ARTIFICIAL INTELLIGENCE IN CAREER AND WORKFORCE DEVELOPMENT — CANADA

A national research study undertaken by the Canadian Career Development Foundation (CCDF), dmh associates & CareerChatUK, England — **June 2026**

ccdf Canadian Career Development Foundation
dmh associates policy | research | practice
career chat PIONEERING INNOVATION
SUPPORTED IN PART BY: MAGNET

CONTENTS

EXECUTIVE SUMMARY

ABOUT THIS STUDY

This report presents the findings of a national and international research study examining how artificial intelligence (AI) is being integrated into career development practice in Canada. The study was undertaken between October 2025 and March 2026 by the Canadian Career Development Foundation (CCDF) in close partnership with dmh associates - a UK-based international careers research organization - and CareerChatUK -a UK-based AI chatbot technology company specializing in career development. The study was supported in part by Magnet, a Centre of Innovation at Toronto Metropolitan University (TMU).

The research was designed to generate practitioner-grounded, Canada-specific evidence to inform policy, professional development, and ethical frameworks for responsible AI use in career services. It comprised three phases: an international literature review examining 154 sources; a national survey of 370 Canadian career development practitioners (CDPs), offered in English and French and distributed through the CCDF, and CERIC; and three focus group sessions with 38 practitioners facilitated by the CCDF in January and March 2026.

THE CONTEXT: A PROFESSION IN TRANSITION

AI is reshaping the world of work at a pace and scale without recent precedent. World Economic Forum projections indicate that 97% of Canadian companies expect AI and information processing technologies to transform their operations by 2030. McKinsey estimates that up to 30% of current work hours could be automated by the same date, generating millions of occupational transitions - with women 1.5 times more likely than men to need to change occupations, and lower-wage workers up to 14 times more likely than higher-wage counterparts. The populations most exposed to this disruption are precisely those that Canada's publicly funded career development system exists to serve.

In this context, the career development profession faces a dual challenge: helping clients navigate an AI-transformed world of work, whilst simultaneously integrating AI tools responsibly into its own practice. The research reported here shows that Canadian CDPs are engaging with both challenges with seriousness and professional integrity - but without yet having access to the training, guidelines, governance frameworks, or infrastructure that responsible AI adoption requires. The profession is, in a phrase that recurred throughout the research, in transition. The direction and quality of that transition will have lasting consequences for the people it serves.

370
CDP's surveyed nationally

38
Focus group participants

154
Sources in literature review

3
Research phases

KEY FINDINGS

ADOPTION IS ALREADY WIDESPREAD - BUT PREPARATION HAS NOT KEPT PACE

Eighty-four per cent of survey respondents currently use AI tools in their work with clients or students, and 74% have used AI to support their own career development. Yet 82% have received no formal training in AI and ethics, and 47% describe themselves as uncomfortable using AI tools in their practice. This gap between adoption and preparation is the defining professional development challenge of the current moment.

84%
Use AI with clients

82%
No formal training

47%
Uncomfortable using AI

1%
Fully trust AI for client needs

THE HUMAN CORE OF CAREER DEVELOPMENT IS NON-NEGOTIABLE

Practitioners across the survey and all three focus group sessions identified the same set of functions as irreducibly human: self-awareness and values clarification; career decision-making as identity work involving lived experience, culture, and trauma; crisis response and mental health support; personalized action planning; and the relational quality - empathy, professional presence, and accountability - that makes career guidance meaningful. This is not resistance to technology; it is professional expertise. The same practitioners who drew these boundaries with clarity also identified substantial and genuine uses for AI: administrative efficiency, labour market information synthesis, session preparation, accessibility and translation support, and client preparation for interviews.

> ***"The work that we do as CDPs is very personal work, individualized, and specific to each individual; skills specific and story specific; and it requires a human touch, empathy, compassion, trauma-informed practice - you cannot get that from AI tools."***

ETHICAL CONCERN IS HIGH, INFORMED, AND AHEAD OF GOVERNANCE

Ninety-five per cent of survey respondents hold concerns about the ethical dimensions of AI in career development - including bias, accuracy, privacy, over-reliance, and the erosion of human capability. Only 7% believe that AI systems used in career development are currently fair and unbiased. These concerns reflect professional knowledge, not technophobia: practitioners with direct experience of AI outputs that misrepresent clients, reproduce social biases, or produce plausible but inaccurate information are well-placed to assess the risks. Professional and regulatory frameworks have not yet caught up with either the pace of adoption or the depth of practitioner concern.

THE PROFESSION IS ASKING TO BE A PARTNER, NOT A RECIPIENT

The most consistent and forceful message from practitioners - across all three focus groups and throughout the survey commentary - is the demand for genuine professional participation in decisions about AI in career services. CDPs wish to co-design the tools developed for their sector, to be involved in procurement and policy decisions before they are made, and to be recognized as the authoritative professional voice on what responsible AI integration in career development looks like. The proposal, developed in the focus groups, for a voluntary certification framework requiring developers to demonstrate competency in career development before their tools are endorsed for sector use reflects the seriousness and specificity of this demand.

EQUITY IS THE CENTRAL TEST OF RESPONSIBLE AI IN CAREER DEVELOPMENT

The populations served by Canada's career development system - newcomers, Indigenous peoples, racialized communities, people with disabilities, and workers in lower-wage roles most exposed to automation - have the most to gain from AI used wisely and the most to lose from AI used badly. Practitioners consistently framed equity not as one consideration among many but as the fundamental test of whether AI is being used responsibly. Particular concern was expressed about the risk of AI being used as a cost-reduction mechanism - reducing the qualified human workforce - in services where the relational quality of professional practice is not incidental but essential.

> ***"Clients are often coming with multiple, complex and intersectional challenges, including trauma. Trust in systems is often low, and isolation and hopelessness is high. AI strips away the sense that anyone cares about the clients being served."***

SUMMARY OF RECOMMENDATIONS

The full recommendations in Section 6 are addressed to five audiences. The headline recommendations for each are summarised below.

Professional Bodies and Training Providers	• Develop sector-specific professional guidelines for AI use in career development. • Integrate AI literacy into CDP certification as a core competency requirement. • Build a funded, sustained, and sector-specific professional development infrastructure, including a national community of practice. • Explore a voluntary certification framework for AI developers working in career services.
Technology Developers	• Co-design with CDPs from inception - not consultation at testing stage. • Build tools that reflect the relational, contextual, and narrative complexity of career development - not just skills matching. • Prioritize administrative efficiency and LMI access; resist replacing relational professional functions. • Take 'do no harm' seriously as a design constraint, with rigorous bias testing and transparent accountability mechanisms
Employers and Service Providers	• Develop clear internal AI use policies, consent processes, and quality assurance frameworks. • Invest in staff AI training as a professional obligation, not an optional benefit. • Resist pressure to reduce CDP staffing based on AI efficiency arguments. • Engage practitioners meaningfully in AI procurement decisions.
Governments and Funders	• Invest in the profession - training, guidelines, communities of practice - not only in technology. • Develop national policy and regulatory frameworks for AI in career development, co-designed with the sector. • Require evidence of outcomes before endorsing or mandating AI tools for career services. • Apply rigorous equity impact assessments to every investment decision involving AI in career services. • Do not use AI as a mechanism for reducing career development service capacity or the qualified professional workforce.
Career Development Practitioners	• Engage with AI actively and critically - neither avoid nor defer to it. • Maintain and articulate the professional boundaries that the research findings have clearly identified. • Seek out sector-specific professional development and advocate for better provision. • Contribute experience and practice knowledge to the collective evidence base the sector needs.

CONCLUSION

The career development profession in Canada is navigating the AI transition from a position of considerable professional strength. The practitioners who participated in this research demonstrated sophisticated ethical reasoning, clear professional values, and a grounded understanding of both the opportunities and the risks that AI presents for their work and their clients. They are not asking to be protected from technology; they are asking for the conditions - training, guidelines, governance, co-design, and investment - under which they can use it responsibly in service of the people they are professionally committed to supporting.

Meeting that ask is not only a professional development challenge. It is a matter of public policy, professional ethics, and social justice. The recommendations in this report are offered in the conviction that those conditions can be created - and that creating them, as a matter of urgency, and is the right response to one of the most significant transitions the career development profession has faced.

For further information about this research, please contact:

Sareena Hopkins, CEO at the Canadian Career Development Foundation.
Email: s.hopkins@ccdf.ca

Dr Deirdre Hughes, Executive Director, dmh associates and CareerChatUK, England.
Email: deirdre@cicichat.co.uk

1. INTRODUCTION

1.1 CONTEXT AND RATIONALE

Artificial intelligence is reshaping the landscape of work, learning, and career development at a pace and scale without recent precedent. As large language models become embedded in job search platforms, recruitment processes, labour market information systems, and professional tools across every sector, the implications for career development practitioners (CDPs) - the professionals who support individuals through the complex, often uncertain navigation of learning and work transitions - are profound and immediate. CDPs are simultaneously being asked to help clients understand and adapt to an AI-transformed world of work, to integrate AI tools into their own professional practice, and to do so without yet having access to the sector-specific training, professional guidelines, or the evidence base that responsible adoption requires.

This challenge is not unique to Canada, but it has particular dimensions and urgency in the Canadian context. Canada is home to a large and diverse career development workforce operating across post-secondary education, publicly funded employment services, non-profit and community organizations, government, and the private sector. It serves a population that is itself highly diverse - including significant newcomer communities, Indigenous peoples, racialized communities, people with disabilities, and workers in sectors facing rapid technological disruption. And it does so at a moment when, according to World Economic Forum projections, 97% of Canadian companies expect AI and information processing technologies to transform their operations by 2030, generating substantial demand for exactly the kind of support that career development professionals provide.

Public optimism about AI in Canada stands at only 40% - amongst the lowest of any country surveyed in the Stanford AI Index 2025 - reflecting a cautious national culture that is both an asset (it creates space for careful, evidence-based adoption) and a potential barrier (it may slow the profession's engagement with tools that, used responsibly, could significantly extend its reach and impact). Canadian career development professionals are adopting AI at high rates but largely without formal training, clear professional guidelines, or structured support. The profession is, in a phrase that recurs throughout this research, in transition - and the direction and quality of that transition will shape both the profession and the people it serves for years to come.

It is against this backdrop that the present study was commissioned. The research reported here seeks to generate the practitioner-grounded, Canada-specific evidence that responsible AI integration in career services requires: not abstract commentary on technology trends, but a detailed, empirically grounded account of how Canadian CDPs are currently engaging with AI, what they need to do so effectively and ethically, and what policy, professional, and investment decisions would best support them.

1.2 COMMISSIONING, PARTNERSHIP, AND AUTHORSHIP

This research was undertaken between October 2025 and March 2026 by the Canadian Career Development Foundation (CCDF) in close partnership with dmh associates & CareerChatUK and supported in part by Magnet. The study was designed as a genuinely collaborative, international-national partnership, bringing together AI chatbot technology expertise, international careers research capability, and deep knowledge of the Canadian career development sector.

The four organizations involved in this study, and their respective roles, are set out in the table below.

Organization	Role in this study
Canadian Career Development Foundation (CCDF)	National partner organization. The CCDF is Canada's leading national organization in career development. It co-designed the national survey, facilitated outreach through its national networks, and led the design and delivery of all three focus group sessions reported in Section 4.
dmh associates & CareerChatUK	UK partner organization. dmh associates is an international specialist careers research organization. CareerChatUK is a UK-award-winning careers chatbot technology innovator. Deirdre Hughes & Chris Percy led the literature review, co-designed the national survey, analyzed survey results, and co-authored this report.
Magnet at Toronto Metropolitan University (TMU)	Supporting organization. Magnet is a Centre of Innovation focused on the Future of Work, housed at Toronto Metropolitan University. As a convening organization, Magnet supported this study as part of its commitment to encourage exploration of AI integration into workforce development.

The literature review and international scan were led by Dr Chris Percy working closely with Dr Deirdre Hughes. The national survey was co-designed by dmh associates and CareerChatUK in partnership with the CCDF, with analysis and reporting undertaken by dmh associates and CareerChatUK. The focus group sessions were designed and facilitated by the CCDF, with Gillian Johnston and Sareena Hopkins leading all three sessions. The integrated analysis and synthesis presented in Sections 5 and 6 of this report, and the executive summary, were produced jointly by all partner organizations.

1.3 RESEARCH AIMS AND QUESTIONS

The overarching aim of this study is to generate actionable insights into the emerging practices, implementation strategies, and ethical considerations surrounding AI use by Canadian career development practitioners, with a particular emphasis on equity and inclusion. The study addresses four research questions.

Research questions

1. How are Canadian career development practitioners currently using AI tools in their work, and how confident and prepared do they feel to do so?
2. What do Canadian CDPs identify as the most significant benefits, risks, and ethical concerns associated with AI in career development practice?
3. What professional development, governance, and infrastructure do Canadian CDPs need to integrate AI into their practice responsibly and effectively?
4. What should developers, professional bodies, governments, and funders prioritize as they make decisions about AI in career development services?

These questions are addressed through three complementary research phases. The international literature review (Section 2) situates Canadian developments within the broader global evidence base. The national survey (Section 3) provides quantitative and qualitative data on the current state of AI adoption, readiness, and concern across the Canadian CDP workforce. The focus groups (Section 4) provide depth and texture, enabling practitioners to articulate in their own words what they most need, most value, and most fear as AI becomes an increasingly present feature of their professional landscape. The cross-cutting synthesis (Section 5) draws together the findings of all three phases, and the recommendations in Section 6 are grounded in that integrated analysis.

1.4 METHODOLOGY

Literature review and international scan

The literature review was conducted between October and November 2025 using a systematic search of two major academic databases - Scopus and ERIC - supplemented by a targeted review of grey literature from major international organizations. The search strategy identified sources at the intersection of AI technology and career development relevance, covering all English-language academic publications since 2020 and recent grey literature from global consulting firms, international policy organizations, research institutions, and national career development bodies. A tiered prioritization process identified 54 high-priority, 74 medium-priority, and 55 low-priority academic sources for review, along with 23 grey literature papers. Canadian-context papers were automatically upgraded one priority level to ensure depth of coverage. In total, 154 relevant papers are referenced in the full review, which is available as a companion document to this report. A summary of the review's ten key themes is presented in Section 2.

National survey of Canadian CDPs

The national survey was designed by dmh associates and CareerChatUK, in partnership with the CCDF, and distributed through the CCDF's national networks as well as through CERIC. The survey was offered in both English and French, reflecting Canada's bilingual character, and ran between October and December 2025. A total of 370 completed responses were received - 340 in English and 30 in French - from practitioners working across all Canadian regions and the full range of career development settings. The survey combined closed quantitative questions with substantial open-text comment fields, generating a rich mixed-methods dataset. Analysis of the combined English and French data was undertaken by dmh associates and CareerChatUK. Survey findings are reported in Section 3.

Focus groups

Three focus group sessions were designed and facilitated by the CCDF between January and March 2026. The first session was held as a pre-conference workshop at CANNEXUS 2026 in Ottawa on 25 January 2026, with 25 participants. Two subsequent virtual sessions were held on 2 March and 6 March 2026, with six and seven participants respectively. All three sessions were facilitated by Gillian Johnston and Sareena Hopkins of the CCDF. A consistent question framework was used across all three sessions, covering the scope of beneficial AI applications, the boundaries of human professional expertise, professional development needs, and recommendations for developers, governments, and funders. Focus group findings are reported in Section 4.

Integration and synthesis

The findings of all three phases were integrated through a process of thematic synthesis, identifying the points of convergence, divergence, and productive tension between the international literature, the quantitative and qualitative survey data, and the focus group contributions. This synthesis is reported in Section 5 and forms the evidential basis for the recommendations in Section 6.

1.5 EQUITY AND INCLUSION

Equity and inclusion are not peripheral concerns in this research - they are central to its design, its analysis, and its implications. Canada's career development sector serves a highly diverse population. Many individuals most in need of support face compounding vulnerabilities: newcomers navigating unfamiliar labour markets; Indigenous peoples connecting cultural identity with working life; racialized communities facing structural barriers; people with disabilities; and lower-wage workers most exposed to automation. These are precisely the groups most at risk from irresponsible AI deployment - and most likely to benefit when AI is deployed well.

Throughout this report, equity considerations are woven into the analysis rather than treated as a separate category. The literature review examines evidence on who is most exposed to workforce disruption and what this means for demand for career guidance. The survey data includes self-reported demographic information enabling analysis across gender, disability status, Indigenous identity, and equity-deserving group membership. The focus group findings include sustained engagement with questions of bias, accessibility, Indigenous perspectives, rural-urban disparities, and the risk that AI-driven services may serve already-advantaged populations more effectively than those facing multiple barriers. The recommendations in Section 6 include a dedicated equity strand addressing what equitable AI integration in career development services requires in practice.

1.6 STRUCTURE OF THIS REPORT

This report is organized as follows:

Section 2 presents a condensed synthesis of the international literature review, covering ten key themes from 154 sources. Section 3 reports the findings of the national survey of 370 Canadian CDPs. Section 4 presents the narrative findings of three focus group sessions with 38 Canadian CDPs. Section 5 draws together findings across all three phases in a cross-cutting thematic synthesis. Section 6 sets out recommendations addressed to five distinct audiences: career development practitioners themselves; professional bodies and training providers; technology developers; employers and service providers; and governments and funders. An executive summary is provided at the front of the report for readers who require a concise overview of the key findings and recommendations.

The survey instrument in English and French, and the focus group discussion guide are available as appendices on request from the research team. The full literature review with complete reference list is available as a separate companion document.

2. LITERATURE REVIEW FINDINGS

This section presents a condensed synthesis of the international literature review undertaken as the first phase of this research study, conducted between October and November 2025 by Dr Chris Percy and Deirdre Hughes of dmh associates and CareerChatUK. The full review - which examined 154 sources drawn from a systematic search of the Scopus and ERIC databases, supplemented by grey literature from major international organizations - is available as a companion document. What follows is a thematic summary designed to provide the international and Canadian context within which the survey and focus group findings reported in Sections 3 and 4 should be read.

The review identified ten cross-cutting themes in the literature on AI, large language models (LLMs), and career development published since 2020. These themes are summarised in the box below and discussed in turn.

Ten Themes from the Literature

1. Technology and use case evolution: from early machine learning to large language models
2. A dramatic increase in research volume, particularly since 2024
3. Empirical evidence remains limited, but early findings are cautiously positive
4. Growing and diversifying usage of generative AI for career-related purposes
5. Limited integration of AI into career development professional practice
6. Rapid adoption alongside significant professional readiness gaps
7. Substantial workforce transitions and skills evolution underway
8. Human-AI collaboration as the dominant emerging model
9. Ethical governance and implementation challenges are pressing and largely unresolved
10. Geographic and economic disparities in AI access and optimism

2.1 TECHNOLOGY AND USE CASE EVOLUTION

The published literature reflects a rapid and ongoing evolution in the types of AI technology being applied to career development. Earlier papers - predominantly from 2020 to 2023 - focus on machine learning algorithms, natural language processing, and recommender systems designed to provide career advice, suggest courses, or match candidates to vacancies. These technologies draw primarily on established career guidance frameworks such as Holland's RIASEC typology, codifying existing professional principles into automated systems rather than generating new insight from data. Their evaluations are typically constrained by the quality of training data and the absence of longitudinal outcome data linking recommendations to long-term career results.

From 2024 onwards, the literature shifts markedly towards large language models and conversational AI.

Studies describe a new generation of tools - including AI chatbots, future-self simulation systems, AI-powered mock interview platforms, and LLM-driven mentorship matching services - that engage users in open-ended dialogue rather than selecting from predefined option sets. This shift represents a qualitative change in what AI can offer career development: the capacity to respond to the complexity, ambiguity, and individuality of a person's career situation in ways that earlier rule-based systems could not. At the same time, the literature notes that most of these newer tools are built on general-purpose, sector-neutral LLM providers such as OpenAI, Meta, and Anthropic, rather than on systems designed specifically for career development purposes - a gap that practitioners in this study consistently identified as a priority to address.

2.2 A RAPIDLY EXPANDING RESEARCH FIELD

A key finding from the systematic review is the sheer pace at which the field is growing. From 2020 to 2022, the review identified only two to seven papers per year meeting the eligibility criteria. By 2025, this had reached 60 publications in a single year - a tenfold increase in five years. Fifty per cent of the high-priority papers in the review were published in 2025 alone. This explosion in research volume reflects the broader acceleration of AI development and deployment, and it has direct implications for professional practice: the knowledge landscape is changing faster than the professional development infrastructure can currently track or transmit.

The prevailing stance in the literature is one of qualified optimism. Most papers discuss the potential benefits of AI tool adoption - personalization, efficiency, accessibility, scale - whilst acknowledging risks that are treated primarily as factors to be managed rather than grounds for restricting adoption. This framing mirrors the conditional, pragmatic engagement that the practitioners surveyed and interviewed in this study consistently described.

2.3 EMPIRICAL EVIDENCE: LIMITED BUT CAUTIOUSLY POSITIVE

Despite the volume of publications, the empirical base for strong claims about AI's effectiveness in career development remains thin. The most common paper types introduce new tools at concept or prototype stage, discuss general frameworks for AI use, or identify ethical concerns - without providing robust evaluative evidence about outcomes. Where user satisfaction data exists, it tends to be positive, but user satisfaction is a weak proxy for the career outcomes that ultimately matter.

A small number of higher-quality evaluative studies offer more substantive evidence. Deshmukh and Bajaj (2024) report, using a randomized A/B testing methodology with 5,000 users, that a retrieval-augmented generation (RAG) system for personalized job recommendations produced a 26% increase in recommendation relevance, a 46% increase in user engagement, and a 41% increase in job application rates relative to a keyword-based baseline. Jin (2025) reports job-matching accuracy improving from 72% to 91% using AI-enhanced placement forecasting, alongside a 35% increase in six-month job placement rates. Xiao and Zheng (2025) found that regular ChatGPT use significantly enhanced students' employment confidence and proactive job-seeking behaviours. These results are encouraging, but they reflect specific implementations in specific contexts, and the review cautions against extrapolating from them to the sector as a whole without replication and broader evaluation.

Research transparency is also a concern: very few studies publish their underlying datasets or code, making independent verification of findings difficult. The review highlights this as a structural weakness in a field that needs robust, replicable evidence to guide responsible implementation.

2.4 GROWING AND DIVERSIFYING USAGE

Emerging evidence from multiple national contexts suggests that generative AI has already achieved substantial penetration among both career service users and career development professionals - often outrunning the professional guidance structures intended to support safe and effective use. A study of 502 Chinese college students found regular generative AI use for career exploration, with uptake shaped by technological readiness, efficacy beliefs, and perceived output quality. UK higher education research (Freeman, 2025) found that 92% of UK students now use generative AI in their studies, with staff AI literacy having more than doubled in one year. Earlier UK research by Hughes, Percy and Tolond (2024) found that while a third of students were open to using AI for careers advice, concerns about trust, data privacy, and unequal access remained significant, and more than half of students at that time had not heard of or did not understand LLMs.

In workplace settings, studies of HR AI adoption (Rajeev et al., 2025; Kim et al., 2025) document growing use of AI for resume screening, candidate sourcing, and turnover prediction - with direct implications for the clients that CDPs serve, who are increasingly navigating AI-mediated recruitment processes that they may not fully understand. Kochling et al. (2025) found that perceived fairness was lowest when an AI agent was the sole decision-maker in career-related assessments, reinforcing the importance of human oversight in high-stakes applications.

2.5 LIMITED INTEGRATION INTO PROFESSIONAL PRACTICE

Despite the volume of tools being developed and the growing usage documented above, the review finds that very few AI applications have been designed with the specific professional context of career development practice in mind. Most tools are presented as standalone interventions - a chatbot a student might use independently, an LMI aggregator a practitioner might point a client towards - rather than as instruments integrated into a professional service delivery model. Donald and Straby's (2025) three-phase narrative storytelling approach, which uses AI-generated reports as the basis for deepened human dialogue between practitioner and client, is highlighted in the review as a rare example of genuine professional integration.

This gap between tool development and professional integration is significant. It means that the default mode of AI use in career services is supplementary and informal rather than structured and professionally guided - with all the risks of inconsistent quality, inadequate safeguarding, and missed opportunity that this implies. It also means that the professional expertise of CDPs is not being systematically leveraged in the design of AI tools - a concern that practitioners in the focus groups of this study articulated with considerable force.

2.6 RAPID ADOPTION, SIGNIFICANT READINESS GAPS

The literature consistently describes a profession in transition - one in which adoption has outpaced preparation

across multiple national contexts. The CERIC 2024 Survey of Career Service Professionals found Canadian CDPs ranging from enthusiastic early adopters to those who are hesitant due to ethical concerns, knowledge gaps, and uncertainty about professional implications. These findings are directly echoed in the national survey reported in Section 3 of this study, where 82% of respondents reported having received no formal AI training, despite 84% already using AI in their practice.

Research by Cukier et al. (2024) at Toronto Metropolitan University, drawing on a longitudinal survey of 5,855 Canadians, found that approximately three in ten employed adults were using AI at work - mostly without employer-provided training or organizational guidance. The study identified age-related gaps in AI familiarity, a narrowing gender gap, and a significant disconnect between post-secondary curricula and employer needs. Notably, it also found that immigrants, racialized, and Indigenous respondents demonstrated higher familiarity with AI tools than white and Canadian-born participants - challenging assumptions about where the AI skills deficit lies and pointing to structural barriers in hiring rather than capability gaps in diverse communities.

Accenture (2025) reports that despite widespread AI adoption, only 13% of organizations achieve significant value from their AI investments - a figure attributed in part to the imbalance in investment portfolios: organizations typically spend three times more on technology than on the people, processes, and leadership development needed to use it effectively. This structural imbalance has direct relevance for public investment decisions in career development services.

2.7 WORKFORCE TRANSITIONS AND THE SCALE OF NEED FOR CAREER GUIDANCE

One of the most consequential findings from the grey literature concerns the scale of workforce disruption that AI and related technologies are expected to generate - and the corresponding expansion of need for career development support. McKinsey Global Institute projections suggest that up to 30% of current work hours in the United States could be automated by 2030, requiring approximately 12 million occupational transitions. A companion analysis for Europe projects a similar number of transitions, at double the pre-pandemic pace. The equity dimensions of this disruption are stark: women are 1.5 times more likely than men to need to change occupations as a result of automation, and lower-wage workers face up to 14 times the likelihood of needing transitions compared to higher-wage counterparts.

The World Economic Forum's Future of Jobs Report 2025 - drawing on perspectives from over 1,000 global employers representing 14 million workers - projects that 97% of Canadian companies expect AI and information processing technologies to transform their operations by 2030. Demand is rising for roles including AI and Machine Learning Specialists, Security Management Specialists, and Software Developers, whilst routine office support and basic customer service roles continue to decline. The skills employers need are shifting towards technological literacy, critical thinking, creativity, and social-emotional intelligence - skills that career development professionals are uniquely positioned to help individuals develop and articulate.

The implication is not that career guidance faces displacement, but that it faces dramatically expanded demand. The populations most likely to need intensive career support - lower-wage workers, older adults, women in administrative roles, workers in sectors most exposed to automation - are precisely the populations that publicly funded career services are designed to reach. This context gives urgency to the investment and policy questions raised throughout this report.

2.8 HUMAN-AI COLLABORATION AS THE EMERGING MODEL

Across the international literature, from major consultancies to academic researchers, a consistent model is emerging: AI as amplifier of human capability rather than replacement for it. PwC's AI Jobs Barometer 2025, drawing on analysis of close to one billion job advertisements across six continents, provides large-scale empirical evidence that AI exposure makes workers more valuable rather than less - even in highly automatable occupations. McKinsey and Accenture similarly emphasize that the most successful AI implementations are those that keep humans in the loop, combining the speed and scale of AI with the judgement, empathy, and contextual understanding of skilled professionals.

Microsoft's 2025 Work Trend Index describes the emergence of what it terms the 'frontier firm' - one in which every employee works alongside AI agents, requiring new skills in human-agent collaboration and a shift in professional identity from task executor to AI orchestrator. Canadian data from the same study found that 76% of Canadian organizational leaders are confident in using AI agents to expand capacity, whilst 42% are prioritizing upskilling of existing workers over external hiring to meet emerging skills needs.

For career development, this framing has significant implications. CDPs are not competing with AI; they are being called upon to help clients navigate an AI-augmented world of work - developing their AI literacy, building their confidence in human-AI collaboration, and ensuring that technology amplifies rather than diminishes their capacity to achieve their career goals. This is a genuinely new professional challenge, and one that the sector's training and professional development infrastructure has not yet fully risen to meet.

2.9 ETHICAL GOVERNANCE: PRESSING CHALLENGES, INCOMPLETE FRAMEWORKS

The literature identifies a wide range of ethical challenges in the deployment of AI for career development and finds that the professional and regulatory frameworks available to address them are not yet adequate. Key challenges include: algorithmic bias and its potential to reproduce or amplify existing inequalities in access to opportunity; privacy concerns relating to the collection, storage, and use of personal career data; transparency and accountability gaps in systems where the basis for recommendations is opaque to both practitioners and clients; and the risk of professional liability when AI-generated advice leads to harmful outcomes.

Ooi and Wilkinson (2025) examine the ethical codes of professional counselling bodies in four countries and find that none are fully equipped to govern the use of generative AI in practice - a finding that applies equally to career development professional standards frameworks. Bilderback et al. (2025) propose an ethics-led framework for AI-driven workforce development that embeds principles of justice, transparency, autonomy, and privacy from the design stage rather than as retrospective safeguards. Attard-Frost et al. (2024) find that across 84 Canadian AI governance initiatives, most focus on industry and innovation rather than on workforce development services, AI education, or digital infrastructure - a significant gap in the national policy landscape.

Nesta's Public AI Assurance project (2025) highlights a disconnect between policy ambitions for responsible AI adoption in public services and public skepticism about AI's societal impact. Evidence from public polling reveals significant pessimism about AI, particularly in developed economies, suggesting that technical compliance with ethical standards is insufficient without genuine engagement with public concerns about fairness, accountability, and the distribution of AI's benefits and risks.

2.10 GEOGRAPHIC AND ECONOMIC DISPARITIES

The final theme from the literature review concerns the very uneven global distribution of AI optimism, access, and benefit. The Stanford AI Index 2025 (Maslej et al., 2025) reports that in countries such as China (83%), Indonesia (80%), and Thailand (77%), strong majorities view AI as more beneficial than harmful. In Canada, only 40% hold this view - one of the lowest rates in the study, alongside the United States (39%). Whilst optimism has grown by 8 percentage points in Canada since 2022, the country's relatively cautious stance has implications for the political and professional environment within which AI adoption in career services is being negotiated.

Deloitte's analysis of sovereign AI strategies (2024) highlights Canada's three-pillar approach - improving private investment, developing public infrastructure, and enhancing access to compute resources - and the role of domestic partnerships with firms such as Cohere in building national AI capacity. The IAEVG's Digital Developments in Guidance magazine (2025) documents examples of AI innovation in guidance across multiple global contexts, alongside the structural inequalities - in infrastructure, skills, and governance - that risk widening rather than narrowing the gap between AI-advantaged and AI-disadvantaged populations.

For career development, the equity implications of geographic and economic disparities in AI access are direct: if AI-mediated career information and guidance tools are most accessible to those who are already advantaged - those with reliable internet access, digital literacy, and the confidence to engage with new technologies - then AI risks compounding the very inequalities that career guidance systems are designed to address. This concern runs through the survey and focus group findings in the sections that follow, and it frames the equity-centred recommendations presented in Section 6 of this report.

2.11 SUMMARY: WHAT THE LITERATURE TELLS US

The international literature presents a picture of a field in rapid, consequential transition. AI technologies are becoming more sophisticated, more widely used, and more deeply embedded in both career-related services and the workplaces that career development practitioners help their clients navigate. The evidence base for AI's effectiveness is growing but remains limited: positive early findings from well-designed evaluations are encouraging, but the sector-specific, longitudinal, and comparative evidence needed to guide responsible implementation at scale is not yet available.

Five findings from the literature are particularly relevant as context for the Canadian practitioner data that follows.

- **Firstly,** adoption is outpacing preparation across all national contexts studied - Canada is not unusual in this respect, but the scale of the readiness gap documented in this study is nonetheless a cause for concern.
- **Secondly,** the most effective AI implementations keep humans in the loop - professional expertise and relational capability remain essential, and AI is most valuable as an amplifier of that expertise.
- **Thirdly,** AI tools are overwhelmingly designed without meaningful input from career development professionals - the sector's knowledge base is not being systematically leveraged in the technology it is being asked to adopt.

- **Fourthly,** the ethical frameworks needed to govern responsible AI use in career development are incomplete, and urgent development work is needed at both professional and policy levels.
- **Finally,** the scale of workforce disruption projected for the coming decade will dramatically expand the demand for high-quality career guidance - making investment in the profession's overall capacity, as well as its specific capacity to use AI responsibly, not merely a professional development question, but a matter of public policy significance.

3. SURVEY FINDINGS: CANADIAN CAREER DEVELOPMENT PRACTITIONERS AND AI

The national survey of Canadian career development practitioners (CDPs) was conducted between October and December 2025 as the second phase of this research study. Distributed through the networks of the Canadian Career Development Foundation (CCDF), Magnet at Toronto Metropolitan University (TMU), and CERIC, the survey received 370 completed responses - 340 in English and 30 in French - from practitioners working across all regions of Canada and the full range of career development settings. This section presents a thematic analysis of the survey findings, drawing on both the quantitative response data and the substantial body of qualitative commentary that respondents contributed throughout the instrument.

The survey sample was predominantly female (81%), predominantly employed full-time (85%), and represented a broad cross-section of sectors: post-secondary education (29%), non-profit and community organizations (28%), government (14%), publicly funded employment services (13%), and the private sector (12%). Fourteen per cent of respondents identified as a person with a disability, and a small number identified as Indigenous (First Nations, Metis, or Inuit), BIPOC, LGBTQ+, or immigrant. This demographic profile broadly reflects the composition of the Canadian career development workforce, whilst underscoring the importance of equity considerations in any analysis of AI adoption and readiness.

The picture that emerges from the survey data is not one of a profession divided between enthusiastic adopters and resistant skeptics. Rather, it is a picture of a profession engaged in a careful, conditional negotiation with a rapidly changing technological environment - one that recognizes AI's genuine utility whilst maintaining a clear-eyed understanding of its limitations and risks. As the data show, the prevailing disposition is neither technophobic nor uncritically techno-optimistic, but one of cautious, conditional engagement - grounded in professional values, ethical awareness, and a strong commitment to the human relationship at the heart of career development practice.

Summary of key findings

- 84% of respondents currently use AI tools in their work with clients or students, at least sometimes.
- 74% have personally used AI tools to support their own career development or job search.
- Only 47% report feeling comfortable using AI tools in their practice.
- Only 1% fully trust AI to evaluate client needs accurately.
- 95% are concerned about ethical issues in the use of AI in career development.
- Only 18% have received any formal training in AI and ethics.
- 86% consider it extremely or quite important that AI tools be developed and evaluated by career development experts.
- 56% believe clear guidelines on AI use in career development are extremely important.

3.1 FAMILIARITY, ADOPTION, AND PERSONAL USE

The survey opened by asking practitioners about their familiarity with AI and their direct experience of using it, both personally and professionally. The responses reveal a profession in which AI has already become part of the working landscape for the substantial majority - but in which familiarity and adoption have significantly outpaced comfort, preparation, and structured professional guidance.

When asked what comes to mind when they think of AI, respondents offered a wide-ranging set of associations that reflect both engagement and ambivalence. The largest single cluster of responses - 24% of the 346 who answered this open question - characterized AI as a tool: something that requires skilled and critical human use, that enhances rather than replaces professional judgement, and that carries genuine value alongside genuine risk. A further 16% emphasised timesaving and efficiency as their primary association, describing AI as a means of handling the routine, repetitive, and information-intensive aspects of their work more quickly. Seven per cent identified AI's research and synthesis functions as its most salient feature, and 5% cited the capacity to process large volumes of data rapidly.

These positive associations were consistently qualified. Across the open-text responses, themes of caution, critical evaluation, and the need for human oversight recurred throughout. Concerns raised included the risk of AI misinformation and hallucination (cited by 3% as a primary concern), privacy and data security risks (1%), the potential for misuse (raised repeatedly in qualitative commentary), the environmental impact of AI infrastructure (3%), and - with particular resonance for a profession whose purpose is to support human capability - the risk that over-reliance on AI may erode the critical thinking, creativity, and interpersonal capacities that both practitioners and their clients need to develop.

> ***"It's a Swiss Army Knife. It can be a coach, an editor, a writer, a coder, a personal assistant, a designer, and so much more. Without critical thinking, understanding of the limitations of AI - hallucinations, sycophantic behaviour - and knowledge of how data is used, there's also a lot of risk involved."***

Seventy-four per cent of respondents reported having personally used AI tools to support their own career development or job search, indicating a high level of direct personal experience with the technology. When asked about professional use - whether they currently use AI tools in their work with clients or students - 84% reported doing so, with 34% using AI often and 49.5% sometimes. Just 4% reported having no interest in using AI in their practice. This level of adoption is striking, particularly when set against the data on comfort levels and formal training discussed below.

3.2 COMFORT, TRUST, AND THE READINESS GAP

Despite widespread and growing adoption, a significant proportion of practitioners reported discomfort with using AI tools in their practice. Thirty-one per cent described themselves as somewhat uncomfortable, and a further 16% as very uncomfortable - meaning that approaching half of respondents (47%) are using AI in their professional work at a level of comfort that falls below 'quite comfortable.' Only 16% described themselves as very comfortable. This gap between adoption and comfort is a defining feature of the current moment in the profession, and one of the most important findings of this survey.

The data on trust reinforce this picture. Respondents were asked to what extent they trust AI systems to evaluate the needs of clients or students accurately. The responses were unambiguous: 17% expressed no trust at all, 66% expressed only partial trust, 16% trusted AI quite a bit, and just 1% - four individuals in a sample of 370 - reported total trust. Taken together, 83% of practitioners report some level of reservation about AI's ability to assess client needs reliably. This level of skepticism is not irrational: it reflects an informed awareness of the documented limitations of AI systems in contexts requiring nuanced human judgement, and it aligns closely with the findings from the focus groups reported in the previous section.

The qualitative commentary throughout the survey consistently framed trust not as binary but as conditional and contextual. Respondents drew clear distinctions between tasks where AI can be trusted to generate a useful starting point - information synthesis, administrative drafting, labour market summaries - and those where its outputs require critical verification, professional oversight, or should not be used at all. The concept of AI as a tool that is as good as the person using it recurred throughout, reflecting an understanding that the quality of AI-assisted practice depends heavily on the practitioner's own expertise, critical literacy, and professional judgement.

> ***"AI is only as good as the user. We need to provide it with the content to which we are wanting fulsomely explored. It is important that we also review and edit their responses again as we teach them the specifics of our work, client needs, etc. If taught and used responsibly it is a valuable resource."***

3.3 HOW PRACTITIONERS AND CLIENTS ARE USING AI

The survey asked practitioners both about their own uses of AI in practice and about the uses they observe among their clients. The two sets of findings tell a complementary story: AI is being integrated most readily into the informational and administrative dimensions of both practitioner and client activity, whilst the more deeply personal and relational dimensions of career development remain human territory.

Practitioner uses

Across the survey's questions on AI tool use and the extensive qualitative commentary, the most commonly reported practitioner applications clustered around three functions. The first and most prominent was administrative support: drafting communications, summarising information, generating templates, note-taking, and reducing the time spent on routine documentation. The second was information gathering and synthesis: using AI to research labour market information, generate occupation descriptions, summarise large volumes of material, and produce accessible summaries for clients with different comprehension levels. The third was preparation support: using AI before client sessions to brainstorm questions, develop tailored prompts, and think through approaches to specific client circumstances.

Additional uses reported by respondents included mock interview facilitation, resume and cover letter assistance (as a starting point for practitioner-guided refinement rather than as a finished product), translation for newcomer clients, and the use of AI tools such as Otter for session note-taking - subject, in each case, to data security conditions. Practitioners also reported using AI to support their own professional development: researching emerging practices, staying current with labour market trends, and developing workshop and programme materials.

The tools most commonly referenced by respondents were ChatGPT (OpenAI), Microsoft 365 Copilot, Google Gemini, Claude (Anthropic), and Perplexity. A smaller number referenced specialist tools, including Otter (for transcription), Canva (for design), and Google Translate (for language support). Several respondents noted that their organizations had not yet established clear policies governing which tools were approved for professional use - a governance gap that will be discussed further in Section 3.5.

Client uses

Practitioners' observations of client AI use centre on the job application process. Resume and cover letter assistance was the most frequently identified client application, with 61% of respondents reporting that their clients use AI at least quite a bit for this purpose, generating a weighted average score of 2.83 out of 4. Interview and application preparation was the second most common use, reported by 51% of respondents as used quite a bit, with a weighted average of 2.54.

Notably, a significant proportion of respondents acknowledged that they simply do not know how their clients are using AI outside of sessions - 'this topic does not come up in conversation' was a representative comment. This itself reflects something important: that client AI use is largely unobserved and unguided, occurring outside the professional relationship and without the critical scaffolding that practitioners believe is essential to its responsible use.

The qualitative commentary on client use ranged from cautious concern to considered optimism. Several practitioners expressed worry that clients were submitting AI-generated application documents without adequate review or personalization, potentially undermining both their authenticity and their effectiveness in screening processes increasingly managed by applicant tracking systems. Others saw AI-assisted applications as an equity opportunity - a way for clients who lack the social capital, language confidence, or professional networks to access guidance to generate a viable starting point.

> ***"Utilizing AI and results can be a fantastic way to start conversations or generate ideas for new directions - the point is that it's a starting place, not a be-all, end-all approach."***

3.4 CONCERNS: ETHICS, BIAS, AND THE LIMITS OF AI

The survey data paint a picture of a profession that is ethically alert to the risks of AI - and, in many respects, ahead of the broader societal conversation on these issues. Across multiple questions, practitioners expressed high levels of concern about AI accuracy, bias, privacy, over-reliance, and the potential erosion of human capabilities that they regard as essential both to their own practice and to their clients' long-term development.

Risk and concern

When asked about their overall level of concern regarding risks involved in clients and students using AI tools, 87% of respondents expressed at least some concern: 11% extremely concerned, 26% very concerned, and 50% somewhat concerned. Only 11% reported no concern at all. The qualitative commentary associated with this question - 206 responses - reveals the nature of practitioners' concerns in considerable detail.

The most prominent themes in this commentary were the accuracy of AI-generated information and the need for critical verification; the risk that clients who cannot distinguish reliable from unreliable outputs will make consequential decisions based on flawed information; privacy and data security concerns, particularly in relation to clients who may not understand how their personal information is processed and stored when they use AI tools; and the risk of over-reliance - both by clients who may use AI as a substitute for developing their own capabilities, and by practitioners who may default to AI-assisted interventions without adequate professional reflection on whether they are appropriate.

> ***"The work that we do as CDPs is very personal work, individualized, and specific to each individual; skills specific and story specific; and it requires a human touch, empathy, compassion, trauma-informed practice - you cannot get that from AI tools."***

A distinctive concern raised by several respondents concerned what might be described as a competency displacement paradox: if AI does the analytical and expressive work - drafting the cover letter, synthesizing the self-assessment, structuring the interview response - are clients being deprived of the opportunity to develop precisely the skills that employment will require of them? This question, which also surfaced in the focus groups, points to a deeper tension in the relationship between AI-assisted career support and the developmental purposes of career guidance.

Ethics and bias

On the question of ethical concerns specifically - including bias, fairness, and transparency in AI systems - the survey found very high levels of practitioner concern. Twenty-eight per cent reported being extremely concerned, 33% quite concerned, and 34% somewhat concerned. Just 5% expressed no concern at all. In other words, 95% of the survey sample hold some level of ethical concern about AI in career development.

This concern is grounded in an informed understanding of how bias operates in AI systems. When asked whether they believe AI systems used in career development are currently fair and unbiased, just 7% said yes. Thirty-seven per cent said no, and 56% said they were unsure — reflecting an honest acknowledgement that this is a complex question, but also a significant knowledge gap that training and professional development must address.

The qualitative responses on bias were prevalent in the survey. Practitioners drew on real experience. They described AI chatbots that simply tell users what they want to hear, rather than offering honest challenge or alternative perspectives. They pointed to resume screening tools known to disadvantage candidates from particular demographic groups. They raised concerns about LLMs trained on internet data that reflects existing social inequalities and misinformation. Time and again, respondents returned to the same concern: that AI risks reproducing and deepening the structural disadvantages already faced by newcomers, racialized communities, Indigenous peoples, and people with disabilities. This is not abstract academic concern - it reflects the profession's lived commitment to equity.

> ***"One of the first pitfalls I share with clients is that chatbots are notoriously biased and often discriminate. So I cannot say yes to any system using chatbot technology. It requires some basic literacy and then competence in effective communication - good prompt engineering as an example."***

The question of guidelines

The survey asked directly about the need for clear guidelines governing AI use in career development. The response was unambiguous: 56% considered clear guidelines to be extremely important, and a further 30% considered them quite important. Eighty-six per cent of respondents thus place the establishment of professional guidelines among their highest priorities for the sector. When asked who should bear primary responsibility for ensuring ethical AI use, respondents cited professional bodies, government, and - importantly - individual practitioners themselves, reflecting an understanding that responsibility in this area is distributed rather than singular.

A theme that emerged in the qualitative data was the perceived inadequacy of existing ethical frameworks to address the specific challenges posed by AI in career development contexts. Professional codes of conduct developed for human-delivered practice do not readily translate to AI-augmented environments, and practitioners expressed a clear need for sector-specific guidance that addresses the concrete dilemmas they face in daily practice: how to handle client data, when to disclose AI use, how to evaluate the reliability of AI-generated information, and how to support clients who are themselves using AI without guidance.

3.5 PREPAREDNESS AND PROFESSIONAL DEVELOPMENT NEEDS

One with the most direct implications for policy and investment - is the gap between the scale of AI adoption and the level of professional preparation that underpins it. The data reveal a profession using powerful, consequential tools without adequate training, clear guidance, or structured support.

When asked how well prepared they feel to support clients and students in using AI tools for career development, just 9% described themselves as extremely prepared, and 22% as quite prepared. Fifty-two per cent felt only somewhat prepared, and 17% not at all prepared. Taken together, 69% of respondents report a level of preparedness that falls below 'quite prepared' - a significant professional development gap in a context where 84% are already using AI with clients.

The training data are illuminating. Only 18% of respondents reported having received any formal training or education related to AI and ethics in career development. Eighty-two per cent have not. This means that the overwhelming majority of Canadian CDPs currently using AI in their practice are doing so on the basis of self-directed learning, peer knowledge-sharing, or simple trial and error - without the structured, evidence-based professional preparation that responsible practice in a complex and rapidly evolving technological environment requires.

Practitioners' self-assessed competency levels reflect this gap. When asked to rate their competence across a range of AI-related capabilities, the area rated highest was the effective use of standard AI tools such as ChatGPT and basic prompt writing, with a weighted average of 2.67 out of 4. Competencies relating to the ethical dimensions of AI use, bias detection, data governance, and the integration of AI into professional practice frameworks were rated lower. The qualitative commentary in this section included repeated calls for training: '

Need some training!', 'This should be an entire course in future CDP programmes', 'Hands-on training in specific tools being used most commonly in the field.'

What practitioners want from professional development

The survey asked practitioners directly about the kinds of training and resources they would like to support their ethical and effective use of AI. The responses pointed consistently towards three priorities. First, sector-specific and practically grounded training: resources that go beyond generic AI literacy and are directly applicable to career development contexts, client populations, and professional practice frameworks. Second, training in prompt engineering and critical evaluation: the technical skills required to use AI effectively and the analytical skills required to evaluate its outputs critically - with an honest account of AI's limitations built in. Third, structured and ongoing learning communities: communities of practice, peer exchange forums, and living resources that evolve as the technology and the professional context evolve, rather than one-off training events that rapidly become outdated.

The importance that practitioners place on expert-led, sector-evaluated AI tools was also clearly evident. When asked how important it is that an AI tool has been developed and evaluated by career development experts before being recommended for use, 43% said this is extremely important and 35% quite important - meaning that 78% place high priority on expert evaluation as a condition of professional endorsement. This finding has direct implications for the development and procurement of AI tools for career services, and reinforces the call - articulated consistently in the focus groups - for CDPs to be recognized as co-designers and evaluators of any tools intended for use in their sector.

> ***"Tools need to be developed by experts in AI, of course, but need to involve the input of career development professionals at all stages and be tested by CDPs before launch."***

3.6 WHAT CDPS WANT AI TO DO

The survey asked practitioners what they most want from AI in career development, and what they believe their clients and students most want. The responses reinforce the picture that emerged from the focus groups: practitioners want AI to amplify their capacity to deliver high-quality, human-centred practice - not be a substitute for it.

For practitioners themselves, the most valued AI capabilities were information access and synthesis, administrative efficiency, and support for the preparatory dimensions of practice - the research, questioning, and session planning that precede and support the human encounter with the client. There was consistent interest in AI that can provide better, faster access to labour market information and occupational data; support the development of tailored resources and materials; handle the documentation and communication burden that occupies a disproportionate share of practitioners' professional time; and help practitioners stay current in a rapidly changing labour market environment.

For clients, practitioners identified the primary AI use cases as job application assistance (resumes, cover letters, employer research), interview preparation, information searching, and - for some client groups - translation and accessibility support. Across both sets of responses, a consistent pattern emerged: AI is valued as a starting point, a research accelerator, and an administrative assistant. It is not valued - and should not be used - as a substitute for the professional relationship, the exercise of professional judgement, or the deeply personal work of career exploration and decision-making.

3.7 SUMMARY OF KEY SURVEY FINDINGS

The national survey of 370 Canadian CDPs yields five headline findings that will be carried forward into the cross-cutting synthesis of this report.

- **Firstly,** AI adoption is already widespread, but preparedness has not kept pace. Eighty-four per cent of practitioners are using AI in their work with clients, yet 82% have received no formal training. This gap between adoption and preparation is the defining professional development challenge of the current moment.
- **Secondly**, trust is conditional and well-calibrated. Practitioners trust AI for bounded, verifiable tasks - information synthesis, administrative support, initial drafting - and distrust it for tasks requiring nuanced human judgement, contextual sensitivity, and professional accountability. This differentiated trust, grounded in professional expertise and ethical awareness, is an asset to be supported and built upon, not a barrier to be overcome.
- **Thirdly,** ethical concern is high and informed. Ninety-five per cent of practitioners hold concerns about AI ethics, and their concerns - bias, privacy, accuracy, over-reliance, the erosion of human capability - reflect an accurate and sophisticated understanding of the documented limitations of current AI systems.
- **Fourthly,** the human dimensions of practice are regarded as non-negotiable. Practitioners are consistent and clear that AI should not encroach upon the relational, identity-focused, trauma-informed, and deeply individualized work that defines quality career development practice. This professional consensus is a foundation for responsible AI governance in the sector.
- **Finally,** practitioners are asking for leadership. The strong demand for clear guidelines, expert-evaluated tools, sector-specific training, and structured communities of practice reflects a profession that is ready and willing to engage constructively with AI - but that needs adequate professional infrastructure to do so responsibly. This is an urgent call for action from professional bodies, educational institutions, employers, and governments alike.

4. FOCUS GROUP FINDINGS

Three focus groups were conducted between January and March 2026 as the third phase of this national research study. The first session (Focus Group 1) took place as a pre-conference workshop at CANNEXUS 2026 in Ottawa on 25 January 2026, drawing 25 career development practitioners (CDPs) alongside facilitators from the Canadian Career Development Foundation (CCDF) and its associates. Two subsequent virtual sessions were held on 2 March and 6 March 2026, with six and seven participants respectively, again facilitated by CCDF colleagues and associates. Across the three sessions, a total of 38 CDPs participated - practitioners drawn from diverse organizational contexts, geographic settings, and areas of specialisation within the Canadian career development sector.

Each session was structured around four core question areas: the scope of beneficial AI application within career services; the boundaries of human professional expertise that AI should not cross; the professional development infrastructure practitioners require; and the principles they wish to see govern both AI development and public investment in the sector. The discussion guide was held constant across all three groups, enabling thematic comparison while allowing each session to pursue the lines of inquiry that resonated most strongly with its participants.

The findings presented in this section synthesise the contributions of all three sessions thematically. Where the three groups converged, this is noted; where particular sessions generated distinctive emphases or additional depth, these are identified accordingly. Direct quotations are drawn from facilitator notes and are reproduced verbatim. The analysis does not attribute individual contributions to named participants.

4.1 WHERE AI CAN ADD VALUE: PRACTITIONER AND CLIENT BENEFITS

Across all three focus groups, practitioners identified a substantial and varied range of tasks where they believed AI could augment the quality, efficiency, and reach of career services. Their responses clustered into four broad areas of application: administrative burden reduction, information access and synthesis, client-facing support tools, and organizational and workforce functions.

Administrative efficiency

The most emphatic and consistent response across all three sessions concerned the administrative load that CDPs carry - a dimension of practice that participants felt was consuming time and energy that should be directed towards clients. In Focus Group 1, the word ‘ADMIN’ appeared in facilitator notes with double exclamation marks, reflecting the energy with which participants raised the theme. Practitioners across all three sessions identified tasks including organizing email and to-do lists, streamlining case note documentation and client charting, generating expense claims, summarizing programme guidelines (such as Better Jobs Ontario applications), and producing policy summaries. Focus Group 2 extended this to include automatic client follow-up reminders - with AI providing the CDP with a précis of previous session notes, decisions, and agreed actions - as well as generating and refining job postings and producing labour market summaries for routine reference. Focus Group 3 added strategic planning support, organizational work plans, identification of key performance indicators, and the linking of programmes and funding proposals to existing benchmarks.

The consensus across sessions was not that administrative efficiency was merely convenient, but that it was ethically significant: time released from bureaucratic tasks is time available for the relational, human-centred work that defines quality career development practice. As participants in Focus Group 2 noted, AI could also support session recording and the automated generation of case notes - provided, they emphasized, that data security and client consent were robustly guaranteed.

Labour market information and research

A second area of strong consensus concerned access to labour market information (LMI) and occupational data. Participants across all three sessions described the challenge of navigating large, fragmented, and sometimes inaccessible information environments, and saw AI as a means of making this work faster, more comprehensive, and more usable. Focus Group 1 identified LMI research as a priority benefit area, with particular interest in AI that could connect clients to employers and opportunities aligned with their values and goals and generate targeted interview questions for specific roles. Focus Group 2 elaborated on this, describing the value of AI in generating thoughtful, occupation-specific questions that prompt clients to consider whether a particular career direction is genuinely right for them - turning LMI from a passive reference function into an active tool for reflection.

Focus Group 3 offered the most detailed and practically oriented account of LMI applications. Participants described AI's potential to access and curate LMI, provide occupational descriptions, identify occupations with similar requirements, and - critically - serve populations who lack the professional networks through which labour market knowledge is often distributed informally. Participants also flagged the potential for AI to analyze web analytics and frequently-asked-question patterns, enabling organizations to ensure that their services are genuinely responsive to the needs clients are actually presenting. One distinctive contribution from Focus Group 3 concerned the analysis of small, local labour markets - a recognized gap in publicly available data - with participants calling on developers to build tools capable of generating employment trend data at the community level rather than relying solely on jurisdictional or national datasets.

Client-facing support tools

All three groups identified a range of client-facing applications where AI could extend the reach and quality of career services, subject to important caveats discussed in the following sub-section. Resume and cover letter assistance was the most commonly cited application across all sessions, though with a consistent qualification: AI should serve as a starting point - generating an initial structure or identifying transferable skills - rather than producing a finished document. Participants in Focus Group 3 used the phrase 'as a starting point (only!)' explicitly. Interview preparation emerged as another well-supported application, including the use of audio-enabled AI tools (such as ChatGPT's voice mode) to provide clients with a live, simulated interview practice environment.

Beyond these core job-search functions, practitioners identified a broader set of ways AI could support clients who face particular access barriers. Focus Group 3 highlighted the ability of AI to rewrite text at adjusted reading levels, improving accessibility for clients with lower literacy, and noted the potential of translation functions for serving newcomer populations. Participants in this session also identified AI's potential to reduce systemic bias in career exploration - for example, by generating occupation suggestions that extend beyond the client's existing frame of reference and prompt consideration of non-traditional or underexplored options, provided that prompts are engineered carefully.

Practitioners across sessions also described the value of AI in supporting the preparation phase of their own work - using AI before sessions to brainstorm questions, explore relevant LMI, and develop prompts tailored to a particular client's circumstances. Focus Group 2 was particularly articulate on this: AI used in this preparatory mode functions as a thinking partner for the practitioner, not as a replacement for the professional relationship itself.

Organizational and workforce functions

Focus Group 3 generated the widest range of organizational applications, reflecting the seniority and organizational leadership roles of some of its participants. In addition to the strategic planning and benchmarking functions noted above, participants described AI's value in developing visual and marketing materials, formatting workshop content from raw inputs, summarising large volumes of information for different comprehension levels, and supporting newer practitioners in building confidence around session structure and professional questioning techniques. The observation that AI could function as a tool for capacity-building within the workforce - not only for clients but for CDPs themselves, particularly those at earlier stages of their careers - was a distinctive contribution of this session.

4.2 THE BOUNDARIES OF PROFESSIONAL EXPERTISE: WHERE AI MUST NOT GO

If practitioners' enthusiasm for the administrative and informational applications of AI was strong and relatively consistent, their articulation of where AI should not be used was even more emphatic - and equally consistent across all three sessions. A clear and coherent picture emerged of the domains that participants regard as irreducibly human, grounded in a shared understanding of what career development practice actually is.

Identity, values, and career decision-making

The most fundamental boundary identified across all three groups concerned the core of career development work: supporting individuals in understanding who they are, what they value, and what direction they wish to take in their working lives. This was characterized not as a practical or technical function but as identity work - complex, contextual, deeply personal, and profoundly dependent on the quality of the human relationship between practitioner and client.

Focus Group 2 was particularly unambiguous on this point. Participants stated that AI should never be used to make final life decisions and that AI should stay away from career decision-making entirely, describing it as 'identity work that involves many factors, such as values, lived experience, culture, and trauma.' The same group emphasized that the deeper dimensions of self-exploration - what they described as 'excavating values, strengths, skills, and stories' - emerge through honest, human interaction and cannot be replicated by an automated system. AI, they noted, cannot reliably read human cues or understand silences, and it cannot be held accountable in the way a professional can.

> ***"Stay away from deeper inner work - excavating values, strengths, skills, and stories. These emerge through honest, human interaction."***

Focus Group 3 raised a complementary but distinct concern: that over-reliance on AI for tasks that jobseekers themselves need to perform - problem-solving, analysis, synthesis, and self-articulation - risks undermining the

very capabilities those clients will need to demonstrate once employed. The question participants posed was a serious one: if AI does the cognitive and reflective work for clients, are we inadvertently preparing them poorly for the realities of the workplace?

Mental health, crisis, and complex need

All three groups were unequivocal that AI should not be used in situations involving mental health presentations, emotional crisis, or the complex intersectional challenges that many clients bring to career services. Focus Group 1 flagged over-reliance on AI in mental health contexts as a specific risk, calling for purpose-built tools with clearly defined scope limitations. Focus Group 2 stated explicitly that AI should not be used in 'crisis situations, mental health issues, or situations requiring quick decisions.' Focus Group 3 added the specific caution against using AI with clients presenting with severe mental illness, particularly paranoia, where the nature of the condition makes the non-human quality of AI interaction potentially harmful.

Underlying these specific cautions was a broader point about the nature of the client population that CDPs serve. As participants in Focus Group 3 articulated when addressing their message to governments and funders: clients frequently arrive carrying multiple, complex, and intersectional challenges - including trauma, low trust in systems, isolation, and hopelessness. In this context, the relational quality of the CDP-client encounter is not incidental to the service; it is the service.

> ***"Clients are often coming with multiple, complex and intersectional challenges, including trauma. Trust in systems is often low, and isolation and hopelessness is high. AI strips away the sense that anyone cares about the clients being served."***

Personalized action planning and accountability

A further boundary that emerged with considerable clarity across sessions concerned the development of individual action plans and follow-up strategies. Focus Group 2 argued that these should never be delegated to AI, on the grounds that each plan must be genuinely tailored to the unique strengths, constraints, and circumstances of the individual client - a level of personalization that requires professional judgement, ongoing relational knowledge of the client, and the kind of accountability that only a human professional can provide. Focus Group 1 expressed concern that AI should avoid what participants called 'the human connection' - that the warmth, presence, and relational depth of human contact should be preserved and deepened, not replaced.

Participants in Focus Group 2 also articulated a broader methodological principle: when AI is used primarily for speed and ease - 'a smooth ride' - the result is likely to be shallow. The depth and complexity that characterize quality career development practice cannot be compressed into efficiency without something important being lost.

The risk of over-reliance and misuse

A theme that cut across both the 'where AI can help' and 'where it should not go' discussions was the risk of over-reliance - both by practitioners and by clients. Focus Group 1 noted explicitly the need to 'avoid over reliance' and called for purpose-built tools with clearly articulated scope limitations, rather than general-purpose AI systems applied indiscriminately to career development contexts. Focus Group 3 reinforced this with the observation that practitioners should not default to AI-assisted interventions simply because the tool is available: 'the client may not need a resume' was offered as a representative example of the danger of solution-first rather than needs-first thinking.

On the client side, participants expressed concern about the tendency of AI-generated documents - resumes in particular - to misrepresent the client's genuine voice, cultural identity, and authentic self-expression. Focus Group 2 noted that AI-produced resume content 'is often not a true reflection of the client's capacity, culture, or way of expressing themselves,' and that the human CDP must retain oversight to ensure that any AI-assisted output is genuinely aligned with the client's real strengths and circumstances.

4.3 PROFESSIONAL DEVELOPMENT: WHAT PRACTITIONERS NEED AND HOW THEY WANT IT

All three focus groups devoted considerable attention to the professional development landscape, both mapping the training resources that participants had found useful and identifying the significant gaps that remain. The picture that emerged was of a sector in which generic AI training is abundant but sector-specific, practically applicable professional development is scarce - and where the formats currently available do not always suit either the learning needs of CDPs or the practical realities of their working lives.

Existing resources

Focus Group 1 produced the most extensive catalogue of existing training resources that participants had encountered and found relevant. These included Alpha+ (an Ontario-based practitioner training programme with a digital literacy component), Nations First (a British Columbia initiative designed for Indigenous contexts), a Canadian AI stewardship practice programme, and the US-based Centre for Humane Technology. Participants also referenced Actionsight.ai (leadership-focused AI training), One Life Tools (which integrates reflection tools with AI), RAISE TM, the CAIDP AI Policy Clinic, TechnologyHelps and Cyber-AI, and GoodBot/Braintranr (which applies systems thinking to AI). Focus Groups 2 and 3 pointed to conferences and webinars - including resources from LMIC and TMU's Future of Work series - as well as Angela Hoyt's courses, Keon micro-credentials in AI from Newfoundland and Labrador, and Harvard's 'AI for Everyone' on Coursera. The breadth of this list reflects both genuine practitioner engagement with available resources and the absence of a recognized, sector-authoritative training offer.

Training gaps

Despite the volume of available material, participants across all three sessions were clear that the existing offer does not meet their needs. The central complaint - articulated most sharply in Focus Group 2 - was that generic AI training is 'overwhelming and not readily applicable.' What practitioners want is training that is specifically focused on the use of AI in career services: contextualized, sector-grounded, and directly translatable to their daily practice.

More specifically, participants identified a cluster of interconnected training needs. Prompt engineering - understanding how to construct, refine, and critically evaluate AI prompts in the specific context of career development practice - was identified across all three sessions as a significant and largely unmet need. Focus Group 3 extended this to include an understanding of how particular word choices within a prompt can produce materially different outputs, and the cultivation of what participants described as an 'AI mindset' - approaching AI as a sophisticated personal assistant rather than a search engine. Critically, participants emphasized that training in prompt engineering should include an honest account of AI's limitations: Focus

Group 3 suggested that practitioners need to understand that most AI responses will be 'wrongish' - a reasonable starting point, but never a finished product.

Beyond prompt engineering, participants identified gaps in training on: the ethical use of AI in practice; how AI affects and should affect the scope of professional practice; how to talk about AI with clients in ways that are accurate, reassuring, and empowering rather than confusing; what over-reliance actually looks like in practice; and how to integrate AI into existing service delivery models without undermining the relational quality of the work. Focus Group 1 also raised the need for leadership-specific AI training that goes beyond information security - addressing the organizational, cultural, and professional implications of AI adoption for team leaders and service managers.

Preferred formats

On the question of training format, there was broad consensus on certain principles and some productive differences of emphasis. All three groups agreed that training must be hands-on and practically grounded: abstract or conceptual instruction about AI is insufficient if it is not accompanied by direct experience of using tools in career-specific scenarios. Focus Group 3 was particularly clear on this, calling for training delivered with computers open, navigating AI platforms together through realistic career development scenarios, with prompt 'cheat sheets' and concrete resources that practitioners can apply immediately.

Focus Group 3 also raised the importance of designing training experiences in which early success is built in - ensuring that practitioners develop confidence and competence before encountering AI's inevitable failures, so that they are not discouraged at the point of first difficulty. Acknowledging the strong feelings that practitioners hold about AI - including anxiety, skepticism, and in some cases resistance - was identified as an important part of effective training design.

Focus Group 1 emphasized that training must be role-specific and content-specific: a one-size-fits-all approach fails to account for the range of settings, client groups, and professional roles within the career development sector. The group also called for ongoing learning structures - post-training follow-up, implementation support, and communities of practice - rather than one-off events that provide information without embedding capability.

This theme of ongoing, evolving community-based learning was developed most fully in Focus Group 2, whose participants articulated a vision of a 'living handbook, blog, newsletter, case studies, or community of practice - something ideally interactive, but minimally organic and able to grow and evolve as the use of AI in our sector grows and evolves.' The image of a living resource - one that can respond to new developments, new questions, and new concerns as they arise - stands in deliberate contrast to the static, time-limited training event. Asynchronous formats were valued for their accessibility, but participants emphasized that these should be complemented by opportunities for role-play, small-group interaction, and community engagement.

4.4 MESSAGES TO DEVELOPERS, GOVERNMENTS, AND FUNDERS

The final section of each focus group session invited participants to address directly the two groups of external actors whose decisions most shape the AI landscape in which CDPs operate: the technology developers who build the tools, and the governments and funders who commission and resource them. The responses

were among the most sustained and forceful of the three sessions, reflecting the depth of frustration that practitioners feel about being excluded from decisions that directly affect their work and their clients.

Messages to developers

The central message to developers was consistent and clear across all three groups: practitioners are not passive end-users of AI solutions developed elsewhere, and they do not wish to be treated as such. The experience of having 'solutions' created for the career development sector without meaningful consultation or collaboration with those who actually do the work was a source of significant professional frustration. Focus Group 1 articulated this with particular clarity, calling for end users to be involved from the earliest stages of design - not as token consultants at the point of testing, but as co-designers and co-learners from inception. The group proposed that developers should be required to demonstrate a basic level of competency in career development theory and practice before being considered a trusted source for the sector, and floated the idea of a voluntary certification framework for AI developers working in career services, backed by adherence to a code of ethics and endorsed by a credible third-party body.

Focus Group 2 offered a detailed account of what sector-grounded development would look like in practice. Participants articulated the need for tools that understand 'the complex, intersectional nature of practice' - that issues of readiness, values, trauma, and people's stories require professional discernment that cannot be automated. They named the tendency of AI systems to agree with and validate users rather than offer the challenging alternative perspectives that are central to good career guidance. They called on developers to recognize that 'meaning making cannot be automated - career development is interpretive, narrative-based, and grounded in co-construction.' And they asked developers to take the principle of 'do no harm' seriously as a design constraint, not merely as a rhetorical commitment.

> ***"Stop trying to come up with AI solutions that are intended to replace the roles that are vitally human and depend on professional judgement and expertise."***

All three groups articulated a consistent, positive vision of what a good AI tool for the sector would look like: a secure, sector-specific AI assistant - variously described as a 'tailored ChatGPT' or 'practice companion' - built on and informed by the career development profession's own competency frameworks, theoretical foundations, code of ethics, training materials, and service delivery models. Such a tool would support CDPs in their preparatory and administrative work, provide reliable LMI access, and generate career-development-specific prompts and resources. It would not attempt to replace the professional relationship, make decisions for clients, or operate beyond its defined scope. Focus Group 3 also highlighted the particular value of AI that can access proprietary or hard-to-find data - noting that tools which simply reorganize information already available in the public domain add little value - and called for better AI capability in relation to small, local labour markets, where data gaps are currently most acute.

Focus Group 1 raised a structural concern that is worth highlighting independently: that much AI development in career services is driven by institutional directives to reduce costs, rather than by a genuine desire to improve the quality and reach of services. Participants questioned whether some AI 'solutions' exist because they should exist - because they genuinely meet an identified need - or simply because the technology makes them possible. This distinction between technology-led and needs-led development was identified as fundamental to responsible innovation in the sector.

Messages to governments and funders

The messages directed at governments and funders were, if anything, even more direct. Across all three groups, a consistent set of concerns emerged about the way in which public investment in AI is currently being framed and directed, and a clear set of priorities for how it should be reoriented.

The most fundamental concern - articulated with particular force in Focus Group 3 - was that AI is already generating a 'false sense of capacity' in the career development system: that funders are accepting diluted, AI-driven service provision as an adequate substitute for competent, human-centred professional practice, with consequent harm to clients and degradation of service quality and outcomes. Participants called on governments to resist this trend - to be 'more critical and discerning' in their assessment of AI-powered services, and to insist on evidence of outcomes rather than being 'blinded by this shiny new object.'

> ***"AI is building a false sense of capacity and a general degradation in the quality and outcomes of service provision. Funders are accepting diluted and sometimes harmful AI-driven approaches as 'good enough' proxies for quality, professional services using evidence-based approaches."***

A closely related concern, raised in Focus Group 2, was the tendency of some policymakers and technology specialists to develop AI solutions for career services without the involvement of CDPs - treating the profession as an object of policy rather than a partner in it. Participants called for practitioners to be recognized as vital partners in procurement, development, and adoption decisions; for career development expertise to inform policy before decisions are made rather than after; and for funders to invest in the professional development of CDPs themselves, rather than directing resources primarily towards technological infrastructure.

On the question of investment priorities, all three groups were in accord: training - sustained, funded, sector-specific, and ongoing - should be the primary focus of government investment in AI-readiness for the workforce development sector. One-off training events are insufficient; practitioners need continuous professional development that evolves as the technology and the policy environment evolve. Equity must be central to any investment framework: this means attending to accessibility, addressing the significant differences between urban and rural realities, respecting Indigenous perspectives, and ensuring that AI tools and the training to use them are available across the full diversity of the sector.

Participants were equally unequivocal on what they did not want to see. Focus Group 2 called on governments not to use AI as a mechanism for reducing the labour and salary costs of service delivery: AI can enhance practice, but it cannot replace the qualified human professionals through whom quality career development is delivered. Focus Group 1 called for policy frameworks that move beyond 'one-for-all solutions,' recognizing that different client populations, geographic contexts, and service models require different AI applications and different governance frameworks. Focus Group 3 noted the failure of common AI-driven intake and assessment tools that have 'miserably failed' in the sector - 'notoriously wrong and doing significant harm' - as evidence of what happens when technology is deployed without adequate professional oversight, co-design, or evidence of effectiveness.

Finally, participants across sessions called for robust national policy frameworks governing how and when AI may and may not be used in career services - frameworks developed with the sector, not for it, and backed by the kind of oversight mechanisms that the public interest demands when significant life decisions are at stake.

4.5 SUMMARY OF KEY FOCUS GROUP FINDINGS

Across the three focus group sessions, a coherent and strongly held picture emerged of how Canadian CDPs understand the role of AI in their sector. Five overarching findings are drawn from this analysis.

- **Firstly,** practitioners are neither AI-phobic nor uncritically enthusiastic. They approach AI as a pragmatic professional community - identifying specific, bounded applications where AI adds genuine value while articulating, with equal clarity and force, the professional and ethical boundaries that must be maintained. Their engagement with the technology is conditional, contextual, and grounded in an understanding of what quality career development practice actually requires.
- **Secondly,** the greatest and most unambiguous benefit of AI in the sector lies in administrative efficiency and information access - tasks that consume professional time without adding relational value. This is also the area of lowest risk, and the one where investment in AI tools is most clearly justified on current evidence.
- **Thirdly,** the core of career development practice - identity work, values clarification, career decision-making, crisis response, and the cultivation of the working alliance - is regarded as irreducibly human. Practitioners articulated this boundary not as a defensive reaction to technology but as a positive account of what their profession offers and why it matters. AI that encroaches on this territory is seen as harmful, not merely unhelpful.
- **Fourthly,** practitioners are under-resourced for responsible AI adoption. The professional development infrastructure needed to support confident, ethical, and effective AI integration does not currently exist in sector-specific form. Generic training is abundant but inadequate; sector-specific, practically grounded, and ongoing development is urgently needed and must be funded.
- **Finally,** the sector's relationship with both developers and funders is characterized by a significant power imbalance that practitioners wish to see redressed. CDPs want to be recognized as co-designers of AI solutions, as partners in policy development, and as the authoritative voice on what responsible AI integration in career services looks like - not as recipients of solutions designed by others on their behal

5. CROSS-CUTTING SYNTHESIS

This section draws together the findings of the three research phases reported in Sections 2, 3, and 4 - the international literature review, the national survey of 370 Canadian CDPs, and the three focus group sessions with 38 practitioners - to identify the themes that cut across all three sources of evidence. Where the Canadian practitioner data confirms what the international literature would lead us to expect, this is noted. Where it extends, nuances, or challenges the picture the literature presents, this is highlighted. Five cross-cutting findings emerge from this synthesis, each of which informs the recommendations in Section 6.

5.1 ADOPTION HAS OUTPACED PREPARATION - AND THE GAP IS A PROFESSIONAL RISK, NOT JUST A TRAINING DEFICIT

Across all three evidence sources, the most consistent and urgent finding is the gulf between the pace of AI adoption and the level of professional preparation underpinning it. The international literature documents this pattern across multiple national contexts. The Canadian survey makes it quantitatively precise: 84% of practitioners are already using AI in their work with clients, yet 82% have received no formal training. The focus groups give it texture and weight, with practitioners describing the experience of navigating a rapidly changing technological landscape largely through self-directed learning, peer exchange, and trial and error.

This gap matters because it is not simply a knowledge deficit that more training will resolve. It reflects a structural failure to invest in the professional infrastructure - guidelines, governance frameworks, evaluation standards, communities of practice, accredited training programmes - that responsible AI adoption in a human services profession requires. The international literature is clear on this point: Accenture's finding that organizations typically spend three times more on technology than on people in AI transformation describes a pattern that is replicated, in different form, when career development services adopt AI tools without commensurate investment in professional development and practice standards.

The survey data add a further dimension: the gap between adoption and preparation is not evenly distributed. Practitioners who are already uncomfortable with AI - 47% of the survey sample - are using tools they do not fully trust, in contexts they are not sure are appropriate, without the guidance they need to make confident professional judgements. This is a condition that exposes both practitioners and clients to risk. The focus groups were explicit about what is at stake: when AI is used without adequate professional preparation, the quality and safety of services to clients - many of whom are navigating complex, high-stakes life transitions - cannot be assured.

5.2 THE HUMAN CORE OF CAREER DEVELOPMENT PRACTICE IS NON-NEGOTIABLE - AND CANADIAN PRACTITIONERS CAN ARTICULATE PRECISELY WHY

The international literature identifies the importance of human oversight in AI-mediated career services, but does so at a relatively abstract level - emphasizing the need for 'humans in the loop' and the limitation of AI in contexts requiring nuanced judgement. The Canadian practitioner data in this study go considerably

further, offering a detailed, professionally grounded, and philosophically coherent account of why the human dimensions of career development cannot be automated.

Across the survey and all three focus group sessions, practitioners consistently and emphatically identified the same set of functions as irreducibly human: supporting clients in self-awareness and values clarification; career decision-making as identity work involving lived experience, culture, and trauma; crisis response and mental health support; the development of genuine accountability between practitioner and client; and the relational quality - the warmth, presence, attentiveness, and professional empathy - that makes a career guidance encounter meaningful rather than merely informational.

> ***"The work that we do as CDPs is very personal work, individualized, and specific to each individual; skills specific and story specific; and it requires a human touch, empathy, compassion, trauma-informed practice - you cannot get that from AI tools."***

This is not resistance to technology. The same practitioners who drew these boundaries with clarity also identified substantive and enthusiastic uses for AI: administrative efficiency, labour market information synthesis, client preparation support, accessibility and translation functions, and the preparatory dimensions of professional practice. The picture is not of a profession divided between technophiles and technophobes, but of professionals applying the kind of contextual, proportionate judgement that the responsible use of any powerful tool requires. The literature's framing of human-AI collaboration as the emerging model is confirmed by the Canadian data - but Canadian practitioners articulate it more precisely, more professionally, and with more ethical depth than much of the published literature.

One distinctive contribution of the Canadian data concerns what the focus groups termed the competency displacement paradox: if AI performs the analytical and expressive tasks - drafting the cover letter, synthesizing the self-assessment, structuring the interview response - clients may be deprived of the opportunity to develop precisely the capabilities that employment will require of them. This concern, which does not feature prominently in the international literature, reflects a sophisticated professional understanding of career development as a developmental process, not merely an instrumental one.

5.3 ETHICAL CONCERN IS HIGH, WELL-FOUNDED, AND AHEAD OF GOVERNANCE - AT EVERY LEVEL

The international literature documents the inadequacy of existing ethical and governance frameworks for AI in career development. The Canadian survey and focus group data show that practitioners are not waiting for those frameworks to arrive: 95% of survey respondents hold ethical concerns about AI in career development, and their concerns - about bias, privacy, accuracy, over-reliance, the erosion of human capability, and the risk of harm to vulnerable clients - reflect an informed, sophisticated, and professionally grounded understanding of the issues at stake.

The survey finding that only 7% of practitioners believe AI systems used in career development are currently fair and unbiased - with 56% unsure and 37% actively disagreeing. It reflects not ignorance but awareness: practitioners with direct experience of AI-generated outputs that misrepresent clients, reproduce existing biases, or produce plausible-sounding but inaccurate information are well-positioned to assess the risks. Their skepticism is an asset to be channelled, not a barrier to be overcome.

The focus groups gave voice to concerns that go beyond the individual practice level. Practitioners directed sharp and specific criticism at the broader system within which AI is being deployed in career services: at funders who accept diluted, AI-driven service provision as an adequate substitute for competent professional practice; at developers who create tools for the sector without meaningful consultation with those who work in it; and at governments that have not yet developed the regulatory and oversight frameworks that the high-stakes nature of career guidance - where significant life decisions are at stake - requires. These systemic concerns align closely with the international literature's account of ethical governance failures, but they are grounded in concrete Canadian experience rather than abstract principle.

> ***"Funders are accepting diluted and sometimes harmful AI-driven approaches as good enough proxies for quality, professional services using evidence-based approaches."***

The survey's finding that 86% of practitioners consider expert-led evaluation of AI tools to be extremely or quite important, and that 56% regard clear professional guidelines as extremely important, reflects a profession that is not asking to avoid AI but asking for the conditions - standards, oversight, expertise-grounded evaluation - under which it can be used responsibly. This is a constructive professional demand that policy and professional bodies have an obligation to meet.

5.4 THE PROFESSION IS ASKING TO BE A PARTNER, NOT A RECIPIENT - IN TOOL DEVELOPMENT, POLICY, AND GOVERNANCE

One of the most consistent and forceful themes across the Canadian practitioner data - present in the survey commentary and developed at length in all three focus group sessions - is the demand for genuine professional participation in decisions about AI in career services. This demand has three distinct but related dimensions: co-design of AI tools, involvement in policy and procurement decisions, and authority over professional standards and guidelines.

On tool development, practitioners across all three focus groups described the frustration of having AI 'solutions' created for their sector without meaningful consultation with those who do the work. They articulated a clear and positive vision of what co-designed tools would look like: sector-specific, grounded in the profession's own competency frameworks and theoretical foundations, code of ethics and service delivery models, and built with the understanding that career development is relational, interpretive, narrative-based, and contextually complex in ways that off-the-shelf AI cannot capture. The demand is not for bespoke technology for its own sake, but for tools that reflect professional knowledge rather than overriding it.

The international literature largely treats this as an implementation challenge - how to involve end-users in AI development. The Canadian practitioners frame it as a professional rights question. They are subject matter experts whose knowledge is essential to responsible innovation in the sector, and they wish to be recognized as such. Focus Group 1 proposed concrete mechanisms: voluntary certification for AI developers working in career services, demonstrated competency in career development as a condition of trusted developer status, and third-party endorsement frameworks. These are sophisticated governance proposals that go well beyond the generic calls for 'stakeholder consultation' that appear in much of the policy literature.

On policy and procurement, the message was equally direct: governments and funders should not commission, procure, or mandate AI tools for career services without the active involvement of CDPs in the design and evaluation process. Practitioners are not positioned as one stakeholder among many; they are identified as the authoritative professional voice on what responsible AI integration in career development looks like, and they expect to be treated as such. The survey finding that 78% consider expert evaluation of AI tools extremely or quite important before those tools are recommended for use reflects the same professional expectation, quantitatively expressed.

5.5: EQUITY IS NOT A DIMENSION OF AI IN CAREER DEVELOPMENT - IT IS THE CENTRAL TEST OF WHETHER AI IS BEING USED RESPONSIBLY

The international literature identifies equity and inclusion as important considerations in AI deployment - documenting geographic disparities in AI access and optimism, the gendered and socioeconomic distribution of automation risk, and the potential for algorithmic bias to reproduce structural inequalities. The Canadian practitioner data treat equity not as one consideration among many but as the fundamental test of responsible AI use in career services.

This framing reflects the nature of the work. Career development practitioners, particularly those working in publicly funded employment services, non-profit organizations, and community settings, disproportionately serve the populations most exposed to the risks that irresponsible AI creates: newcomers navigating unfamiliar systems in a second or third language; Indigenous peoples whose cultural contexts and ways of knowing are poorly represented in AI training data; racialised communities facing structural barriers that AI systems may reproduce or amplify; people with disabilities who may be both most helped by accessible AI tools and most harmed by discriminatory screening algorithms; and lower-wage workers in the occupations most exposed to automation, who are simultaneously the clients most in need of intensive career support and the least resourced to access it.

Across the survey and focus groups, practitioners articulated equity concerns at multiple levels. At the client level: that AI-generated documents - resumes, cover letters - may misrepresent clients whose cultural identity, language, and self-expression are poorly served by AI trained predominantly on majority-culture data. At the service level: that AI-driven career services may create a two-tier system in which those with digital literacy, reliable internet access, and the confidence to engage with new technologies receive AI-augmented support, while those without these resources receive diminished human services. At the system level: that the pressure to deploy AI as a cost-reduction measure - reducing the headcount of qualified CDPs - will damage service quality most severely for the clients who most need intensive human support.

The focus group contribution on this point was particularly sharp. Practitioners working with clients carrying complex, intersectional challenges - trauma, low trust in systems, isolation, and hopelessness - described AI as stripping away the sense that anyone cares about the people being served. This is not a generic concern about technology replacing human contact; it is a specific professional judgement about the populations served, the nature of the working alliance required to serve them effectively, and the harm that would result from its removal. It deserves to be treated with the weight it carries.

The survey's demographic data add a further dimension. Cukier et al.'s finding that immigrants, racialized, and Indigenous respondents demonstrated higher AI familiarity than white and Canadian-born participants challenges deficit framings of equity-deserving communities in relation to AI. The readiness gap is not located where it is often assumed to be. The more pressing equity question is whether the governance, procurement, and deployment decisions being made about AI in career services are adequately centred on the needs and rights of the populations career development exists to serve.

5.6 IMPLICATIONS FOR POLICY AND PRACTICE

Taken together, these five cross-cutting findings yield a clear and consistent set of implications. The readiness gap is structural and requires structural solutions: funded, sustained, sector-specific professional development; accredited training embedded in CDP certification programmes; and national professional guidelines developed with the sector rather than for it. The boundaries of human professional expertise are well-understood by practitioners and should be codified in governance frameworks that protect those boundaries rather than eroding them in the name of efficiency. Ethical governance frameworks are urgently needed and must be developed with the active participation of CDPs, drawing on their professional knowledge of the specific risks and contexts involved. Practitioners must be recognized as co-designers of AI tools and co-authors of AI policy, not recipients of solutions developed without them. And equity must be the central test of responsible AI deployment in career services - applied rigorously, at the level of individual tools, service models, and national investment decisions.

These implications are developed into specific, audience-targeted recommendations in **Section 6.**

6. RECOMMENDATIONS

The recommendations presented in this section are grounded in the five cross-cutting findings of Section 5 and reflect the integrated evidence of the international literature review, the national survey of 370 Canadian career development practitioners, and the three focus group sessions. They are addressed to five distinct audiences whose decisions and investments will most shape the trajectory of AI in Canadian career development services: professional bodies and training providers; technology developers; employers and service providers; governments and funders; and career development practitioners themselves. Across all five audiences, the recommendations are underpinned by a single organizing principle: AI in career development should amplify professional expertise and extend equitable access to high-quality guidance - not be a substitute for qualified human practice, reduce investment in the profession, or serve as a cost-reduction mechanism at the expense of the people career development systems exist to serve.

6.1 RECOMMENDATIONS FOR PROFESSIONAL BODIES AND TRAINING PROVIDERS

Professional bodies and training providers bear primary responsibility for the infrastructure that enables responsible AI adoption across the career development workforce. The research findings are unambiguous about the scale and urgency of the gap that currently exists between the level of AI adoption and the level of professional preparation supporting it.

Develop sector-specific professional guidelines for AI use

The 86% of survey respondents who regard clear guidelines on AI use in career development as extremely or quite important are making a professional demand that national bodies have an obligation to meet. Guidelines must go beyond generic AI ethics principles and address the specific dilemmas that CDPs face in daily practice: when AI use is and is not appropriate in client interactions; how to handle, store, and disclose data generated through AI-assisted sessions; how to evaluate the reliability of AI-generated labour market information; how to support clients who are using AI without professional guidance; and how to maintain professional accountability in an AI-augmented practice environment. Guidelines should be developed with, rather than for, the practitioner community, drawing on the professional knowledge that CDPs have already demonstrated throughout this research.

Integrate AI literacy into CDP certification as a core competency

AI literacy - encompassing both the technical competencies required to use AI tools effectively and the ethical, critical, and professional competencies required to use them responsibly - should be embedded in the core competency frameworks that govern CDP certification and accreditation. This is not an optional enhancement but a baseline professional requirement, given that 84% of practitioners are already using AI with clients. Initial integration should be followed by a commitment to ongoing review, ensuring that competency frameworks keep pace with the rapid evolution of the technology. As one survey respondent put it directly: *'this should be an entire course in future CDP programmes.'*

Build a funded, sustained, and sector-specific professional development infrastructure

The most consistent message from both the survey and the focus groups on professional development is the inadequacy of generic AI training and the urgent need for resources that are sector-specific, practically grounded, and designed to evolve as the technology and practice context evolve. Professional bodies should invest in: accredited training programmes covering prompt engineering in career development contexts, critical evaluation of AI outputs, ethical decision-making frameworks, and AI use with specific equity-deserving populations; a national community of practice enabling peer learning and real-time knowledge exchange; and a living resource - whether a handbook, a regularly updated portal, or a curated content hub - that can respond to new developments as they arise. One-off training events are insufficient; the infrastructure required is ongoing, adaptive, and professionally authoritative.

Establish a voluntary certification framework for AI developers in career services

Drawing on the specific proposal developed in Focus Group 1, professional bodies should explore the development of a voluntary certification framework for AI developers working in the career development sector. Certification should require demonstrated competency in career development theory, practice, and ethics; adherence to a sector-specific code of conduct; and endorsement by a credible third-party body. This framework would create a mechanism for the profession to exercise quality assurance over the tools being developed in its name and would provide practitioners with a basis for informed procurement decisions.

6.2 RECOMMENDATIONS FOR TECHNOLOGY DEVELOPERS

Technology developers working in the career development space are asked, by practitioners across this study, to do something that requires both professional humility and genuine commitment: to understand career development before building solutions for it, and to treat CDPs as co-designers rather than passive recipients.

Co-design with CDPs from inception, not consultation at testing

The research findings document widespread practitioner frustration with AI tools developed for the career development sector without meaningful professional input. Co-design is not a token gesture of stakeholder engagement at the point of user-testing; it is the sustained, iterative involvement of practitioners with subject matter expertise and lived experience of the work from the earliest stages of conception through to evaluation and ongoing improvement. Developers should build teams that include career development professionals, not only as advisers but as co-architects of the tools being created.

Build for the complexity of career development, not the simplicity of skills matching

The most important message from practitioners to developers in this study is that career development is relational, contextual, narrative-based, and grounded in the co-construction of meaning between practitioner and client. It is not skills matching. It is not resume generation. It involves identity work, values clarification, trauma-informed practice, cultural responsiveness, and the kind of accountability that only a human professional relationship can sustain. AI tools that treat career development as a data optimization problem will fail professionally, regardless of their technical sophistication. Developers must understand what career development actually is - its theoretical foundations, its professional competency framework, its ethical code - before they can build tools that genuinely support it.

Prioritize administrative and informational functions; resist replacing relational ones

The clearest finding from both the survey and the focus groups about what practitioners want AI to do is: reduce the administrative burden and improve access to information. Developers should focus their investment in these areas - which have the highest potential benefit, the lowest professional risk, and the strongest practitioner support - before attempting to automate functions that are irreducibly human. Tools that generate case notes from session recordings (with appropriate consent and security), provide labour market summaries, support session preparation, and enable translation and accessibility are valued and needed. Tools that attempt to substitute for the professional relationship, make career decisions on behalf of clients, or provide personalized guidance without professional oversight are not.

Take 'do no harm' seriously as a design constraint

The principle of 'do no harm' was raised in every focus group session and across the survey commentary. For developers, this means: building in explicit safeguards around confidentiality, consent, and data security from the design stage; subjecting AI outputs to rigorous bias testing before deployment, with particular attention to the demographic groups most served by career development practitioners; building in transparency mechanisms so that practitioners and clients can understand the basis for AI-generated information and recommendations; and establishing feedback and continuous improvement mechanisms that allow practitioners to report problems and have them addressed. Nothing should move from pilot to adoption without rigorous vetting by frontline CDPs and, where appropriate, by the client populations the tool is designed to serve.

6.3 RECOMMENDATIONS FOR EMPLOYERS AND SERVICE PROVIDERS

Organizations that employ career development practitioners - whether in post-secondary institutions, publicly funded employment services, non-profit organizations, or the private sector - have both an obligation and an opportunity to shape responsible AI adoption at the service delivery level.

Develop clear internal AI use policies and governance structures

Survey respondents consistently described a governance gap at the organizational level: uncertainty about which tools are approved for professional use, what data handling requirements apply, and what quality assurance processes govern AI-assisted work. Employers should address this gap by developing clear, specific, and regularly reviewed internal AI use policies that specify which tools are approved for staff use and why; data privacy protocols and client consent processes that meet applicable legal and professional standards; quality assurance processes for AI-assisted work, including human review requirements; and channels through which practitioners can raise concerns about AI tools and have them addressed.

Invest in staff training as a professional obligation, not an optional benefit

The expectation that practitioners will self-educate on complex AI ethics, technical capabilities, and professional implications is unrealistic and professionally irresponsible. Organizations should fund and require training for all staff using AI in their work with clients, covering both the technical and ethical dimensions of AI use. This is not a discretionary investment; it is a prerequisite for the quality and safety of services delivered to clients, many of whom are in vulnerable circumstances.

Resist pressure to reduce staffing based on AI efficiency arguments

Both the survey data and the focus group findings include clear warnings about the risk that AI will be used as a justification for reducing the number of qualified CDPs in service delivery. Employers should resist this pressure. AI can make CDPs more effective; it cannot replace the professional relationship, the exercise of professional judgement, or the human presence that quality career development requires. Any reallocation of resources freed by AI efficiency gains should be directed towards improving service quality - more time with clients, better professional development, stronger equity and access outcomes - not towards reducing the qualified workforce.

Engage practitioners in AI procurement decisions

Decisions about which AI tools to adopt, how to deploy them, and how to evaluate their effectiveness should not be made by technology teams or organizational leadership without meaningful input from the CDPs who will use them. Practitioners have the professional knowledge required to assess whether a tool is fit for purpose in career development contexts, and their involvement in procurement decisions is essential to responsible adoption. Organizations should establish clear processes for practitioner input into AI procurement, including structured evaluation periods and formal mechanisms for feedback.

6.4 RECOMMENDATIONS FOR GOVERNMENTS AND FUNDERS

Governments and funders occupy a uniquely influential position in the career development AI landscape: they commission services, fund tools, set policy frameworks, and make investment decisions whose effects cascade throughout the system. The research findings include some of the most direct and urgent practitioner messages for this audience.

Invest in the profession, not just the technology

The most fundamental reorientation required of government and funder investment in AI for career services is a shift in balance from technology to people. The international literature finds that organizations typically spend three times more on technology than on the people and processes required to use it effectively. Public investment decisions should reflect the opposite priority: sustained, funded professional development for CDPs; support for the development of sector-specific training programmes and professional guidelines; and investment in communities of practice and knowledge-sharing infrastructure. Technology investment that is not matched by equivalent investment in professional capacity will not deliver the outcomes that justify it.

Develop national policy and regulatory frameworks for AI in career development

The research findings document a significant gap in national AI governance as it applies to career development services. Governments should develop clear regulatory frameworks that address: privacy requirements for AI tools used in career development contexts; bias detection and mitigation requirements; transparency obligations regarding AI use in publicly funded services; accountability mechanisms when AI-assisted services produce harmful outcomes; and minimum standards for the professional qualifications of those delivering AI-augmented career guidance. These frameworks must be developed in genuine partnership with the career development sector - not as top-down impositions, but as collaboratively developed standards that reflect professional knowledge of what responsible AI integration requires.

Require evidence of outcomes before endorsing or mandating AI tools

Focus group practitioners were unambiguous on this point: governments and funders are accepting AI-driven approaches as adequate substitutes for qualified professional practice without demanding evidence that they deliver comparable outcomes. Public investment in AI tools for career services should be conditional on rigorous evaluation evidence demonstrating effectiveness for the client populations being served - not on cost savings, efficiency arguments, or vendor promises. Where evidence does not yet exist, investment should be directed towards generating it, through properly designed evaluations that track client outcomes over time and across demographic groups.

Centre equity in every investment decision

Every government and funder investment decision about AI in career development should be subject to a rigorous equity impact assessment addressing: whether the tool or service is accessible to all client groups, including those with lower digital literacy, limited internet access, language barriers, or disabilities; whether the AI systems involved have been tested for bias with respect to the demographic groups the service is designed to serve, including newcomers, racialised communities, Indigenous peoples, and people with disabilities; whether the investment will narrow or widen existing gaps in access to high-quality career guidance; and whether the voices of equity-deserving communities have been meaningfully included in the design and evaluation process.

Do not use AI as a mechanism for reducing career development service capacity

The clearest message from practitioners to governments and funders across all three focus group sessions is this: do not use AI as a way to decrease the labour and salary costs of career development service delivery. The clients who most need intensive, human-centred career support are those facing multiple, complex, intersectional challenges - trauma, low trust in systems, language barriers, discrimination. For these clients, the relational quality of the CDP encounter is not a luxury; it is the core of the service. Investment in AI should expand the reach and quality of career development services, not serve as a justification for reducing the qualified human workforce that delivers them.

6.5 RECOMMENDATIONS FOR CAREER DEVELOPMENT PRACTITIONERS

The recommendations above are addressed to the institutions and systems that shape the environment in which CDPs work. But practitioners themselves also have agency in determining how AI is adopted in their sector - and the research findings make clear that many are already exercising it thoughtfully, critically, and with professional integrity.

Engage with AI actively and critically - neither avoid nor defer to it

The evidence from this study suggests that the most effective professional response to AI is neither wholesale adoption nor wholesale avoidance, but engaged, critical, and contextually sensitive use: trying AI tools in low-risk contexts, evaluating their outputs with professional scrutiny, building prompt engineering skills through practice, and developing a personal and team-level understanding of where AI adds genuine value and where it does not. Practitioners who disengage from AI entirely risk being excluded from conversations about its governance and development that will affect their clients regardless. Those who adopt AI uncritically risk exposing clients to harms they have the professional expertise to prevent.

Maintain and articulate the professional boundaries that matter

The research findings demonstrate that Canadian CDPs have a sophisticated, coherent, and professionally grounded understanding of where AI should and should not be used in career development practice. This professional knowledge is a resource - for individual practice decisions, for team and organizational policy development, for engagement with developers and policymakers, and for the broader public conversation about AI in human services. Practitioners should document, share, and advocate for the professional boundaries they have articulated in this research, treating them not as resistance to technology but as evidence of professional expertise.

Seek out and build sector-specific professional development

In the absence of a comprehensive sector-specific training infrastructure - which this research identifies as urgently needed but not yet in place - practitioners have an individual responsibility to seek out the best available resources and to advocate within their organizations and professional networks for the development of better ones. Resources identified by practitioners in this study include sector-specific webinars and conference sessions, communities of practice, and emerging AI literacy programmes with career development relevance. Practitioners should also advocate actively, through their professional associations and employer channels, for funded training, clear guidelines, and the governance infrastructure that responsible AI adoption requires.

Contribute to the knowledge base the sector needs

The international literature review found that the sector-specific, practitioner-grounded evidence base for responsible AI use in career development is thin. CDPs who are developing effective approaches to AI integration - whether in session preparation, administrative efficiency, client preparation for AI-transformed labour markets, or equity-centred AI use with specific populations - have knowledge that the sector needs. Contributing to communities of practice, sharing case studies through professional associations, and supporting research efforts like this one are all ways in which individual practitioners can help build the collective knowledge that will enable the sector to navigate the AI transition well.

6.6 A FINAL NOTE

The career development profession in Canada is navigating the AI transition from a position of considerable professional strength. The practitioners who participated in this research demonstrated sophisticated ethical reasoning, clear professional values, and a grounded understanding of both the opportunities and the risks that AI presents for their work and their clients. They are not asking to be protected from AI; they are asking for the conditions - training, guidelines, governance, co-design, and investment - under which they can use it responsibly in the service of the people they are professionally committed to supporting.

Meeting that ask is not only a professional development challenge. It is a matter of public policy, professional ethics, and social justice. The populations served by Canada's career development system - newcomers, Indigenous peoples, racialized communities, people with disabilities, workers facing automation-driven displacement - have the most to gain from career guidance that uses AI wisely, and the most to lose from career guidance that uses it badly. The recommendations in this paper are designed to help ensure that the former, rather than the latter, is the direction of travel.

SUPPORTED IN PART BY: