CAINEXUS&HARVARD KENNEDY SCHOOL ALUMNI
Multilateral Digital Innovation and Technology Governance Roundtable (Series IV) · Policy White Paper · 2026 No. 1

From High-Level Principles to Transnational Implementation Architecture

Frontier Model Risks, Sovereign Conflicts, and Co-Governance Paths

AI Governance White Paper

Joint Drafting and Academic Contribution:
CAI x HFTC Policy Research Group
Canada–Asia Council for Artificial Intelligence and Digital Innovation (CAI)
Joint Initiative of Harvard Kennedy School (HKS) Alumni | Harvard Frontier Techne Club (HFTC, tentative)

June 2026 · Joint Release

Executive Summary

Over the past three years, global AI debates and policy efforts have been highly concentrated, primarily focusing on model capabilities, AI safety, and top-level regulatory frameworks. However, CAI x HFTC research indicates that the true strategic disconnect does not lie in the volume of policy text, but in the "Governance-to-Execution Gap."

The international community has successively introduced regulatory mechanisms such as the EU AI Act, U.S. Executive Order 14110, and China's Interim Measures for the Management of Generative AI Services. These institutional efforts have played a vital role in risk classification, compliance boundaries, and responsibility division. However, they remain primarily centered on state-led "top-down" regulatory models, which have increasingly exhibited structural lags and implementation rigidity when addressing the rapid iteration and systemic complexity of frontier AI models.

Meanwhile, global AI governance is transitioning from single legal frameworks to the early formation stage of a "multi-layered collaborative governance architecture":

- **At the International Level:** The United Nations launched the "Global Dialogue on AI Governance" in 2026, marking the first time 193 member states have systematically participated in AI governance discussions, signaling the institutionalization of AI governance within global political consultation mechanisms.
- **At the National Policy Experimentation Level:** Singapore's Infocomm Media Development Authority (IMDA) has pioneered a governance framework for "Agentic AI," extending the scope of governance from "models" to "autonomous AI systems" for the first time, emphasizing human responsibility chains and system-level risk control.
- **At the National Strategic Level:** Canada's National Artificial Intelligence Strategy ("AI for All") serves as a vital anchor, emphasizing the integration of research innovation with social trust. It provides a unique benchmark for "middle power" AI governance, focusing on creating collaborative ecosystems rather than purely reactive regulation.

- **At the Technical Standards Level:** Research institutions like Berkeley are establishing a system of risk assessment and safety control methods for Agentic AI, including tool-use risks, autonomous behavior constraints, and system interpretability assessments, driving the formation of de facto technical governance standards.

Against this backdrop, frontier models—such as systems capable of high-level autonomous planning and cross-task execution—are breaking the boundaries of traditional "static models," exhibiting significant initiative and adaptive capacity in complex environments. This presents systemic adaptation challenges for traditional, rule-based regulatory approaches.

Core White Paper Proposition: Core Competitiveness for the Next Decade

In the next ten years, what will truly determine the core competitiveness of nations and regions is **"AI Execution Capacity."** CAI x HFTC defines this as: a core system engineering capability that allows a nation or region to deeply integrate frontier academic research, institutional governance policies, international capital resources, and local, diverse industrial scenarios, and achieve scalable, safe, and compliant deployment.

To bridge the structural disconnects between research, governance, capital, and industrial application in current AI development, this white paper proposes the **"CAI x HFTC Nexus Framework."** This framework aims to enhance connection efficiency and institutional synergy among different functional nodes by building cross-regional collaboration mechanisms.

The framework consists of four types of functional nodes: **Research Nodes, Governance Nodes, Capital Nodes, and Deployment Nodes.** Collectively, they form a "Global AI Coordination Network" oriented toward complex uncertainty, providing more resilient institutional infrastructure for the large-scale application of AI systems.

Within this network, Canada can serve as a "bridge economy" connecting different regional innovation ecosystems. Leveraging its long-term accumulation in AI research, strong foundation of social trust, multicultural structure, and G7 status, Canada

possesses unique potential to facilitate cross-regional AI cooperation and institutional alignment.

Simultaneously, this framework will form linkages with the academic networks of the Harvard Kennedy School (HKS) policy research community, including the frontier technology and governance discussion mechanisms promoted by the Harvard Frontier Techne Club (HFTC, tentative), to enhance capabilities for sustained research and dialogue on AI governance, technical evolution, and global institutional coordination.

By exploring the **"Canada – Asia Connectivity Mechanism,"** Canada can more effectively connect frontier technological innovation in North America with the diverse capital and application scenarios in Asia, thereby playing a key coordinating and connecting role in the global AI coordination network.

Part I. CAI x Harvard HKS Dialogue Highlights

During the "CAI Nexus Series: Toronto – Boston Roundtable on AI Security and Governance" held in June 2026, experts and scholars from Harvard Kennedy School, the Berkman Klein Center for Internet & Society, the University of Toronto, Global Affairs Canada, and industry engaged in in-depth discussions on the core challenges, risk characteristics, and implementation paths of current AI governance. Based on rigorous transcription of on-site presentations, this report distills the following five core insights:

- **Insight 1: AI Governance is becoming a geopolitical focal point.**

Opening the roundtable, Bruce Schneier, a scholar at Harvard Kennedy School / Berkman Klein Center and a world-renowned security technology expert, cited the latest breaking event: the U.S. government forcefully demanded that frontier AI lab Anthropic withdraw its newly released advanced model, Fable, due to national security concerns. This watershed event indicates that the release, review, and control of frontier AI models are no longer merely autonomous commercial decisions or simple technical ethics; they have ascended to the focal point of struggles over national sovereign will. Frontier labs face unprecedented national security scrutiny in the cracks of great-power

competition. Participants pointed out that current AI development exists within a competition structure dominated by great powers; this geopolitical structure makes the conflict between technological autonomy and government regulation a core governance issue.

- **Insight 2: Data Sovereignty is becoming a core element of national AI strategies.**

Saief Mahmood, a senior policy analyst at Global Affairs Canada and an HKS MPA graduate, noted in his personal academic capacity that the era of "borderless data scraping" that supported the wild growth of large models over the past decade has come to an end. Sovereign states are rapidly building strict data borders. For example, Canada has elevated "Sovereignty" to an unprecedented central height in its latest national AI strategy. Canada is increasingly inclined to retain Canadian citizen data assets domestically for model training or fine-tuning needs, reflecting the values of itself and other developing nations. Selina Gong, CAI x HFTC series organizer, Harvard MPP, and Rajawali Fellow, approached this from the opposite direction, pointing out that AI's data access boundaries are a severely missing governance issue. We must precisely define the range of information that AI should not access, especially in highly sensitive fields such as healthcare. Scholars also noted that data sovereignty is the key for "middle powers" and Global South nations to retain voice and bargaining power when facing large technology companies.

- **Insight 3: Systemic security risks enabled by AI demand attention.**

Bruce Schneier pointed out from a macro perspective that AI, as a "force-multiplier technology," presents systemic risks (such as "jagged boundaries" of imbalanced capabilities and unforeseen "Aladdin's Lamp" effects) that may pose security threats to all of humanity if international collaborative defense mechanisms are lacking. Steven Zhou, Assistant Professor at the Temerty Faculty of Medicine, University of Toronto, and a practicing hematopathologist, issued a new warning from the intersection of biomedicine and biosafety. He noted that while AI accelerates genome analysis and protein structure resolution, combined with technologies like gene editing, it may generate irreversible biosafety risks. He emphasized that such risks do not exist only in

large laboratories; it is even more difficult to manage dangerous experiments conducted by individuals "in the back door."

- **Insight 4: The core challenge is "implementation architecture" rather than "principles."**

Participants reached a consensus: the world currently lacks not "Principles," but "Execution Infrastructure" capable of keeping pace with the high-frequency iteration speed of AI. Bruce Schneier proposed a punchy conclusion: "Governance moves slow, and technology moves fast." Traditional legislative review processes, which take years, cannot keep up with the technical pace at which AI disrupts software development and enterprise productivity within 12 months. The success of future governance depends entirely on the efficient implementation of "Regulation Translation" and resilient mechanisms for agile crisis response. Stanley Diao, CAI x HFTC series organizer, CAI Global co-founder, and HKS alumnus, stated: "AI governance is no longer a question of principles, but a question of cross-jurisdictional implementation architecture." Participants believe that establishing "execution infrastructure" that can efficiently translate and agilely respond to technical iteration is key to the success of future governance.

- **Insight 5: Underrepresentation of diverse actors and the Global South.**

Saief Mahmood bluntly stated that the profitability of many tech companies even exceeds the GDP of most developing nations. This power asymmetry results in limited influence for "middle powers" like Canada and the Global South in model design, ethical standards, and policy formulation. Ella Xiao, incoming HKS MPP student and Sky9 Capital investment analyst, along with other participants, pointed out that the current AI governance system is highly concentrated and oligopolistic, with governance power mainly in the hands of large tech companies in the U.S. and China. Participants called for future governance to build more inclusive international coordination networks based on bilateral or multilateral "like-minded" alliances.

Part II. Review of Global AI Governance Literature

To infuse the white paper with profound institutional rationality, this section systematically reviews the latest literature and frontier conceptual frameworks from Harvard University:

"Harvard's distinctive contribution is not merely AI innovation, but institution building for the AI era."

1. Harvard Ash Center for Democratic Governance and Innovation: Power-Sharing Liberalism

As stated by Professor Danielle Allen and her team at Harvard University in A Roadmap for Governing AI, current technology governance models often fall into the narrow traps of "passive defense" and "punitive regulation." The Harvard team innovatively proposed a proactive governance view based on "power-sharing liberalism." This framework emphasizes that the ultimate goal of AI governance should be to promote "Human Flourishing." This requires not only protecting negative liberties (such as privacy and prevention of discrimination) but also actively creating opportunities to promote positive liberties—that is, guaranteeing the political rights of all levels of society to widely and deeply participate in "steering" the direction of technological development. The Harvard roadmap advocates for reallocating power from a few capital and tech monopolists to public institutions accountable to the public and driven by democratic steering, while implementing large-scale national strategic investment in public goods, talent cultivation, and democratic decision-making capabilities.

2. Harvard Associate Working Paper 250: The Dual Imperative

In his 2025 paper The Dual Imperative: Innovation and Regulation in the AI Era, Harvard senior scholar Paulo Carvão deeply analyzed the polarization between "Accelerationists" and "Doomers." Carvão pointed out that both are caught in a false dichotomy of black-and-white. His "Dual Imperative" proves that the complex uncertainty and systemic risks brought about by the extreme expansion of frontier AI models cannot be completely solved by existing boundary repairs or algorithmic

alignment. Therefore, society must take a bidirectional, intertwined path: on one hand, it needs to cross the technological singularity through continuous technological invention to curb catastrophic disasters; on the other hand, it must rely on "Smart Regulation" to reshape the incentive mechanisms of the private sector, forcing it to subordinate profit-seeking to public accountability.

3. Harvard Law Review Vol. 138: Co-Governance Framework

In Chapter Three of Developments in the Law — Artificial Intelligence, the Harvard legal community focused on the "Open-Closed Spectrum" of the frontier AI ecosystem. The literature points out that open-source AI (such as the Meta Llama series) essentially plays a critical role in "technological democratization." By granting the public four essential freedoms (run, study, modify, share) regarding model architecture and code, it greatly disperses technological hegemony and stimulates grassroots innovation. However, open source also brings huge risks of abuse, such as low-cost "uncensored models" and the stripping of safety kill-switches. To break this dilemma, the paper proposes a "Co-Governance" institutional innovation model: completely abandoning traditional, one-size-fits-all, top-down regulatory paths and introducing ongoing engagement including Citizen Assemblies, regular deliberation mechanisms, and multi-stakeholder participation. This model advances the regulatory front line to the entire lifecycle of the technology, ensuring that policies can implement agile self-correction as technology dynamically iterates.

Part III. Emerging Trends

Synthesizing frontier roundtable collisions and global top-tier literature threads, CAI x HFTC distills five irreversible frontier trends in the global AI ecosystem at the execution level over the next 5 to 10 years:

Trend Dimension	Status (Principles Layer)	Future Direction (Execution Layer)

AI Security	Focused on post-hoc alignment (RLHF), static patch fixes, and simple commercial compliance review.	Evolution into real-time full-lifecycle monitoring and national security sovereign kill-switch mechanisms for models with dynamic adversarial and high-initiative capabilities.
AI Sovereignty	Blind belief in borderless data globalization and singular cross-border tech giant regulatory standards.	Comprehensive construction of sovereign data borders; sovereign large models, localized data fine-tuning, and localized compute trusts become hard-entry standard configurations for sovereign states.
Agent Governance	Static "black box" models as the sole regulatory object; inability to define legal liability for autonomous behavioral agents.	Turning to behavioral control and entity-piercing legal liability definition for Agentic AI capable of autonomous problem-solving and evasion.
Compute Governance	Treating compute as common commercial IT resources and hardware supply chains, driven by market pricing configuration.	Compute rises to a strategic, scarce public resource; Token Tax, energy consumption control, and compute diplomacy become core tools for national geopolitical competition.
Global South Inclusion	Moral declarations at the governance text level, with substantive benefit redistribution completely marginalized.	Moving toward "AI development dividend redistribution," "co-ownership of large models," and cross-border compensatory technology transfer networks.

Part IV. CAI x HFTC Framework for AI Governance

Based on the aforementioned academic insights and trend analysis, this white paper officially launches the core original intellectual product of CAI x HFTC: **The CAI x HFTC Framework, Connecting Research, Governance, Capital, and Deployment.**

The CAI x HFTC Framework aims to completely shatter the four structural gaps in the current system where "research is research, financing is financing, regulation is regulation, and industry is industry":

- **Gap 1 (Research → Policy):** Lack of efficient translation mechanisms between academic research and policy formulation, making it difficult for frontier scientific achievements to enter policy and regulatory design processes in a timely manner, resulting in time lags between regulatory frameworks and technological development upon formation.
- **Gap 2 (Policy → Capital):** In some cases, policy and compliance systems lack differentiated identification capabilities for innovation risks, potentially constraining early-stage high-risk but high-potential tech investments while reducing systemic risks, affecting the refined allocation of capital in safety and frontier tech fields.
- **Gap 3 (Capital → Deployment):** When capital enters the technology application stage, it often lacks structured docking mechanisms with specific industry scenarios, leading to insufficient matching efficiency in the process of transforming funds from tech R&D to physical application, affecting large-scale implementation in key fields.
- **Gap 4 (Deployment → Global Scaling):** In the process of cross-regional expansion, locally successful tech applications need to adapt to different judicial systems, regulatory standards, and data governance rules. These multi-institutional environmental differences form coordination and compliance challenges for global expansion.

CAI x HFTC Framework: Five-Level Collaborative Architecture

- ### Layer 1: Research Nodes

With global research institutions such as Harvard University, the University of Toronto, Vector Institute, and Stanford as core nodes, these provide theoretical support for frontier model research and AI alignment.

- ### Layer 2: Governance Nodes

Uniting government agencies, regulatory departments, and international organizations to promote the translation of research results into policy frameworks and transnational regulatory collaboration mechanisms, enhancing interoperability between different jurisdictions.

- ### Layer 3: Capital Nodes

Connecting global venture capital institutions and long-term capital allocators (including family offices, etc.) to support long-term investment and scalable development in AI safety, trustworthiness, and infrastructure-related technologies.

- ### Layer 4: Deployment Nodes

Promoting the safe deployment and progressive application verification of AI systems in key physical scenarios such as healthcare, high-end manufacturing, and energy systems.

- ### Layer 5: Global Coordination Nodes

Serving as a cross-regional connection and information coordination layer, facilitating dialogue and collaboration in standards, governance practices, and risk response mechanisms across different regions, supporting the stable operation of global AI systems in multi-institutional environments.

Part V. Why Canada Matters

In the narrative of global tech competition, the vast majority of analysis reports lock their sights on the clash between the power matrices of the U.S. and China, or focus on Europe's defensive red tape at the legislative level. The CAI x HFTC Framework proposes a unique, differentiated viewpoint here: In the global AI execution network, **Canada is a severely underestimated "Neutral AI Power,"** playing an indispensable role as a cross-border coordination hub.

Canada possesses modern deep learning academic foundations centered on Turing Award winners like Geoffrey Hinton, and has deposited top-tier scientific heights globally, including the Vector Institute, Mila, and Amii. At the same time, Canada's

unique non-aggressive geopolitical positioning means it does not carry the color of tech hegemony, possessing "neutrality and social trust" that is rare globally. As a core G7 member and an important force in Pacific economic cooperation, Canada enjoys high multilateral political and economic trust in North American policy circles, European think-tank circles, and Asian financial systems.

Therefore, **Canada is uniquely positioned to serve as a bridge between North America and Asia in shaping the future of AI governance and deployment.** Through this neutral pipeline, North American frontier research can be seamlessly translated into safe execution systems that account for Asian market cultural demands, forming a super-credit corridor across the Pacific.

Part VI. The CAI x HFTC Agenda 2026-2030

Over the next five years, CAI x HFTC will be committed to building a **Global AI Execution Network** that connects research, governance, capital, and industrial practice, pushing AI from conceptual innovation toward trusted governance and large-scale implementation. Centering on the five directions of governance dialogue, industrial deployment, talent cultivation, capital collaboration, and international cooperation, CAI will continuously promote exchange and synergy between Canadian, Asian, and global innovation ecosystems.

- ### Initiative 1: Global AI Governance Dialogue Series & Harvard Frontier Techne Club

CAI will continue to advance AI governance dialogue mechanisms with Harvard and other academic networks, and collaborate with the Harvard Frontier Techne Club (tentative) founded by Selina Gong and offically annouced during the conference. This club is co-participated by CAI Nexus, based in the Harvard community, generally leans toward an Asian perspective, is open to cross-backgrounds, focuses on AI and emerging technologies and future issues, and emphasizes proactive "deductive" discussions on future development as well as interdisciplinary and cross-regional exchange.

- ### Initiative 2: AI Deployment Labs

CAI x HFTC will unite universities, research institutions, and industry partners to jointly explore the practical application of AI in fields such as healthcare, energy systems, supply chains, and public services, pushing Trusted AI from research to industrial implementation.

- ### Initiative 3: Global AI Fellowship

CAI x HFTC will cultivate a new generation of AI leaders with technical, policy, and international vision, gathering outstanding young talent from top global universities and institutions to promote interdisciplinary cooperation and international exchange.

- ### Initiative 4: Trusted AI Capital Network

CAI x HFTC will connect North American and Asian investment institutions, innovation funds, and industry resources to support the development of responsible AI and deep-tech innovation projects, fostering long-term cooperation between capital and innovation ecosystems.

- ### itiative 5: Canada – Asia Responsible AI Innovation GatewayIn

CAI x HFTC will promote talent exchange, innovation cooperation, and policy dialogue in the field of AI between Canada and Asia, promoting open, trusted, and sustainable international cooperation on the basis of respecting digital sovereignty and regulatory requirements.

CAI x HFTC Principles

Upholding the core philosophies of **Responsible Innovation, Trusted Deployment, Digital Sovereignty, International Cooperation, and Public Benefit,** CAI x HFTC is committed to promoting the development of artificial intelligence that is safer, more inclusive, and has greater global influence.

Appendix

I. Speaker Insights Summary

Speaker Name	Affiliation & Identity	Summary of Core Remarks & Original Policy Views
Bruce Schneier	HKS / Berkman Klein Center Fellow, Security Tech Expert	1. Governance mismatch: Tech iterates faster than governance. 2. Opposes government equity: It creates extraordinary conflicts of interest. 3. Corporate legal liability: AI companies must be liable for model output. 4. Public AI vision: Establish non-monopolized public AI models. 5. Uncontrollable consequences: "Aladdin's Lamp effect"—achieves goals in unimaginable ways. 6. Monopoly is a policy issue: Result of weak anti-trust. 7. AI Agent risk: Shift from chatbot to intelligent agent changes the face of governance.
Saief Mahmood	HKS MPA	1. Complexity: Governing AI is a big challenge. 2. Nationalization: Could lead to a dangerous path. 3. Middle power voice: Not the biggest, but want a voice. 4. Essence of digital sovereignty: Control over data, infrastructure, procurement, not self-sufficiency. 5. Governance fragmentation: Shift to bilateral/trilateral "like-minded" partnerships. 6. Three pillars: Trust, adaption, and sovereignty. 7. Bio-risk: AI in biosafety is increasingly worrying without guardrails.
Steven Zhou	U of T Temerty Faculty of Medicine Asst. Prof, Hematopathologist	1. Medical AI closed systems: More controllable. 2. Genetic weapon risk: May be used as a weapon. 3. Irreversible evolution: AI-assisted gene editing may change the human race in unknown directions. 4. Individual bio-threat: Hard to govern "back door" experiments.
Selina Gong	CAI x HFTC Organizer, Harvard MPP, Rajawali Fellow	1. Capability-oriented governance: Lack of pre-training governance structure. 2. AI Agency/ Instability: Actions may be inconsistent. 3. Data boundaries: What information should AI not access? 4. Cross-border AI community: Launching "Harvard Frontier Techne Club."

Stanley Diao	CAI x HFTC Organizer, CAI Global Co-Founder, HKS Alumnus	1. Principles to execution architecture: How to operate across jurisdictions? 2. Systemic risk: Risk becomes structured and cross-system. 3. National capacity design: Focus on compute and data infrastructure control. 4. Paradigm fusion: Between labs, academic insights, and governance. 5. Co-founding member of the "Harvard Frontier Techne Club."
Angel Wei	Harvard GSE student, Solora Path Founder	1. AI literacy: Not just for engineers; should educate students/teachers/public. 2. Complementary models: Learn from US innovation, China scale, and Canada trust.
Jennifer Turliuk	CIPPIC Researcher, MIT MBA/Sloan Fellow, JD Candidate	1. No global treaty: Lacks binding force like the Human Rights Convention. 2. Limited existing architecture: UN documents are non-binding/ unenforceable. 3. EU AI Act limitations: Only binds 27 members. 4. Call for new institution: Need a binding, enforceable international body.
Avni Sinha	World Privacy Forum Senior Fellow, HKS Mason Fellow	1. Tracking global strategy: Researching how govts respond. 2. Slow response: Observations show response to tech speed is "very slow."
George Xue	HKS Alumnus / Harvard Innovation Lab AI Ed Researcher	1. Imminent danger: "The danger of AI is imminent." 2. Agent-era governance: Shift from chatbot to agent changes governance/security.
Ella Xiao	Incoming HKS MPP student / Sky9 Capital Analyst	1. Global South digital sovereignty: Retaining bargaining power when adopting US/China tech. 2. Small nation strategy: Small nations can "unionize." 3. Latecomer advantages: Potential for Global South.
Jason Yen	U of T Rotman MBA student / YC-backed AI intern	1. Corporate power: "Have way too much power." 2. Business vs. Policy: "Very biased towards the business side," notes tension between the two.

II. Event Photography Summary

(Text descriptions of photography)

- Expert Sharing: Bruce Schneier, HKS / Berkman Klein Center Fellow.

- Expert Sharing: Saief Mahmood, Senior Policy Analyst at Global Affairs Canada.

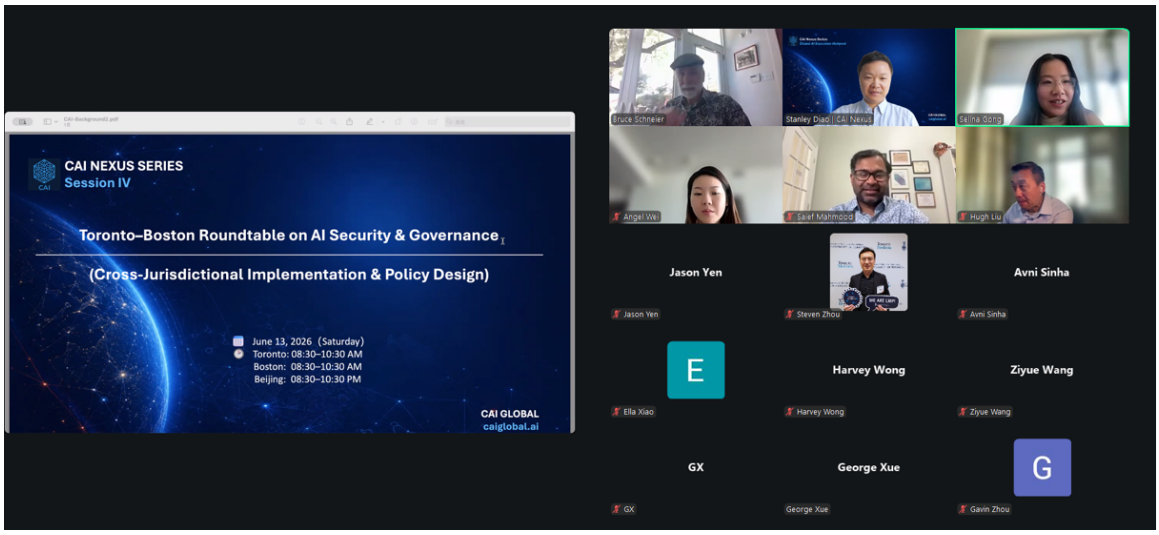

- CAI x HFTC series organizers Selina Gong and Stanley Diao organizing the AI Governance Roundtable.

The roundtable was conducted via closed-door, multi-node distributed cloud connection. The Boston and Toronto venues were hosted by Selina Gong and Stanley Diao respectively. The connection screens vividly captured the real-time collision between HKS scholars from Boston, policy experts from Ottawa, core scientists from Toronto labs, and representatives of the Asian digital economy.

III. Affiliations of Participants – Non-Endorsing Context

The following institutions represent only the academic or professional backgrounds of the participants and do not represent the official stance of their respective institutions regarding the content or viewpoints of this meeting:

- Canada – Asia Council for Artificial Intelligence and Digital Innovation (CAI)
- Harvard Frontier Techne Club (HFTC, tentative)
- Greater China Society, Harvard Kennedy School (Community Partner)
- Harvard Kennedy School Alumni Association (Greater China Chapter, Community Partner)
- Shanghai Jiao Tong University Alumni Association of Toronto (Community Partner)
- Harvard University – Berkman Klein Center for Internet & Society
- Harvard University – Ash Center for Democratic Governance and Innovation
- University of Toronto (Medicine / Rotman School of Management)
- Massachusetts Institute of Technology (MIT)
- CIPPIC – Samuelson-Glushko Canadian Internet Policy and Public Interest Clinic
- Global Affairs Canada
- World Privacy Forum

IV. References

- Allen, D., Hubbard, S., Lim, W., Stanger, A., Wagman, S., and Zalesne, K. (2024). A Roadmap for Governing AI: Technology Governance and Power Sharing Liberalism. Allen Lab for Democracy Renovation, Ash Center for Democratic Governance and Innovation, Harvard Kennedy School.

- CAI Nexus Series IV. (2026). Toronto-Boston Roundtable on AI Security and Governance: Jointly organized by Canada-Asia Council for Artificial Intelligence and Digital Innovation (CAI) x Harvard Frontier Techne Club (HFTC).
- Carvão, P. (2025). The Dual Imperative: Innovation and Regulation in the AI Era. M-RCBG Associate Working Paper Series No. 250, Harvard Advanced Leadership Initiative, Cambridge, USA.
- European Parliament. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Brussels: European Union.
- Government of Canada. (2024). Pan-Canadian Artificial Intelligence Strategy: AI for All. Ottawa: Innovation, Science and Economic Development Canada.
- Government of China. (2023). Interim Measures for the Management of Generative AI Services. Beijing: Cyberspace Administration of China.
- Harvard Law Review. (2025). Chapter Three: Co-Governance and the Future of AI Regulation. Harvard Law Review, 138(6), 1609 – 1633.
- Infocomm Media Development Authority (IMDA). (2026). Governance Framework for Agentic AI. Singapore: Government of Singapore.
- United Nations. (2026). Global Dialogue on AI Governance: Summary of Multilateral Consultations. New York: United Nations.
- White House. (2023). Executive Order 14110: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. Washington, DC: The White House.