mpac Municipal Property Assessment Corporation™

2026 AI Roundtable: Meeting 2

Attendees: *Nicole McNeill, Scott Milne, Julia Lipman, Jeffrey Ma, Tammy Wells-Garrett, Stratos Kaloutas, Mark Daley, Jesslyn Dymond, Warren Ali, Dawn Hall, Chad Cogar, Sam Ip, Simon Hodgett, Arjan Arenja*.

MEETING NOTES:

AI Trust Atlas and Public Perception Research: Jesslyn Dymond presented TELUS's AI Trust Atlas, providing a comprehensive overview of the company's approach to AI governance, its commitment to participatory inclusion of underrepresented groups, and key findings from their annual report examining public perceptions of artificial intelligence. The presentation prompted questions from attendees regarding research methodology and emerging trends in public sentiment toward AI technologies.

Participatory AI Governance

- TELUS's AI governance program is intentionally built around participatory principles
- The organization has made a strategic decision to include underrepresented communities in their research efforts, ensuring that diverse perspectives actively shape AI strategy, guiding principles, and operational practices
- This participatory approach reflects a broader commitment to developing AI governance frameworks that account for the needs and concerns of populations that might otherwise be overlooked in technological development processes

Key Findings from the AI Trust Atlas

- Trust remains a critical factor influencing AI adoption
- The majority of survey respondents expressed a clear preference for companies to conduct thorough reviews of AI systems for potential harm before releasing new tools to the public
- Privacy and security concerns emerged as particularly prominent barriers to everyday AI use
- These concerns were especially pronounced among youth populations and Indigenous peoples
- These findings underscore the importance of addressing privacy considerations and building trust mechanisms as prerequisites for broader AI adoption

Methodology and Inclusion

- TELUS partners with external organizations as well as internal employee resource groups to ensure research samples are both diverse and representative of Canadian society

- The organization intentionally seeks out voices from Indigenous communities, LGBTQ individuals, and other underrepresented groups
- This approach ensures the research captures perspectives that might otherwise be absent from mainstream technology surveys

Trends and Literacy

- Overall concerns about AI have remained relatively steady over time, while adoption rates continue to increase
- An interesting finding emerged regarding AI literacy that groups with less access to technology often demonstrate higher levels of AI literacy and greater awareness of potential risks associated with these technologies
- This counterintuitive finding has prompted further academic research to better understand the underlying factors driving these patterns

Certification and Compliance Efforts

- TELUS is actively pursuing ISO 42001 certification for its AI management systems
- The organization is also working to align its practices with the requirements of the EU AI Act
- Recent organizational changes, including the reintegration of their digital business unit, have broadened the scope of AI systems covered under their governance framework and necessitated expanded compliance efforts

G7 AI Grand Challenge and Scaling Roadmap: Dawn Hall provided the group with a comprehensive update on the G7 AI Grand Challenge, detailing participation statistics, lessons learned from the initiative, and the recent publication of a roadmap designed to support the scaling of AI solutions across government. Attendees raised follow-up questions regarding public availability of resources and the challenges associated with cross-country collaboration.

Grand Challenge Overview

- The initiative received a total of 152 proposals, with the majority originating from Canada
- The review process resulted in:
  - Seven recipients
  - Eleven honourable mentions
  - Three finalists
  - Projects selected address a range of public sector challenges, including:
  - Improvements to public sector communications
  - Optimization of resource allocation processes
  - Enhanced information management capabilities

Lessons Learned and Participation Challenges

- Broader participation may have been limited by:
  - The relatively short two-week submission window, which created time constraints for potential applicants
  - Limited advertising of the opportunity in all G7 languages, which may have reduced awareness among non-English-speaking participants
  - Time zone differences complicated coordination and submission logistics, particularly for countries outside Canada

Challenges adopting solutions across jurisdictions

- Data availability to support deployment of the solutions in different jurisdictions can be an obstacle.
- This is in part due to different data collection and access practices, and different policies, laws and compliance approaches across jurisdictions.
- Dawn noted that these challenges are common across governments, suggesting that addressing these challenges will require coordinated effort

Feedback and Knowledge Sharing

- A feedback mechanism was established to support participants in the Grand Challenge
- Feedback was provided to youth participants, finalists, honourable mentions, recipients, and any other submitters who requested it
- This feedback loop was designed to ensure that even all proposals could benefit from the evaluation process and improve future submissions

Scaling Roadmap Publication

- Dawn described the newly published roadmap for scaling AI solutions, designed to facilitate the reuse and expansion of successful AI implementations across government
- The roadmap is structured around five distinct lifecycle stages that AI solutions can progress through
- It is further organized according to four supporting pillars:
  - Design and engagement
  - Governance and responsible AI
  - Capacity building
  - Measurement and continuous improvement
- This framework is intended to provide government agencies with practical guidance for adopting and expanding AI solutions that have proven effective in initial deployments

AI Strategy, Moonshots, and Organizational Change: The group then moved on and engaged in a substantive discussion regarding the meaning of AI strategy, the concept of "moonshots" as proposed by Ajay Agrawal, and the challenges of coalition building, public acceptance, and systemic change in both government and private sectors. The discussion was informed by Ajay Agrawal's submission paper to Canada's AI Strategy Task Force.

Ajay Agrawal's Moonshot Framework

- The paper advances the central argument that setting a target to become the "best AI country in the world" is meaningless without specifying measurable outcomes
- AI is fundamentally a tool—computational statistics that performs prediction—and successful adoption should be defined not by the quantity of AI use but by applying that tool to problems with measurable results
- The pertinent questions are not whether Canada leads in AI per se, but whether Canadians receive faster healthcare, better education, and stronger defense
- The paper proposes addressing this challenge through moonshots: ambitious, measurable, high-impact goals that force system-wide change
- A single bold target—such as cutting healthcare wait times by 90 percent—requires redesigning entire systems with the latest technology to reduce redundancy and increase productivity
- Governments are uniquely positioned to lead these efforts because they can coordinate across agencies, provinces, and sectors in ways the private sector cannot

Five Proposed Moonshots

The moonshots share specific characteristics: they are ambitious, benefit many Canadians, require AI to succeed, build infrastructure other projects can use, and cannot be delivered by the private sector alone.

- **Healthcare**: Cut referral-to-treatment time by 90 percent using AI for triage, scheduling, and early detection
- Target: Reduce median days referral-to-definitive-treatment from 120 days to 12 days across multiple high-volume pathways in at least two provinces
- Methods: AI triage and imaging, eReferral routing and load-balancing, operating room optimization, virtual pre-operative assessments, and AI-assisted documentation
- Higher fidelity early detection would drive prevention and reduce treatment loads
- **Education**: Reduce the Grade 3 reading non-proficiency rate by 90 percent
- Target: Reduction from approximately 30 percent to 3 percent across participating districts in at least two provinces or territories
- Methods: Adaptive K-3 screening, precision tutoring minutes, teacher copilots that lighten cognitive load, and attendance nudging

- Sustained gains and transparent reporting required
- **Natural Resources**: Shrink wildfire and oil spill response time by 90 percent
- Target: Cut median detection-to-containment time from 10 hours to 1 hour across at least two resource domains in different Canadian regions
- Methods: Multi-sensor anomaly detection, lightning-to-ignition nowcasts, drone tasking, crew pre-positioning, playbook generation, and incident-command copilots
- **Defense**: Reduce time to detect and classify airspace or maritime incursions by 90 percent
- Target: Reduce coverage-weighted time-to-detect-and-classify from 20 minutes to 2 minutes across defined Canadian approaches and domains
- Methods: Multi-sensor fusion and automated tracking at specified probability of detection and controlled false-alarm rates
- **Homelessness**: Cut the time from first contact to stable housing by 90 percent
- Target: Reduce median days-from-first-contact-to-stable-housing from 300 days to 30 days with 180-day retention, operating across Toronto, Vancouver, and Montreal
- Methods: Predictive triage, AI-assisted unit and benefit matching, process mining, and retention risk nudging

System-Level vs. Point-Level Solutions

- The paper emphasizes that AI is only part of the solution; moving the needle requires system-level rather than point-level solutions
- Point solution: A bolt-on AI model that improves one narrow decision inside an existing workflow while leaving upstream and downstream processes unchanged
- System solution: Reorganizes the entire decision flow to capitalize on cheap, accurate prediction—redesigning roles, handoffs, incentives, data rights, IT, and governance so that when the model outputs a prediction, the rest of the system can act on it immediately
- Governments have a comparative advantage in overseeing system-level solutions because they are especially well placed to manage coordination risk
- Estonia's "once-only" design rule was cited as an example: Agencies could not ask people to re-enter data the state already held and services were forced to pull verified information automatically via shared data infrastructure
- Result: Faster delivery, fewer errors, and staff time refocused on edge cases rather than routine verification

Implementation Principles

- Engage the creativity and leadership of the private, public, and not-for-profit communities through requests for proposals with meaningful budgets, selecting proposals based solely on their likelihood of success
- Commit budgets that recognize the importance of the moment

- Economies smaller than Canada such as the UAE and Saudi Arabia committed approximately 50 times more to AI than Canada in 2024
- Ensure lessons learned and operating procedures are well documented and shared freely, with progress reported regularly and publicly

AI Sovereignty

- AI sovereignty is not about building everything at home but about preserving Canadian optionality
- Turning AI into a sustained engine of national wealth
- Guaranteeing assured domestic access to critical inference so essential workloads keep running even if foreign providers falter
- Securing bargaining power by becoming a top-tier producer of at least one indispensable AI-system component

Meeting Discussion: Defining AI Strategy

- Participants debated whether AI strategy should be distinct from organizational strategy
- Consensus emerged that AI should be viewed as a tool to support broader strategic goals rather than as a standalone strategy
- Organizations should focus on their strategic objectives - whether faster healthcare, better education, or stronger defense - and deploy AI as a means to achieve those outcomes

Moonshots and Incrementalism

- Some participants expressed skepticism about the government's ability to deliver system-level solutions given bureaucratic constraints
- Others cited past national initiatives as evidence of potential, suggesting that Canada has historically demonstrated capacity for large-scale coordinated efforts when there is sufficient political will and public support

Coalition Building and Barriers

- Coalition building was highlighted as a key barrier to achieving ambitious goals
- Aligning incentives and resources across teams, organizations, and jurisdictions presents significant complexity
- Examples from healthcare and education illustrated how fragmented governance structures and competing priorities can impede progress on shared objectives, even when the potential benefits are widely recognized

Public Acceptance and Social License

- Concerns about public acceptance of AI in sensitive areas such as mental health and family violence were raised
- The importance of social license was emphasized—the implicit acceptance by the public that certain applications of AI are appropriate and trustworthy
- Government must address systematic issues and build partnerships to earn and maintain public trust in AI deployments that touch on vulnerable populations and sensitive personal circumstances

Canadian Context and Optimism

- Participants reflected on Canada's history of nation-building and the current opportunity for innovation
- There was recognition that rallying around shared goals and leveraging moments of crisis can drive systemic change
- The conversation touched on the importance of cultivating optimism about Canada's potential while remaining realistic about the challenges of implementation

AI Cybersecurity Risks and Organizational Responses: Chad Cogar led a discussion on recent advances in AI-driven cybersecurity threats, referencing developments such as Mythos and Claude Code, supply chain attacks, and organizational strategies for strengthening defenses. Attendees contributed perspectives on the role of government policy in addressing these emerging risks.

AI-Driven Hacking Capabilities

- AI tools like Claude Code have dramatically accelerated the process of discovering and exploiting security vulnerabilities
- These tools enable attackers to automate and scale penetration testing far beyond what traditional manual methods could achieve
- The combination of AI's ability to rapidly analyze code, identify weaknesses, and generate exploit code represents a fundamental shift in the threat landscape
- Compresses timelines that previously required skilled human hackers working over extended periods

Recent High-Profile Incidents

- The group discussed recent high-profile security incidents, including the McKinsey hack where AI agents exploited API vulnerabilities to gain unauthorized access
- This incident and others underscore the broader implications for organizations that rely on complex software dependencies and interconnected systems
- The expanding attack surface created by modern software architectures provides numerous entry points that AI-enabled attackers can systematically probe and exploit

Organizational Mitigation Strategies

- Review and reduce software dependencies to minimize potential attack vectors
- Implement cool-down periods before adopting new software versions which allows time for vulnerabilities to be discovered and patched before deployment into production environments
- Ensure security teams have access to advanced AI tools for proactive testing:
- Enables identification of weaknesses before malicious actors do

Government and Regulatory Role

- Dawn Hall described the government's role in setting responsible AI policies and establishing guardrails for public sector use of AI technologies
- Regulation of private sector AI use falls under other departments, reflecting the distributed nature of regulatory authority in this domain
- The EU AI Act was referenced as an example of broader regulatory efforts to establish frameworks for AI governance, suggesting that similar approaches may eventually be adopted in other jurisdictions

Project Glasswing and Defensive Collaboration

- Chad outlined Project Glasswing, an initiative in which advanced AI models are being shared with major infrastructure companies to identify and patch vulnerabilities before such capabilities become widely available to malicious actors
- This defensive collaboration represents an effort to stay ahead of the threat curve:
- Ensures critical infrastructure operators have access to the same tools that attackers might use
- Chad highlighted the urgency of updating software dependencies and acknowledged the challenges organizations face in maintaining readiness against rapidly evolving AI-enabled threats

MPAC AI Project Updates and Internal Innovation: Stratos and Tammy provided updates on the Municipal Property Assessment Corporation's AI initiatives, covering the Orchestrator platform, pursuit of ISO 42001 certification, the AI Academy educational program, and the Pitch program for sourcing and selecting new AI projects. Attendees asked questions about project selection criteria and opportunities for knowledge sharing with other organizations.

AI Orchestrator and Governance

- MPAC's Orchestrator serves as the central governance and compliance engine for AI projects within the organization
- The platform provides a unified framework for managing AI initiatives, ensuring consistency in oversight and adherence to organizational policies

- Ongoing efforts are focused on aligning the Orchestrator and associated practices with ISO 42001 requirements
- Goal: Establishing a formal centre of excellence for AI governance

AI Academy and Staff Engagement

- The AI Academy initiative aims to educate both business and IT staff about AI technologies and their applications
- The program promotes a "human in the lead" approach, emphasizing that AI tools should augment human decision-making rather than replace it
- The Academy is designed to:
  - Address concerns about AI adoption
  - Improve AI literacy across the organization
  - Ensure staff at all levels understand both the capabilities and limitations of AI systems

Project Selection and Metrics

- Project selection at MPAC is driven by a company-wide Pitch program
- Ideas for AI projects are sourced from employees throughout the organization, reflecting a bottom-up approach to innovation
- Proposals are evaluated through design thinking workshops that help refine concepts and assess feasibility
- Final selection decisions are based on criteria including return on investment, scalability potential, and ethical considerations
- Early metrics from implemented projects indicate significant improvements in outcomes post-implementation, validating the effectiveness of the selection and development process

Cost Management and Tool Access

- The discussion addressed strategies for managing the cost of co-pilot licenses, which can be prohibitive for organization-wide deployment
- The team is exploring open-source alternatives that could provide broader staff access to AI capabilities while controlling costs
- Ensuring secure integration with internal systems remains a priority
- Balancing the desire for accessibility with requirements for data protection and governance compliance

Knowledge Sharing and Scalability

- MPAC expressed willingness to share lessons learned and project selection frameworks with other public sector organizations
- This openness to collaboration reflects the recognition that many organizations face similar challenges in adopting AI and could benefit from shared knowledge

- Stratos and Tammy noted an interesting evolution in the types of proposals received through the Pitch program:
  - Early submissions frequently focused on chatbot automation
  - More recent proposals increasingly emphasize AI-powered augmentation tools that enhance human capabilities rather than simply automating routine tasks

Follow-Up Tasks: The meeting concluded with the identification of several follow-up actions to be completed before the next session.

- **ISO 42001 Certification Collaboration (Jesslyn)**: Arrange a follow-up discussion with her group regarding ISO 42001 certification to share insights and collaborate on the certification process with interested parties
- **Sharing AI Project Selection Criteria and Lessons Learned (Stratos, Tammy)**: Share MPAC's AI project selection criteria, metrics for evaluating project effectiveness, and lessons learned with Arjan and other interested public sector organizations to support their potential adoption of AI initiatives
- **Providing Access to Scaling Roadmap (Dawn)**: Send the link to the published scaling roadmap to interested parties for their reference and use. **The link can be found here**
- **Exploring Off-the-Shelf Chatbot Solutions (Stratos)**: Evaluate and consider implementing open-source or off-the-shelf chatbot front ends, such as Open Web UI, to provide AI chat capabilities to staff without Copilot licenses while ensuring integration with internal orchestration and security controls