a given use case.

Explainability is particularly important for publicly available tools used in decision-making contexts. When LMI is used to inform career, education, or workforce decisions, organizations must be able to clearly explain how and why outputs are generated. Transparency supports understanding, promotes consistent practices, and helps build trust in both AI systems and the products developed using them.

In areas such as job-matching or recruitment, the ability to explain how outputs are generated was seen as especially important for user confidence and institutional trust. For job seekers, knowing when and how they are receiving support from AI systems through transparent practices can empower them to navigate and tailor their use of new tools more confidently.

> **We are trying to balance speed, accuracy, and ease of maintenance when building automated systems.**

For adopters:

Prioritize AI systems that can explain how they generate outputs related to career or employment decisions. Assess whether vendors disclose data sources, model types, and any underlying assumptions: this can help users compare and distinguish between tools and make more informed decisions about how and when to use them.

For developers:

Embed documentation and interpretability tools into AI system design. Clearly identify and explain to clients the risk levels of different models and the trade-offs between accuracy and explainability. For high-stakes applications, ensure there is traceability from input to outcome so users can understand and, when necessary, challenge the results.

6. Acknowledge and mitigate bias

Vendors were clear that while bias cannot be fully eliminated from AI systems, it can and should be acknowledged and mitigated. When it comes to bias related to areas such as gender, race, or geography, vendors described inclusive partnerships and validation processes as important for improving representativeness.

For example, if the data you collect are limited to what is available online, and a model relies exclusively on those data, then it is important to provide a disclaimer. The disclaimer should note that communities or industries that are less likely to post job vacancies online may not be accurately represented in the insights provided.

Some vendors also highlighted the importance of regularly benchmarking outputs across equity-related indicators like gender, race, and geography. Transparency around bias mitigation strategies, beginning at the data-collection stage, was identified as a key practice for strengthening both the quality of the insights and user trust.

For adopters:

Ask vendors how they identify, test, and mitigate bias in their tools. This can shed light on the rigour of a tool's design and development. Evidence of fairness testing, benchmarking, and the use of diverse, reputable data sources are positive signals about how representative the data are.

For developers:

Transparency about data gaps and limitations, including how bias is addressed, is critical—not only for equity, but for assessing how well a dataset represents different segments of the population, given how the data are collected. These factors should be disclosed clearly so users are able to understand what can and cannot be assessed. Incorporating evaluation indicators across different parameters and working with local or community organizations can help reveal and address underrepresented labour realities.

> **There is always a balance between what the data allow you to do and the biases they introduce. Understanding those limitations is essential.**

Considerations

Considerations and signals for trustworthy AI in LMI tools

The six themes outlined here emerged in response to a question that arose repeatedly in our conversations with vendors: How can we build and maintain trust as AI-driven tools become more deeply embedded in Canada's labour market information ecosystem?

While the technologies themselves are evolving quickly, the interviewees noted that many of the core challenges are not new. These reflect long-standing issues in LMI around data quality, interpretation, governance, and appropriate use—but now, they are being amplified by the scale and complexity that AI can introduce.

In this context, the considerations highlighted here are less about finding the most sophisticated technical solutions and more about helping decision-makers identify AI tools that are practical, trustworthy, and fit for purpose. Trust is earned through choices made over time: about how data are sourced and documented, how models are designed and validated, how limitations are communicated, and how human judgment and expertise are integrated alongside automated processes.

Several interviewees pointed to the importance of designing AI-driven LMI tools with governance and accountability in mind from the outset. This includes designing systems that allow for audit and scrutiny, particularly as policy frameworks (such as the Government of Canada's Directive on Automated Decision-Making) continue to evolve. Rather than treating these requirements as external constraints, we should frame them as reference points for shaping more robust, reliable, transparent systems.

Across conversations, there was also a strong emphasis on practicality. Writing clear, accessible documentation, favouring simpler and more interpretable models where possible, being upfront about data gaps or biases, and maintaining human involvement throughout the development and deployment processes were all recommended as ways to keep AI tools aligned with real-world LMI needs. These practices are not guarantees, but they are signals that tools are being developed with an awareness of their broader implications.

Conclusion

At LMIC, we see how the integration of AI and emerging technologies into Canada's LMI space is accelerating. Our conversations with private-sector vendors underscore how strong data-quality practices, careful model-development processes, transparency, and human oversight are central to realizing the long-term benefits of these tools.

This LMIC Insights piece provides a foundation for setting expectations between developers, users, and policy-makers. Our forthcoming work will build on these themes, offering more practical guidance on specific AI technologies used in the LMI space, how these are developed, their limitations, and appropriate use cases. Future work will also include a jurisdictional scan of regulatory and policy approaches internationally and in Canada, situating Canada's role within a broader context.

Together, these findings point to the fundamental conditions for the responsible and effective use of AI in LMI systems. By embedding responsible, human oversight and realistic approaches to bias mitigation, Canada can support a trusted and innovative LMI space.

Acknowledgements

We would like to thank Qiqi Zhang and Annie Kangy for their invaluable insights and contributions as research assistants throughout the development of this report.

References

Adkins-Hackett, L. & Trehan, S. (2024). Connecting skills: Using online job postings to unravel demand for skills in the labour market. Labour Market Information Council (LMIC).

Badman, A., & Kosinski, M. (n.d.). *What is metadata?* IBM. https://www.ibm.com/think/topics/metadata

Canadian Institute for Advanced Research. (n.d.). *The pan-Canadian AI strategy.* https://cifar.ca/ai/#topskipToContent

Castaldo, J., & Norman, P. (2025, June 24). *Canada's first AI minister Evan Solomon on his plan to scale up the industry.* https://www.theglobeandmail.com/business/economy/article-artificial-intelligence-ai-minister-evan-solomon-tech/

CPA Ontario. (2025). *Light, tight, right: Canada's new AI minister aims for balance.* https://www.cpaontario.ca/insights/blog/canada-new-ai-minister-aims-for-balance

Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627–660. https://doi.org/10.5465/annals.2018.0057

Innovation, Science and Economic Development Canada. (2025a). *Canadian sovereign AI compute strategy.* https://ised-isde.canada.ca/site/ised/en/canadian-sovereign-ai-compute-strategy

Innovation, Science and Economic Development Canada. (2025b). *Pan-Canadian Artificial Intelligence Strategy.* https://ised-isde.canada.ca/site/ai-strategy/en

Kosinski, M. (2024, October 29). *What is black box artificial intelligence (AI)?* IBM. https://www.ibm.com/think/topics/black-box-ai

Social Sciences and Humanities Research Council. (2025). *International fellowships to investigate how artificial intelligence is transforming research.* https://sshrc-crsh.canada.ca/en/news/2025/international-fellowships-to-investigate-how-artificial-intelligence-is-transforming-research.aspx

Tzimas, G., Zotos, N., Mourelatos, E., Giotopoulos, K. C., & Zervas, P. (2024). From data to insight: Transforming online job postings into labor-market intelligence. *Information*, 15(8), 496. https://doi.org/10.3390/info15080496

Vankevich, A., & Kalinouskaya, I. (2021). Better understanding of the labour market using big data. Ekonomia I Prawo, 20(3), 677–692. https://doi.org/10.12775/EiP.2021.040