



20 July 2026

# Hong Kong's emerging AI governance framework



**Sam Wu**  
Partner



**Beverly Fu**  
Associate



**T**he Hong Kong government reiterated, in early 2026, that it will continue to regulate artificial intelligence mainly through existing regimes – most notably personal data, financial regulation, copyright and consumer protection – supplemented by sectoral guidance, rather than a single AI code.

In practice, this means the primary sources of “hard law” for AI use remains the Personal Data (Privacy) Ordinance, licensing and conduct requirements under the Securities and Futures Ordinance, anti-money laundering and counter-financing of terrorism rules, and common law duties.

The incremental regulatory focus is now concentrated on governance frameworks, guidelines and supervisory tools that guide institutions in deploying AI responsibly within that existing legal architecture.

## Ethical framework

The Ethical Artificial Intelligence Framework was first developed by the Digital Policy Office (DPO) to guide government departments in planning, designing and implementing AI and big data projects, with a clear focus on incorporating ethical considerations at each stage of the AI lifecycle.

In December 2025, the blueprint was promoted more broadly as a reference for any organisation operating AI systems in Hong Kong.

The Ethical AI Framework articulates 12 ethical AI principles grouped around themes such as

transparency and interpretability, reliability, robustness and security, fairness, diversity and inclusion, human oversight, lawfulness and compliance, data privacy, safety, accountability, beneficial AI, co-operation and openness, sustainability and just transition.

Importantly for governance, the framework proposes a three lines of defence model for AI. Project teams are the first line, responsible for design, risk assessment and initial mitigation. Project steering committees and assurance teams are the second line, responsible for independent review and final approval, with an information technology board/chief information officer (with external advisers where

appropriate) as the third line, responsible for ongoing monitoring of AI strategy and risks.

## Generative AI guideline

In April 2025, the DPO released the Hong Kong Generative Artificial Intelligence Technical and Application Guideline (updated in December 2025) (generative AI guideline), which it continues to refine in 2026.

The generative AI guideline draws on work by the Hong Kong Generative AI Research and Development Centre, seeking to balance AI innovation, application and responsibility tailored to local circumstances.

The generative AI guideline identifies three main types of stakeholders: technology developers (who build and train models); service providers (including platforms that deploy generative AI in customer-facing applications); and service users (who rely on generative AI services for personal or professional purposes).

It then sets out a governance framework built on five pillars – personal data privacy, intellectual property, crime prevention, reliability and trustworthiness, and system security – against which stakeholders are expected to assess their use cases.

One of the most significant features, especially from a policy and supervisory standpoint, is a four-tier risk classification system – unacceptable, high, limited, and low risk – with corresponding regulatory strategies.

For example, “unacceptable” uses are envisaged as subject to full prohibition and legal liability; “high risk” uses call for conformity assessments, human in the loop mechanisms, and real time monitoring; “limited risk” cases emphasise self-assessment and transparency duties; and “low risk” applications may rely on lighter touch self certification.



Although the generative AI guideline does not have direct statutory effect, it provides a structure that is already being reflected in speeches and sectoral guidance across government, and that is likely to influence how regulators characterise and supervise AI-related risks going forward.

## Trusted AI in practice

On 30 January 2026, the DPO effectively updated and consolidated the Hong Kong government’s thinking on AI with a presentation titled “AI Governance in Hong Kong – Trusted AI in Practice: Governance & Third Party Risk Management”.

The presentation reiterates that AI development must be “steered by safety and driven by application”, and positions AI governance as having three main objectives: encouraging innovative AI applications; mitigating AI related risks; and promoting widespread adoption and advancement of AI in Hong Kong.

This latest guidance highlights the Ethical AI Framework and generative AI guideline as the backbone of policy, and maps them against sectoral guidance already issued or in progress – for example, judiciary guidelines, financial market policy statements, and prudential principles for AI use.

The update also gives particular prominence to third party risk management. Organisations are urged to treat AI vendors, model providers and other external partners as integral parts of their risk profile, requiring

Careful due diligence, contractual controls, ongoing monitoring and clear allocation of responsibilities.

The emphasis is on third party AI risk – for example, reliance on a foreign model provider with training data that may expose the organisation to copyright or data privacy risks, or with model updates that could introduce new biases or security issues – rather than just internal model risk.

## Implications

The current AI framework in Hong Kong should be treated as the benchmark for AI governance, with boards and senior management expected to demonstrate clear oversight, a three lines of defence

structure, and project documentation that maps AI use cases against the framework's risk tiers and governance dimensions.

In parallel, companies should tighten vendor and outsourcing arrangements to cover data protection, intellectual property, security, reliability, bias mitigation, incident reporting and audit rights.

Although the framework is non-binding, it is increasingly likely to shape what regulators, counterparties and possibly courts may regard as a reasonable standard for responsible AI deployment in Hong Kong.

---

YYC Legal LLP is in Association with East & Concord Partners (Hong Kong) Law Firm.

---


*This material has been prepared for general informational purposes only and is not intended to be relied upon as professional advice. Please contact us for specific advice.*

*First published in June 2026 YYC Legal - legal trends of China Business Law Journal.*

## KEY CONTACT



**Sam Wu**  
Partner  
samwu@east-concord.com.hk  
+852 3797 8237



**Beverly Fu**  
Associate  
beverlyfu@east-concord.com.hk  
+852 3797 6641

## Contact us



2803-2803A  
China Resources Building  
26 Harbour Road  
Wanchai, Hong Kong

Tel: +852 2816 6888  
Fax: +852 3797 3835  
<https://www.yyc-ec.com/>

© 2026 YYC Legal LLP  
All Rights Reserved