

Artificial Intelligence in Singapore

 imda.gov.sg/about-imda/emerging-technologies-and-research/artificial-intelligence

About Artificial Intelligence

Artificial Intelligence (AI) refers to the study and use of intelligent machine learning to mimic human action and thought. With the availability of Big Data, advances in computing, and the invention of new algorithms, AI has emerged as a disruptive technology in recent years. Singapore is at the forefront of leveraging AI and machine learning to drive innovation and growth in a range of industries.

Whether it's improving healthcare outcomes or enhancing banking access for seniors, AI is already transforming how Singapore lives, works, and connects. Across sectors, businesses and institutions are adopting AI to spark innovation, boost efficiency, and elevate experiences.

Using AI Responsibly

To ensure safe and responsible use of AI, Singapore has led the way with global AI governance efforts through numerous products and initiatives.

Collapse all

Launch your LLM-based applications with confidence. The [Starter Kit \(3.02MB\)](#) provides clear, step-by-step guidance for pre-deployment testing of GenAI applications, consolidating emerging best practices from industry and government into practical, voluntary guidelines. It helps organisations identify relevant risks, apply robust testing methodologies, and assess whether their applications meet baseline safety and reliability. This represents an important first step towards more standardised testing and a stronger AI assurance ecosystem.

It covers five common risks: hallucination and inaccuracy, bias in decision-making, undesirable content, data leakage, and vulnerability to adversarial prompts.

CREX, an external **AI startup**, was [a successful example that leveraged our starter kit](#). They built an AI-powered sustainability reporting platform designed to help hotels in Asia and beyond accurately measure and communicate their carbon emissions.

The 2025 Singapore Conference on AI (SCAI): International Scientific Exchange on AI Safety brought together more than 100 leading AI scientists and experts across geographies to identify the key research priorities in AI safety.

The result, [The Singapore Consensus on Global AI Safety Research Priorities](#), organises AI safety research domains into three broad buckets: creating trustworthy AI systems (Development), evaluating AI systems' risks (Assessment), and monitoring and intervening after deployment (Control).

IMDA, in partnership with Humane Intelligence, conducted the world's first-ever multicultural and multilingual AI safety red teaming exercise focused on Asia-Pacific in November and December 2024. This Challenge brought together more than 350 participants from 9 countries across Asia-Pacific to red team four LLMs (Aya, Claude, Llama, SEA-LION) for cultural bias stereotypes in both English and regional languages.

Read [Executive Summary \(117.74KB\)](#) and [Evaluation Report \(1.41MB\)](#) to find out more!

Singapore released its first edition of the Model AI Governance Framework (“Model Framework”) for broader consultation, adoption and feedback on 23 January 2019. The framework provides detailed and readily implementable guidance to private sector organisations to address key ethical and governance issues when deploying AI solutions. By explaining how AI systems work, building good data accountability practices, and creating open and transparent communication, the framework aims to promote public understanding and trust in technologies.

The second edition of the Model AI Governance Framework was released on 21 January 2020.



Decisions made by AI should be
EXPLAINABLE, TRANSPARENT & FAIR



AI systems should be
HUMAN-CENTRIC



Internal Governance Structures and Measures

- Clear roles and responsibilities in your organisation
- SOPs to monitor and manage risks
- Staff training



Determining the Level of Human Involvement in AI-augmented Decision-making

- Appropriate degree of human involvement
- Minimise the risk of harm to individuals



Operations Management

- Minimise bias in data and model
- Risk-based approach to measures such as explainability, robustness, and regular tuning



Stakeholder Interaction and Communication

- Make AI policies known to users
- Allow users to provide feedback, if possible
- Make communications easy to understand

The second edition includes additional considerations (such as robustness and reproducibility) and refines the original Model Framework for greater relevance and usability. For instance, the section on customer relationship management has been expanded to include considerations on interactions and communications with a broader network of stakeholders. The second edition of the Model Framework continues to take a sector- and technology-agnostic approach that can complement sector-specific requirements and guidelines.

Access the [second edition of the Model Framework](#), and the [primer \(640.36KB\)](#) for more details.

The AI Verify Foundation (AIVF) and IMDA have published the [Model AI Governance Framework for Generative AI](#), expanding on the second edition of the Model AI Governance Framework in 2020 covering Traditional AI. This is the first comprehensive framework pulling together different strands of global conversation surrounding AI governance.

Building on previous model governance frameworks, IMDA published the [Model AI Governance Framework for Agentic AI \(1.03MB\)](#) in 2026. It provides guidance to organisations on how to deploy agents responsibly, while emphasising that humans are ultimately accountable.

The framework gives organisations a structured overview of the risks of agentic AI and emerging best practices in managing these risks. It is targeted at organisations looking to deploy agentic AI, whether through developing AI agents in-house or using third-party agentic solutions. The framework provides organisations with guidance on technical and non-technical measures they need to put in place to deploy agents responsibly, across four dimensions:

- Assessing and bounding the risks upfront by selecting appropriate agentic use cases and placing limits on agents' powers
- Making humans meaningfully accountable for agents by defining significant checkpoints at which human approval is required
- Implementing technical controls and processes throughout the agent lifecycle
- Enabling end-user responsibility through transparency and training.

Intended as a companion guide to the Model Framework, [ISAGO](#) aims to help organisations assess the alignment of their AI governance practices with the Model Framework. It also provides an extensive list of useful industry examples and practices to help organisations implement the Model Framework.

ISAGO is the result of the collaboration with World Economic Forum Centre for the Fourth Industrial Revolution to drive further AI and data innovation. The guide was developed in close consultation with the industry, with contributions from over 60 organisations.

Complementing the Model Framework and ISAGO is a Compendium of Use Cases (**Compendium**) that demonstrates how local and international organisations across different sectors and sizes implemented or aligned their AI governance practices with all sections of the Model Framework. The Compendium also illustrates how the featured organisations have

effectively put in place accountable AI governance practices and benefitted from the use of AI in their line of business. We hope these real-world use cases will inspire other companies to do the same.

[Volume 1](#) features use cases from Callsign, DBS Bank, HSBC, MSD, Ngee Ann Polytechnic, Omada Health, UCARE.AI and Visa Asia Pacific.


[Volume 2](#) contains use cases from the City of Darwin (Australia), Google, Microsoft, Taiger as well as a special section on how AI Singapore implemented our Model Framework in its 100 Experiments projects with IBM, RenalTeam, Sampo Asia Holdings and VersaFleet. Under the guidance of the Advisory Council on the Ethical Use of AI and Data, IMDA and PDPC collaborated with the Lee Kuan Yew Centre for Innovative Cities (LKYCIC), Singapore University of Technology and Design to develop Singapore's first guide on AI-driven job redesign.

Launched on 4 December 2020, the [Guide](#) offers an industry-agnostic, practical framework to help organisations understand how AI can transform existing roles — and how to prepare employees for a digital future. It provides guidance on practical steps in four areas of job redesign:



Transforming Jobs

Assessing the impact of AI on tasks, including whether each task can be automated or augmented by AI or remain in human hands, and deciding which jobs can be transformed within an appropriate time frame.



Charting clear pathways between jobs

Chart task pathways between jobs within an organisation and identify the tasks employees would need to learn to transition from one job to another.



Clearing barriers to Digital Transformation

Suggest ways to address potential challenges and support employees when implementing AI.



Enabling effective communication between employers and employees

Build a shared understanding within the organisation of “why”, “what”, and “how” AI will augment human capabilities and empower employees in their careers.

The Guide supports IMDA’s efforts to build a trusted and progressive AI environment that benefits businesses, employees and consumers. For example, the Model Framework guides organisations to deploy AI responsibly and address consumer concerns. Likewise, the Guide encourages organisations to take a human-centric approach to manage the impact of AI adoption by investing in redesigning jobs and reskilling employees. Access the [primer](#) for more details.



IMDA developed [AI Verify](#), a [AI governance testing framework](#) and software toolkit. The framework outlines 11 governance principles and aligns with international AI standards from the EU, US, and OECD. AI Verify helps organisations validate AI performance through standardised tests across principles such as transparency, explainability, reproducibility, safety, security, robustness, fairness, data governance, accountability, human agency, inclusive growth, and societal and environmental well-being.

AI Verify was developed in consultation with companies of different sizes and sectors — including AWS, DBS, Google, Meta, Microsoft, Singapore Airlines, NCS/LTA, Standard Chartered, UCARE.AI, and XOPA. It was released for international pilot in May 2022 and open-sourced in 2023, with more than 50 companies including Dell, Hitachi, and IBM participating.

With the rise of Generative AI (GenAI), the AI Verify testing framework has been enhanced to address its unique risks. It now supports testing for both Traditional and GenAI use cases.

Collapse all



As AI testing technologies continue to evolve, there is a growing need to bring together global expertise in this space. To meet that need, IMDA established the [AI Verify Foundation](#) — a not-for-profit body that harnesses the open-source community to advance responsible AI testing worldwide. The Foundation has nine premier members (AWS, Dell, Google, IBM, IMDA, Microsoft, Red Hat, Resaro and Salesforce) and more than 180 general members.

The not-for-profit Foundation will:

- Foster a community to contribute to the use and development of AI testing frameworks, code base, standards, and best practices
- Create a neutral platform for open collaboration and idea-sharing on testing and governing AI

- Nurture a network of advocates for AI and drive broad adoption of AI testing through education and outreach

In February 2025, AI Verify Foundation and IMDA launched the [Global AI Assurance Pilot](#) to help codify emerging norms and best practices around technical testing of Generative AI applications. The findings and use cases have been released.

An initiative by IMDA and AI Verify Foundation, [the Sandbox](#) is a testing ground for builders and deployers of GenAI applications to get them tested by specialist technical testers.

The objectives of the Sandbox are to:

- Reduce testing-related barriers to GenAI adoption through provision of practical guidance and facilitating access to specialist testing partners.
- Provide inputs into (eventual) technical testing standards for GenAI applications
- Support the growth of a viable AI assurance market.

Read more from the [media release](#).

Building Singapore's Future with AI

At IMDA, we don't just research AI—we turn it into real-world solutions that make a difference. We develop and deploy transformative AI solutions to help local businesses improve productivity and help local enterprises thrive in the digital economy.

By developing and deploying transformative AI, we are continuously growing Singapore's capabilities and building digital trust, ensuring our enterprises are empowered to succeed.

Collapse all

About the AI Verify Toolkit

The AI Verify Toolkit (AIVT) is an open-source, extensible software solution that enables organisations to assess their AI systems against internationally recognised governance principles. Designed with both technical and procedural evaluations in mind, AIVT helps organisations demonstrate fairness, explainability, robustness, and compliance through an integrated testing framework.

Hosted on GitHub and governed by the AI Verify Foundation's open-source community, AIVT continues to evolve. Organisations benefit from shared improvements, international alignment, and practical support for responsible AI adoption.

Modular Architecture and Sector Focus

Built for flexibility, AIVT supports plugin-based extensions tailored to different sectors.

- Finance (with MAS): Developed by the Monetary Authority of Singapore (MAS), the Veritas Toolkit is now available as a plugin within AIVT. This integration allows financial institutions to assess their AI models against FEAT principles directly within the AI Verify environment, streamlining responsible AI adoption and regulatory compliance in the finance sector
- Competition and Consumer Protection (with CCCS): The Competition and Consumer Commission of Singapore (CCCS) AI Markets (AIM) Toolkit has also been integrated as a plugin in AIVT. This enables businesses to conduct targeted self-assessments of their AI models, specifically addressing anti-competitive risks and consumer protection compliance requirements within the unified AIVT platform

This toolkit is accessible via [GitHub](#).

About Project Moonshot

Project Moonshot is one of the world's first Large Language Model (LLM) evaluation toolkits. It is an open-source solution that combines benchmarking and red-teaming, enabling developers and compliance teams to test and evaluate LLMs and their applications effectively.

Integrates Seamlessly into CI/CD pipelines

The latest version of Project Moonshot can be easily integrated into CI/CD pipelines, providing automated testing for a range of safety risks, including hallucinations, undesirable content, data disclosure, and susceptibility to adversarial prompts. This enables organisations to identify and address vulnerabilities proactively, ensuring AI applications are safer before deployment.

Whether you're a CI/CD developer, application owner, or compliance officer, Moonshot helps you ensure your AI applications meet the highest standards of safety and reliability.

This open-source tool is hosted on [GitHub](#).

The new GPT-Legal Q&A model, developed by IMDA together with SAL, powers the generative AI-powered search engine on LawNet. The system is trained on Singapore's legal context and supported by vast amounts of authoritative data, including judgments, Singapore Law Reports (SLRs), legislation, and textbooks. It delivers real-time, precisely referenced answers to complex legal research questions – starting with contract law. It addresses various legal question archetypes, from factual queries to case-specific and hypothetical scenarios, providing nuanced and context-aware insights that benefit over 10,000 LawNet users, representing 70% of Singapore's legal fraternity.

GPT-Legal Q&A transforms legal research by undertaking the meaning beyond user queries rather than relying on traditional keyword matching. Using advanced language model technology with retrieval-augmented generation (RAG), the system intelligently matches questions with relevant legal content by interpreting context and semantic meaning. This specialised approach delivers more focused and relevant responses, particularly when queries are framed using legal terminology.

The Agentic AGM Demonstrator, developed by IMDA with the Singapore Academy of Law (SAL), is a proof-of-concept showcasing how agentic AI can transform corporate secretarial work. Built around the Annual General Meeting (AGM) process, the demo highlights how AI agents can autonomously plan timelines, generate and circulate statutory documents, and track signatories, while remaining adaptable to regulatory requirements. By embedding autonomy, proactiveness, memory, and tool use, the demonstrator illustrates the potential of agentic AI to offload routine compliance tasks, boost efficiency and accuracy, and address manpower challenges in the corporate services sector—starting with AGMs but scalable to wider corporate secretarial functions.

AI for Businesses

The SMEs Go Digital programme aims to help SMEs use digital technologies and build stronger digital capabilities to seize growth opportunities in the digital economy. Through [CTO-as-a-Service](#) - a one-stop platform that helps SMEs go digital anytime, anywhere – SMEs can easily access a suite of AI and GenAI tools and resources, including:

- Pre-approved [GenAI solutions](#) with up to 50% grant support
- [GenAI Navigator](#) to help SMEs find the right GenAI solution for their needs
- [Complimentary AI events and services](#) from Digital Enterprise Blueprint partners
- [GenAI for SME resource page](#) to help SMEs get started on GenAI
- Inspiring success stories of SMEs that have embraced AI and GenAI
- [Industry Digital Plans \(IDPs\)](#) offering sector-specific guides on using AI and GenAI

IMDA also develops foundational tools to accelerate AI adoption across enterprises, such as the [National Speech Corpus](#) which is a large-scale Singapore English corpus of open speech data.

Collapse all

The [GenAI Playbook](#) is designed to provide structured guidance and support for enterprises at different stages of digital maturity. For those at the early stages of exploration, the Playbook provides an evaluation framework and information to help users make informed choices when they adopt a GenAI solution. As for digitally mature enterprises that require customised solutions, the Playbook provides insights on the tech capabilities and partners they need to ensure successful implementation, and the governance considerations they need to be aware of.

IMDA has two GenAI Sandboxes to help SMEs experiment, innovate and adopt GenAI tools in business areas such as Marketing and Sales, Customer Engagement, Generative Web Design and Talent Acquisition.

Close to 30 pre-approved GenAI solutions, jointly curated with industry and technical experts, were made available for trial with grant support. Over 200 SMEs have been supported under the two sandboxes, using GenAI tools to improve productivity and cut costs. IMDA is

continuously exploring new ways to support enterprises in experimenting with GenAI. Solution providers with GenAI-powered solutions can register your interest [here](#) and stay updated on potential future initiatives.

The [GenAI x Digital Leaders initiative](#) aims to help digitally mature enterprises raise their understanding of GenAI and provide them with access to GenAI expertise and resources to develop and implement customised GenAI solutions with tech partners. Enterprises keen to participate in the initiative can indicate their interest via the [online form](#).

A podcast on generative AI and the future of work: A Singapore perspective



IMDA's Assistant Chief Executive Terence Chia joins McKinsey Singapore partner Sanjna Parasrampurua and digital strategy expert Kathryn Kuhn to explore how generative AI is reshaping work, globally and in Singapore.

The discussion dives into critical questions: How will GenAI impact jobs? What new risks and responsibilities will emerge? And how can companies build a workforce ready to harness AI's potential?

ACE Terence and Kathryn focused on generative AI's potential to unlock innovation and shared solutions to key challenges — such as expanding beyond tech talent to roles like AI risk and compliance officers, and starting AI training at school level to build a skilled pipeline through to the C-suite.

He also shared how Singapore is already tackling many of these hurdles through programmes such as IMDA's TechSkills Accelerator, and its plan to continue to do so. The conversation led to the conclusion that collaboration is vital to ensure that leadership and talent are future-ready, and that companies look to capture the significant value, rather than be distracted by the power of generative AI.





Listen in to the podcast trailer:

[Full podcast on McKinsey's website](#)

Led by IMDA, and in collaboration with SkillsFuture Singapore, Workforce Singapore, the National Trades Union Congress, and industry partners, TeSA takes an integrated approach to skills acquisition and practitioner training. It enables professionals to acquire relevant, in-demand skills for the future of tech.

Explore related tags

[Artificial intelligence](#)





Stay tuned for our newsletter

[Subscribe now](#)