



**The Journal of Robotics,
Artificial Intelligence & Law**

Editor's Note: Physical AI

Victoria Prussen Spears

Physical AI: The New Patent Frontier

Matthew H. Grady

When Mathematicians Push Back: The Leiden Declaration on AI and Mathematics and Why It Matters

Marco Blei

Algorithms as Exhibit A: Practical Steps for Navigating Discovery of AI Evidence in Workplace Disputes

Andrew R. Turnbull

Singapore's Agentic AI Framework: Practical Guidance for Market Entry

Amita Haylock and Justin Lai

South Africa: AI Policy Moves Toward Approval

Ashlin Perumall, Fatima Ismail, and Despina Lazanakis

Artificial Intelligence and Human Resources in the EU

Emmanuel Plasschaert, Sari Depreeuw, Evelien Jamaels, and Sofiane Fergali

Employer AI Literacy Obligations Under the EU AI Act

Lauren Cuyvers, Evelien Jamaels, and Sofiane Fergali

UK Copyright and AI Report: The "Opt-Out" Is Dead, But What Comes Next?

Nick Breen, Henry Birkbeck, Harry Clark, and Oliver "Ollie" Hogg

- 397 Editor's Note: Physical AI**
Victoria Prussen Spears
- 401 Physical AI: The New Patent Frontier**
Matthew H. Grady
- 417 When Mathematicians Push Back: The Leiden Declaration on AI and Mathematics and Why It Matters**
Marco Blei
- 423 Algorithms as Exhibit A: Practical Steps for Navigating Discovery of AI Evidence in Workplace Disputes**
Andrew R. Turnbull
- 431 Singapore's Agentic AI Framework: Practical Guidance for Market Entry**
Amita Haylock and Justin Lai
- 437 South Africa: AI Policy Moves Toward Approval**
Ashlin Perumall, Fatima Ismail, and Despina Lazanakis
- 445 Artificial Intelligence and Human Resources in the EU**
Emmanuel Plasschaert, Sari Depreeuw, Evelien Jamaels, and Sofiane Fergali
- 451 Employer AI Literacy Obligations Under the EU AI Act**
Lauren Cuyvers, Evelien Jamaels, and Sofiane Fergali
- 457 UK Copyright and AI Report: The "Opt-Out" Is Dead, But What Comes Next?**
Nick Breen, Henry Birkbeck, Harry Clark, and Oliver "Ollie" Hogg

EDITOR-IN-CHIEF

Steven A. Meyerowitz

President, Meyerowitz Communications Inc.

EDITOR

Victoria Prussen Spears

Senior Vice President, Meyerowitz Communications Inc.

BOARD OF EDITORS

Jennifer A. Johnson

Partner, Covington & Burling LLP

Paul B. Keller

Partner, Allen & Overy LLP

Garry G. Mathiason

Shareholder, Littler Mendelson P.C.

James A. Sherer

Partner, Baker & Hostetler LLP

Elaine D. Solomon

Partner, Blank Rome LLP

Edward J. Walters

Chief Strategy Officer, vLex

John Frank Weaver

Director, McLane Middleton, Professional Association

THE JOURNAL OF ROBOTICS, ARTIFICIAL INTELLIGENCE & LAW (ISSN 2575-5633 (print) /ISSN 2575-5617 (online) at \$495.00 annually is published six times per year by Full Court Press, a Fastcase, Inc., imprint. Copyright 2026 Fastcase, Inc. No part of this journal may be reproduced in any form—by microfilm, xerography, or otherwise—or incorporated into any information retrieval system without the written permission of the copyright owner. For customer support, please contact Fastcase, Inc., 729 15th Street, NW, Suite 500, Washington, D.C. 20005, 202.999.4777 (phone), or email customer service at support@fastcase.com.

Publishing Staff

Publisher: David Nayer

Production Editor: Sharon D. Ray

Cover Art Design: Juan Bustamante

Cite this publication as:

The Journal of Robotics, Artificial Intelligence & Law (Fastcase)

This publication is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional should be sought.

Copyright © 2026 Full Court Press, an imprint of Fastcase, Inc.

All Rights Reserved.

A Full Court Press, Fastcase, Inc., Publication

Editorial Office

729 15th Street, NW, Suite 500, Washington, D.C. 20005

<https://www.fastcase.com/>

POSTMASTER: Send address changes to THE JOURNAL OF ROBOTICS, ARTIFICIAL INTELLIGENCE & LAW, 729 15th Street, NW, Suite 500, Washington, D.C. 20005.

Articles and Submissions

Direct editorial inquiries and send material for publication to:

Steven A. Meyerowitz, Editor-in-Chief, Meyerowitz Communications Inc.,
26910 Grand Central Parkway, #18R, Floral Park, NY 11005, smeyerowitz@
meyerowitzcommunications.com, 631.291.5541.

Material for publication is welcomed—articles, decisions, or other items of interest to attorneys and law firms, in-house counsel, corporate compliance officers, government agencies and their counsel, senior business executives, scientists, engineers, and anyone interested in the law governing artificial intelligence and robotics. This publication is designed to be accurate and authoritative, but neither the publisher nor the authors are rendering legal, accounting, or other professional services in this publication. If legal or other expert advice is desired, retain the services of an appropriate professional. The articles and columns reflect only the present considerations and views of the authors and do not necessarily reflect those of the firms or organizations with which they are affiliated, any of the former or present clients of the authors or their firms or organizations, or the editors or publisher.

QUESTIONS ABOUT THIS PUBLICATION?

For questions about the Editorial Content appearing in these volumes or reprint permission, please contact:

David Nayer, Publisher, Full Court Press at david.nayer@clio.com or at
202.999.4777

For questions or Sales and Customer Service:

Customer Service

Available 8 a.m.–8 p.m. Eastern Time
866.773.2782 (phone)
support@fastcase.com (email)

Sales

202.999.4777 (phone)
sales@fastcase.com (email)

ISSN 2575-5633 (print)
ISSN 2575-5617 (online)

Singapore's Agentic AI Framework: Practical Guidance for Market Entry

Amita Haylock and Justin Lai*

In this article, the authors examine Singapore's Model AI Governance Framework for Agentic AI, which provides AI developers and deployers with structured guidance on managing the unique risks presented by agentic AI systems.

Singapore continues to position itself as a leading hub for artificial intelligence (AI) innovation. In February 2026, Prime Minister Lawrence Wong announced the establishment of a National AI Council¹ to oversee “AI missions” aimed at transforming four core sectors: advanced manufacturing, connectivity, finance, and healthcare. These sectors combine high-volume workflows with high-stakes environments demanding strong governance. Further, to encourage small and medium-sized enterprises to accelerate AI adoption, the government's Enterprise Innovation Scheme,² originally meant to encourage businesses to engage in research and development, innovation, and capability development activities, will be expanded to permit businesses to claim 400 percent tax deductions on qualifying AI expenditures, capped at S\$50,000 (US\$39,600) annually for 2027 and 2028.

Against this backdrop of commercial momentum, the Infocomm Media Development Authority released the Model AI Governance Framework for Agentic AI³ on 22 January 2026. This Model Agentic AI Framework provides AI developers and deployers with structured guidance on managing the unique risks presented by agentic AI systems—the next evolution of AI technology that holds transformative potential for users and businesses alike.

Understanding Agentic AI

Agentic AI systems are those capable of planning across multiple steps to achieve specified objectives, using AI agents that

possess some degree of independent planning and action-taking. Unlike generative AI, which responds to prompts, AI agents can take actions, adapt to new information, and interact with other agents and systems to complete tasks on behalf of humans. The core components of an AI agent include the model (serving as the central reasoning and planning engine), instructions (defining the agent's role, capabilities, and behavioural constraints), memory (allowing the agent to store and access information from previous interactions), planning and reasoning capabilities (enabling the agent to output a series of steps needed to complete a task), and tools (enabling the agent to interact with other systems, such as writing to files and databases, controlling devices, or performing transactions).

While use cases are rapidly evolving, agents are already transforming the workplace through coding assistants, customer service agents, and automating enterprise productivity workflows. These greater capabilities also bring forth new risks. Agents' access to sensitive data and ability to make changes to their environment—such as updating a customer database or making a payment—may raise concerns. If agents malfunction, they could lead to harmful real-world impacts, including erroneous actions, unauthorised actions, biased or unfair actions, data breaches, and disruption to connected systems.

The Four Dimensions of the Model Agentic AI Framework

The Model Agentic AI Framework builds on Singapore's earlier governance frameworks and is structured around four key dimensions.

1. *Assess and Bound the Risks Up Front.* Organisations should adapt internal structures to account for agent risks, which depend on factors such as domain tolerance for error, access to sensitive data, reversibility of actions, and task complexity. Risk mitigation measures include limiting agent access to the minimum required tools and data, defining standard operating procedures, and designing offline mechanisms for malfunctions. Agents should have unique identities tied to supervising agents or users for

accountability, and threat modelling is recommended to identify security risks, including memory poisoning, tool misuse, and privilege compromise.

2. *Make Humans Meaningfully Accountable.* Agent autonomy complicates traditional responsibility assignments, and multiple actors across the agent life cycle can diffuse accountability. The Model Agentic AI Framework recommends clearly allocating responsibilities internally (across decision makers, product teams, and cybersecurity teams) and externally through contracts addressing security, performance, and data protection. “Human-in-the-loop” must be adapted to address automation bias, including defining checkpoints requiring human approval and implementing regular audits and real-time monitoring.
3. *Implement Technical Controls and Processes.* Technical controls should address planning and reasoning (logging for verification), tools (applying least privilege and limiting database write access), and protocols (whitelisting trusted servers and sandboxing code execution). Testing before deployment is essential for task accuracy, policy compliance, and robustness. Agents should be deployed gradually with continuous monitoring maintained.
4. *Enable End-User Responsibility.* Trustworthy deployment also relies on end-users using agents responsibly. Users should be informed of authorised actions, data handling practices, and their responsibilities. Transparency is key—declaring agent interactions and providing human escalation points. The Model Agentic AI Framework also highlights potential tradecraft impacts: as agents take over entry-level tasks, organisations should ensure staff retain foundational skills through adequate training.

Singapore's Broader AI Governance Ecosystem

The Model Agentic AI Framework does not exist in isolation, but forms part of Singapore's comprehensive approach to AI governance. Singapore does not have legislation specifically governing the use of AI, instead adopting an approach to AI regulation that is pragmatic, sector-specific, and use-case centric. Key frameworks that complement the new agentic guidance include the Model AI

Governance Framework (originally launched in January 2019 and updated in January 2020), the Model AI Governance Framework for Generative AI released in May 2024, and sector-specific guidelines including the FEAT (fairness, ethics, accountability, and transparency) Principles for the financial sector and the AI in Healthcare Guidelines.

AI developers operating in Singapore must also ensure compliance with the Personal Data Protection Act 2012 (PDPA) to the extent that AI systems involve the collection, use, or disclosure of personal data. The Personal Data Protection Commission's Advisory Guidelines on use of Personal Data in AI Recommendation and Decision Systems, published in March 2024, provide practical guidance on lawful data use in AI contexts.

For organisations looking to validate their AI governance practices, Singapore offers AI Verify—an AI governance-testing framework and software toolkit that validates the performance of AI systems against a set of internationally recognised principles through standardised tests. The Implementation and Self-Assessment Guide for Organisations (ISAGO) is also available as a companion guide to help organisations assess the alignment of their AI governance processes with Singapore's frameworks.

Practical Steps for AI Developers Entering Singapore

AI developers seeking to enter and operate in Singapore's AI sector should take the following practical steps to align with the Model Agentive AI Framework and broader regulatory expectations:

- *Conduct Comprehensive Risk Assessments.* Before deploying agentive AI, organisations should systematically identify risks by considering factors such as domain tolerance for error, access to sensitive data, scope and reversibility of actions, level of autonomy, and task complexity. Risk assessment should be ongoing, with the threat model regularly updated.
- *Establish Clear Governance Structures.* Organisations should clearly define the responsibilities of different stakeholders, both within the organisation and with external vendors. This includes establishing chains of accountability and

emphasising adaptive governance so that the organisation can quickly respond to new developments.

- *Design Meaningful Human Oversight.* Define significant checkpoints or action boundaries that require human approval, especially before sensitive actions are executed, including high-stakes actions, irreversible actions, and outlier or atypical behaviour. Implement training for human overseers to identify common failure modes such as inconsistent agent reasoning and agents referring to outdated policies.
- *Implement Robust Technical Controls.* Apply the principle of least privilege to limit tools available to each agent, enforced through robust authentication and authorisation. Ensure that agents are tested for baseline safety and reliability before deployment, and continuously monitored after deployment.
- *Ensure Data Protection Compliance.* Comply with the PDPA when processing personal data through AI systems, including obtaining meaningful consent, practising data minimisation, and ensuring data accuracy.
- *Leverage Available Resources.* Utilise toolkits and resources made publicly available by government bodies, such as AI Verify and ISAGO, to assess alignment with Singapore's AI governance expectations.

Conclusion

Singapore's Model AI Governance Framework for Agentic AI represents a significant milestone in the country's ongoing efforts to balance innovation with responsible AI deployment. Described as a living document, it is developed in collaboration with government agencies and industry stakeholders, and will be continuously updated to keep pace with new developments. Organisations are invited to submit feedback to refine the framework and case studies demonstrating how the framework can be applied for responsible agentic deployment.

As agentic AI continues to evolve, AI developers should remain vigilant, be adaptive, and closely monitor regulatory updates and emerging best practices to ensure continued compliance and build and maintain public trust in their AI systems and products. With

Singapore's strategic focus on AI transformation and its comprehensive governance ecosystem, the city-state offers a compelling environment for AI developers seeking to deploy cutting-edge technology responsibly.

Notes

* The authors, attorneys with Mayer Brown, may be contacted at amita.haylock@mayerbrown.com and justin.lai@mayerbrown.com, respectively.

1. <https://www.channelnewsasia.com/singapore/budget-2026-national-artificial-intelligence-council-ai-lawrence-wong-5925886>.

2. [https://www.iras.gov.sg/schemes/disbursement-schemes/enterprise-innovation-scheme-\(eis\)](https://www.iras.gov.sg/schemes/disbursement-schemes/enterprise-innovation-scheme-(eis)).

3. [https://www.mddi.gov.sg/newsroom/singapore-launches-new-model-ai-governance-framework-for-agentic-ai-/-](https://www.mddi.gov.sg/newsroom/singapore-launches-new-model-ai-governance-framework-for-agentic-ai-/).