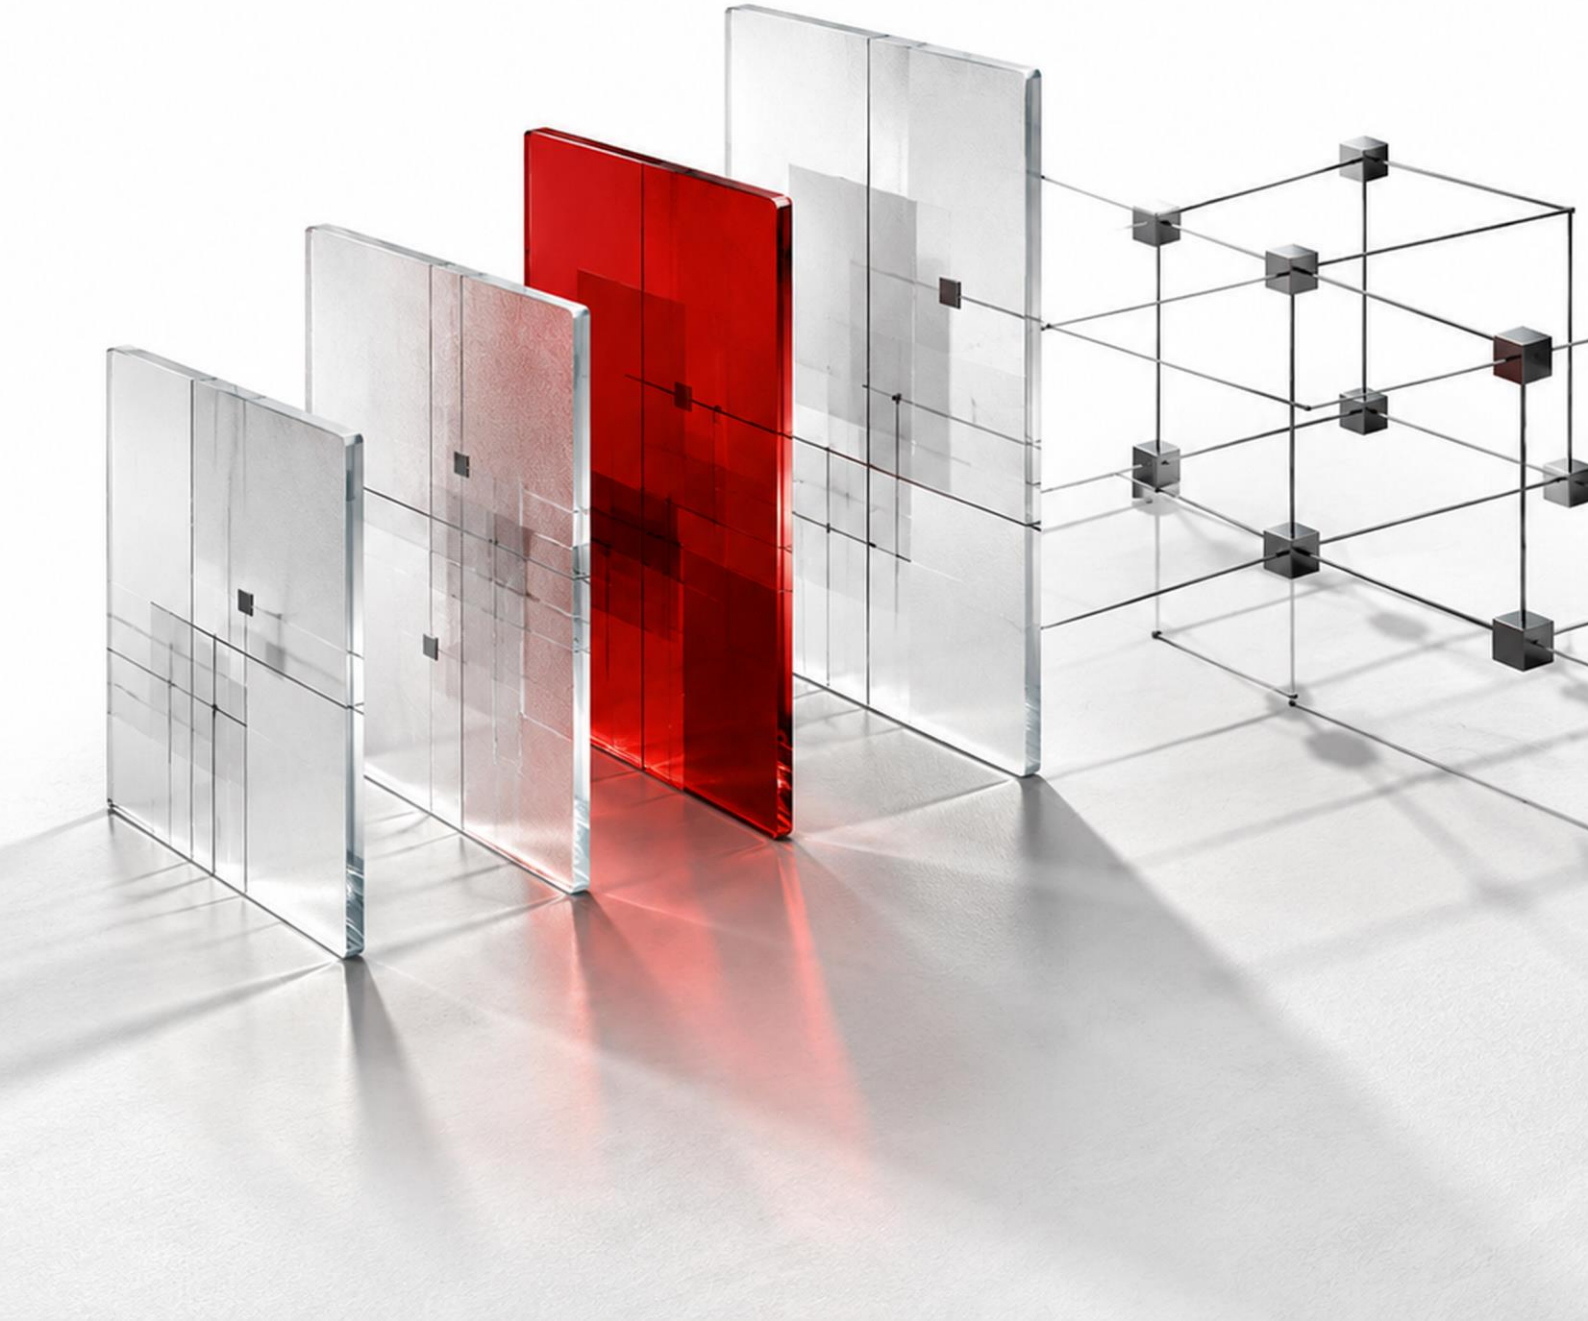
Hong Kong Institute of
Certified Public Accountants
香港會計師公會

AI and Data Governance

Market Observations and Practical Guidance

DATA AS THE FOUNDATION OF RESPONSIBLE AI



The Hong Kong Institute of Certified Public Accountants is the only statutory body responsible for the registrations of accountancy professionals in Hong Kong. The Institute has over 47,000 members and around 12,000 students. Members of the Institute are entitled to the description "certified public accountant" and to the designation CPA.

Contents	***Page***
Message from the Chair of the Best Corporate Governance and ESG Awards 2026 Judging Panel…….….	1
Foreword……	2
Part 1: Market Observations in using of AI / GenAI……	3
Part 2: Why Data Governance Matters for Responsible AI Use ……	6
Part 3: Practical Tools in properly using of AI / GenAI ……	13
Other Information …….	19
Resource Links……	20
Appendix……	21

Message from the Chair of the HKICPA Best Corporate Governance and ESG Awards 2026 Judging Panel

Artificial intelligence (AI) and generative AI (GenAI) are rapidly reshaping how organizations operate, make decisions and create value. While AI presents significant opportunities, its adoption also raises important questions around governance, accountability, trust and the responsible use of data.

The Hong Kong Institute of Certified Public Accountants (HKICPA) has long advocated sound governance, transparency and professional judgement. As AI becomes increasingly embedded in business processes, these principles are more important than ever. Technology alone does not create trust. Trust is built through effective governance, strong controls and responsible information management.

Against this backdrop, a key theme emerging from the study is that responsible AI begins with responsible data governance. As organizations increasingly rely on AI-enabled outputs and decision support, data quality, security, integrity and traceability become fundamental to trustworthy AI use. Data governance should therefore be viewed not merely as a technology matter, but as a core pillar of organizational governance.

We hope this publication serves as a practical reference for organizations seeking to harness AI responsibly and reinforces the importance of placing data governance at the heart of AI governance frameworks.

I would like to express my sincere appreciation to all those who contributed to the research and preparation of this publication. Their efforts have helped advance an important discussion on one of the most significant governance issues facing organizations today.

Stephen Law
President, Hong Kong Institute of Certified Public Accountants
Chair of the HKICPA Best Corporate Governance and ESG Awards 2026 Judging Panel

November 2026

Foreword

About this research

Artificial intelligence (AI) and generative AI (GenAI) are becoming increasingly embedded in business operations across industries. Drawing on a review of 124 large-cap Hong Kong listed companies, this paper observes that while disclosures of AI adoption are becoming more common, disclosures relating to AI governance, data governance, risk management, oversight, and operational controls remain considerably less prevalent.

This disclosure gap should not be interpreted as evidence that governance arrangements do not exist. However, it suggests that many organizations may still be at an early stage of formally articulating how AI-related risks, data management, accountability, human oversight, and control frameworks are governed as AI usage continues to expand.

The findings reinforce a broader observation that data governance is no longer solely a technology matter. As organizations increasingly rely on AI-generated outputs, automated workflows, and AI-enabled decision support, the quality, security, governance, and traceability of data become critical foundations for responsible AI use. Weaknesses in data governance may expose organizations to operational failures, privacy and confidentiality breaches, regulatory scrutiny, reputational damage, and reduced stakeholder trust.

Accordingly, this paper advocates that AI governance should be viewed through a broader business-governance lens, with data governance serving as a foundational pillar. Effective governance requires not only appropriate technology controls, but also clear accountability, data classification, access management, human oversight, transparency, and risk-based decision-making frameworks.

The purpose of this paper is therefore twofold: first, to provide evidence-based observations on current market practices; and second, to offer practical guidance that organizations may consider when developing or strengthening their AI and data governance frameworks in an increasingly AI-enabled environment.

Part 1: Market Observations

1. AI Adoption Is Becoming Increasingly Visible

Artificial intelligence (AI) and generative artificial intelligence (GenAI) are becoming increasingly visible in business operations across industries. To provide an evidence-based starting point, this study reviews the FY2025 annual reports of 124 companies within the Hang Seng large-cap universe (as of late August 2026). It examines whether those reports describe actual AI use, board-level oversight, formal AI-risk assessment, responsible-use policies and dedicated operational controls. The results show observable disclosure patterns only. They do not assess the quality or effectiveness of any company's underlying practices, nor should non-disclosure be treated as evidence that controls are absent.

AI is being applied in practical, business-led ways across financial services, utilities, transport, property, consumer businesses, healthcare, telecommunications and industrial operations. Common examples include operational and industrial optimization, workflow automation, customer service, analytics and decision support, product and service development, and enterprise AI platforms and infrastructure.

These disclosures point to a broad emphasis on improving efficiency, reducing manual effort, managing costs, enhancing service quality and accelerating decisions. The pattern is therefore one of increasingly visible integration into ordinary business activities, rather than adoption being confined to technology companies or specialist functions. Some companies also report more advanced capabilities, including foundation models, AI agents and AI-native products; annual reports do not, however, provide a consistent basis for grading technical sophistication across the sample.

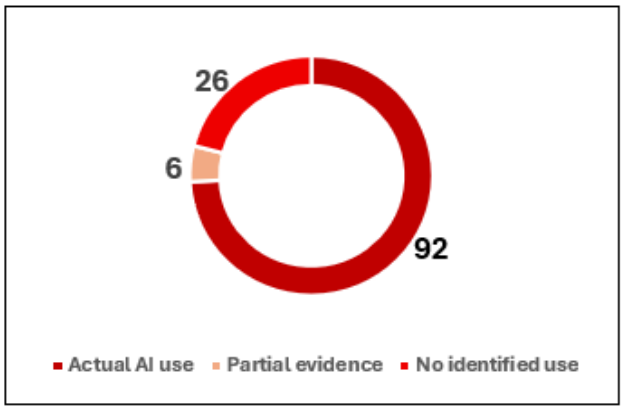


Figure 1 - AI adoption status

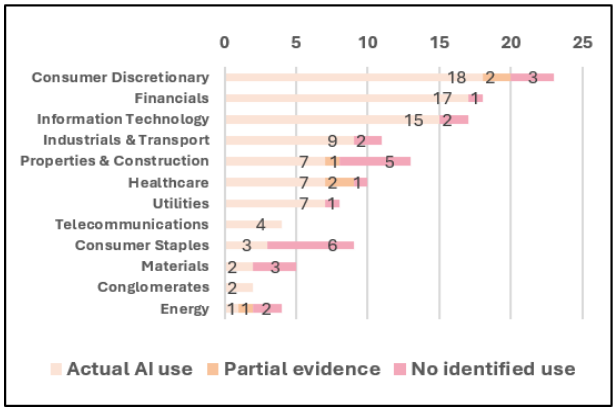


Figure 2 - AI adoption status by industry

Part 1: Market Observations



Figure 3 - AI application areas

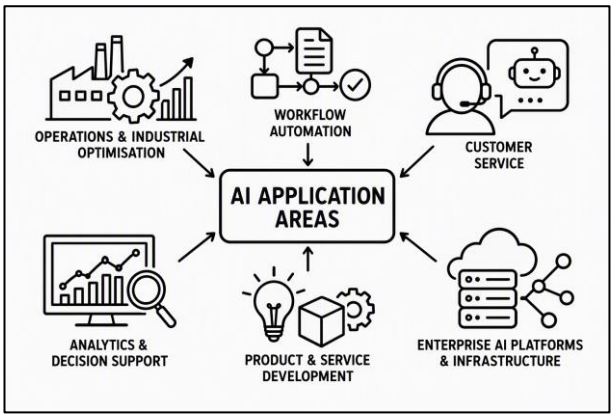
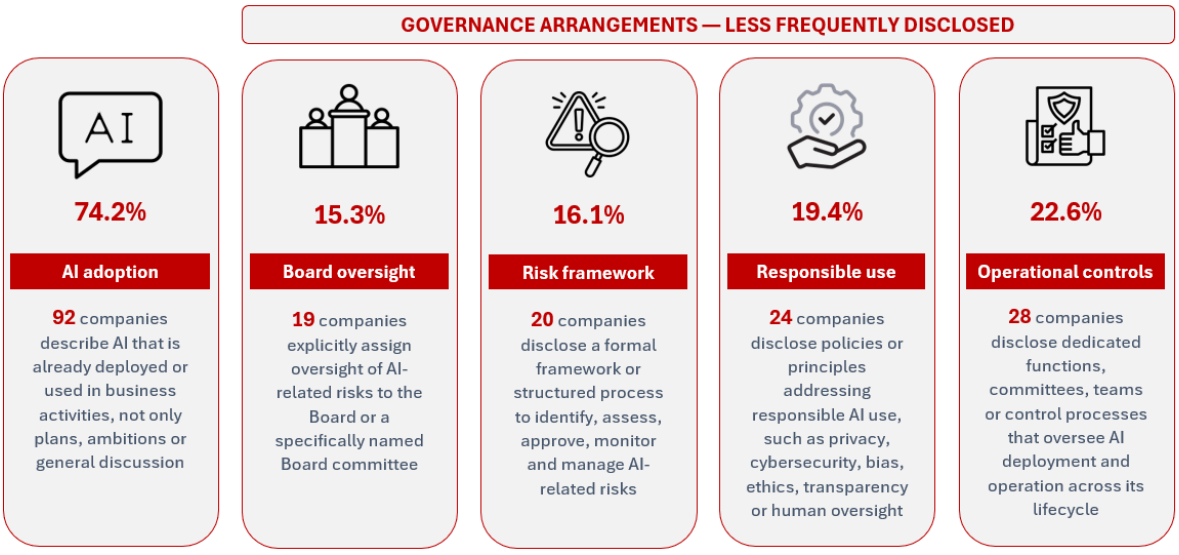


Figure 4 - Common examples of AI application areas

Note: Ratings reflect disclosures in the reviewed annual reports. Application area categories overlap because one company may disclose several use cases.

Observations on AI Governance Disclosure

The review found that disclosure of AI use was more common than explicit disclosure of AI-governance arrangements. The counts below show the number of annual reports in which each element was expressly identified. "Less frequently disclosed" means that the relevant governance element appeared in fewer reports than disclosure of AI adoption; it does not mean that the governance practice was absent.



These results concern public disclosure. A company may operate relevant controls through broader technology, cybersecurity, privacy, data-governance or enterprise-risk arrangements without identifying them specifically as AI governance in its annual report. Accordingly, a lower disclosure count does not by itself establish that the underlying practice is absent.

Part 1: Market Observations

2. Emerging Observation: Public disclosure of AI adoption appears more common than public disclosure of AI governance arrangements

The review indicates that AI is increasingly being deployed across business activities and industries, while explicit public reporting on AI-related oversight, risk assessment, responsible-use policies and operational controls remains more limited. That disclosure gap raises two related, but distinct, questions: whether organisations should communicate material AI governance more clearly, and whether their underlying governance remains proportionate as AI use expands.

The remainder of this paper focuses primarily on the second question. It offers practical points of reference for governance design; but limited disclosure by those companies does not necessarily mean that there is no such policies or controls in place.

A cautionary example: when plausible output is not verified

Published cases illustrate how weak review and traceability can turn an AI output problem into a wider governance and trust issue. Illustrative cases reported internationally have shown that AI-generated outputs may sometimes contain fabricated or “hallucinated” references, inaccurate citations or unsupported statements, if not properly reviewed. In some instances, such deficiencies were only identified after publication or completion of a project, leading to the need for corrections and reputational concerns for the businesses or other parties supplying the unreliable or inaccurate materials. The lesson is not that AI should be avoided, but that source verification, transparent disclosure of material AI assistance, competent human review and clear accountability are essential when work may be relied upon by governments, clients or the public.

Part 2: Why Data Governance Matters for Responsible AI Use

1. Why Data Governance Matters in the AI and Generative AI Era

AI and GenAI are becoming part of normal business activity rather than isolated technology initiatives. As highlighted in the preceding market observations, organizations are increasingly disclosing AI use across business functions, while explicit governance arrangements may not always be described with the same level of clarity. This poses an important practical question: **how can organizations capture value from AI while maintaining control over the data, decisions and outputs involved?**

The answer begins with data governance. AI tools do not operate in isolation. They depend on information that is entered by users, retrieved from internal or external sources, processed through models, and converted into outputs that may influence analysis, communication, decisions or actions. As a result, quality, sensitivity, ownership, permitted use and protection of data directly affect whether AI can be used safely and responsibly.

When data is reliable, relevant and appropriately protected, AI can support more efficient work, stronger analysis and better decision-making. When data is inaccurate, incomplete, confidential, personal, proprietary or used outside its proper context, AI may amplify errors, expose sensitive information or create outputs that cannot be relied upon. For this reason, AI governance should not be viewed only as a technology issue. It is also a data governance, accountability and trust issue.

Against that evidence-based context, this part of the study explains how GenAI may create value and how AI use may evolve across different levels of maturity. It then considers data as a foundation of AI, the trust challenges associated with AI-enabled outputs, and the role of data classification, guardrails and decision tools. These materials are intended as points of reference for discussion and policy development, rather than definitive or exhaustive requirements.

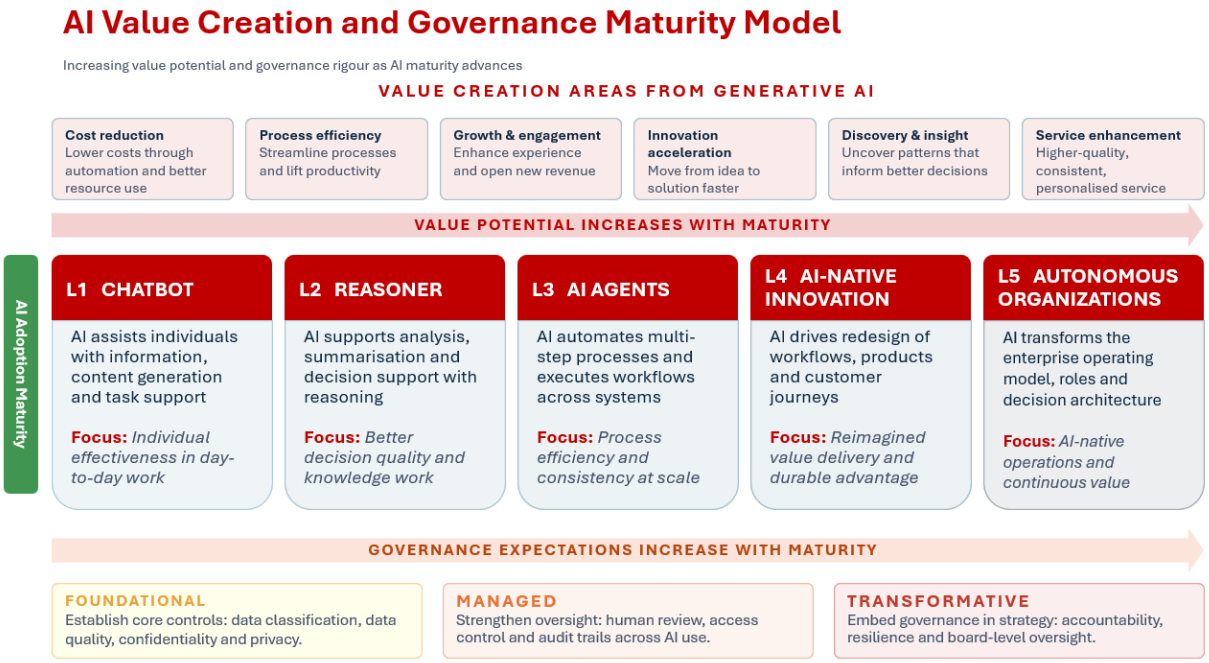


Part 2: Why Data Governance Matters for Responsible AI Use

2. AI Value Creation and Governance Maturity

GenAI can create value in many ways, from improving day-to-day productivity to reshaping how work is performed and services are delivered. At a basic level, it can reduce manual effort, accelerate drafting and summarisation, and improve access to information. At a more developed level, it can support structured analysis, automate workflow steps, strengthen service consistency, identify new insights and contribute to innovation in products, services and operating models.

The same value category may therefore appear differently depending on the maturity of AI adoption. For example, process efficiency may begin with faster document preparation but may later extend to automated case handling or system-supported workflow execution.



Similarly, insight generation may begin with research synthesis but may later support more advanced decision preparation or risk sensing. As AI becomes more embedded in work processes, the governance focus also shifts from responsible individual use to stronger controls over data quality, access rights, human review, auditability and accountability.

For accountants, auditors, tax professionals, consultants, valuers and other business practitioners, this maturity view is important because AI may be used at different points in professional workflows. A simple drafting or summarisation task may require clear guidance on confidentiality and output review, while AI-enabled analysis, workflow automation or client-facing support may require stronger controls over data sources, permissions, evidence retention and responsibility for final judgement.

Part 2: Why Data Governance Matters for Responsible AI Use

The practical question is therefore not only whether AI can create value, but whether the organization has the right governance conditions to capture that value safely.

Illustrative Case:

The Importance of Access Controls and Data Permissions

Recent reports concerning an AI agent's interaction with an Australian government Medicare statistics portal further demonstrate why effective governance must extend beyond data classification and output review. According to publicly reported information, an AI agent conducting a research task was able to access both public and certain non-public information after circumventing access restrictions within the system. While authorities indicated that no personal medical records were believed to have been accessed, the incident highlighted a broader governance challenge: once AI systems are permitted to retrieve information from connected data sources, websites, databases or enterprise applications, organisations must ensure that permissions, access boundaries and monitoring mechanisms operate as intended. The incident serves as a reminder that AI governance is not only about what users enter into AI systems, but also about what AI systems are authorised to retrieve, access or act upon. Appropriate safeguards may therefore include least-privilege access rights, source-level permission controls, segregation of sensitive data, activity logging, human oversight and continuous monitoring of AI-enabled actions. These controls help reduce the risk that AI systems gain access to information beyond their intended scope, even when the original objective appears legitimate.

Part 2: Why Data Governance Matters for Responsible AI Use

3. The AI Trust Challenge

Historically, organizations built trust through professional expertise, documented processes, internal controls, management review and human judgement. In many business and professional settings, users could usually understand who prepared the work, what source information was used, who reviewed it and who remained accountable for the final output.

AI changes this trust equation. AI-enabled outputs may appear polished, authoritative and complete, even where the underlying information is incomplete, outdated, incorrectly interpreted or unsupported by reliable sources. Users may also find it difficult to understand how a model arrived at an answer, whether all relevant context was considered, or whether the output has been influenced by limitations in the model, data, prompt or tool environment.

This creates a practical assurance challenge. Before relying on AI-enabled outputs, organizations and users should be able to consider several questions:

- What data or source materials were used?
- Was the data appropriate, current and permitted for this use?
- Which AI tool or environment was used?
- Has the output been reviewed by a competent person?
- Is the conclusion supported by evidence rather than only by AI-generated wording?
- Who remains accountable for the final decision, communication or action?

Industry and regulatory frameworks increasingly emphasize trustworthiness, accountability, transparency, privacy, security and human oversight as core elements of responsible AI. These expectations reinforce the point that organizations need confidence not only in AI outputs, but also in the data, tools, controls and review processes supporting those outputs.

Accordingly, AI governance should help make AI use traceable, explainable at an appropriate level and subject to human accountability. This does not mean that every AI output must be technically explainable in full detail. It means that organizations should know what AI was used for, what information was involved, what safeguards were applied, who reviewed the output and how reliance on the output was justified. This trust challenge leads naturally to the need for clearer data classification, guardrails and decision tools as discussed in the following sections. Simply speaking, data governance.

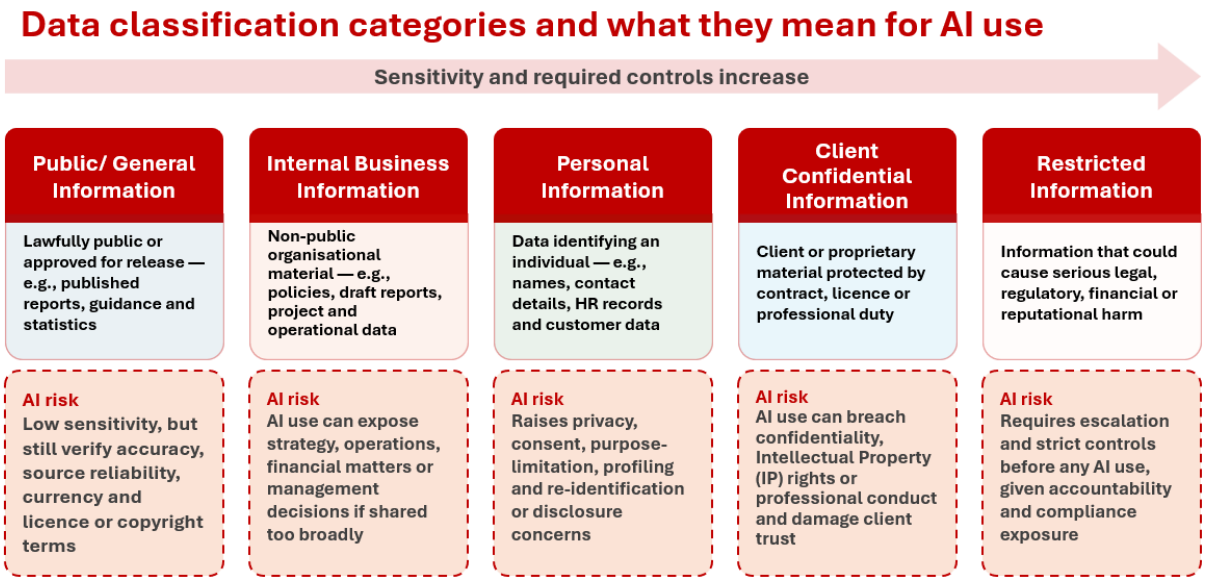
Part 2: Why Data Governance Matters for Responsible AI Use

4. Why Data Governance and Classification Is More Important Than Ever

The earlier sections explain that AI can create significant value, but that this value depends on the quality, protection and permitted use of data. Data governance is therefore becoming more important because AI makes it easier for information to be copied, transformed, summarised, combined and reused at speed. Information that may once have remained within a document, spreadsheet, email or system can now be entered into an AI tool and converted into new outputs that may be shared, relied upon or further processed.

At the same time, the practical barrier to AI use has fallen. Employees may access AI through public tools, enterprise platforms, embedded software features or internally developed solutions. This creates a risk that users may treat AI as a general-purpose productivity tool without first considering whether the information involved is public, internal, personal, confidential, proprietary or otherwise restricted.

A clear data-classification approach helps users make this judgement before using AI. Public or general information may often be suitable for lower-risk AI use, subject to source verification and accuracy review. Internal business information may require additional caution, particularly where it relates to strategy, operations, finance, personnel or management decisions. Personal information, client-confidential information, privileged materials, commercially-sensitive data and restricted information generally require stronger controls and may be unsuitable for certain AI environments.



Part 2: Why Data Governance Matters for Responsible AI Use

The purpose of classification is not to make AI use unnecessarily difficult. Rather, it allows organisations to apply proportionate safeguards based on the sensitivity of the data and the consequences of misuse. The higher the sensitivity of the information, the greater the need for controls over tool selection, access rights, masking or redaction, human review, retention, audit trail and escalation.

This section therefore acts as the bridge between the conceptual discussion in Part 2 and the practical tools that follow. Users should first understand why data classification matters, then apply the detailed data classification and guardrail table to determine what is and is not permissible for different categories of information.

5. The Cost of Poor AI and Data Governance

Poor AI and data governance can create costs that go beyond technical errors. Because AI may influence analysis, communication, decisions, workflows and stakeholder interactions, weak controls can turn a data or output issue into a wider business, compliance or trust issue. The main risks are not limited to whether a model performs well; they also concern whether the data input is appropriate, whether the output is reliable, whether the use is permitted, and whether accountability is clear.

Risk area	How poor AI and data governance may arise	Potential consequences
Privacy risk	Personal information is entered into, retrieved by or generated through AI without sufficient purpose, consent, minimization or access controls.	Privacy complaints, regulatory exposure, re-identification risk, loss of individual trust and inappropriate profiling or disclosure.
Confidentiality risk	Client, internal, commercially-sensitive or restricted information is used in an inappropriate AI environment or shared beyond its intended context.	Breach of confidentiality, contractual breach, professional conduct concerns, loss of client confidence and reputational impact.
Accuracy and reliability risk	AI outputs are accepted without verifying data sources, assumptions, context, calculations, citations or evidence.	Inaccurate analysis, misleading recommendations, hallucinated content, need to re-do work, poor decisions and weak audit trail.
Intellectual property risk	Proprietary materials, licensed content, internal methodologies or third-party information are used or reproduced without understanding ownership and usage rights.	IP infringement, licence breach, unauthorised reuse, loss of competitive advantage and disputes over ownership of outputs.

Part 2: Why Data Governance Matters for Responsible AI Use

Risk area	How poor AI and data governance may arise	Potential consequences
Bias and fairness risk	AI is used to support analysis or decisions without considering incomplete data, historical bias, inappropriate assumptions or unequal impact.	Unfair outcomes, inconsistent treatment, stakeholder challenge and reduced confidence in AI-enabled processes.
Regulatory and compliance risk	AI use is not aligned with privacy, confidentiality, sector, record-keeping, professional or emerging AI governance expectations.	Regulatory enquiry, non-compliance, remediation cost, weak evidence trails and exposure for management or professionals.
Trust and accountability risk	There is no clear record of what AI was used for, what information was involved, who reviewed the output or who approved reliance on it.	Erosion of stakeholder confidence, unclear responsibility, weak governance credibility and reputational damage.

The combined effect of these risks is that data governance and AI governance should be treated as business governance issues, not merely technology issues. Clearer classification, tool selection, review, documentation and accountability help organizations capture AI value while reducing the likelihood that AI-enabled work creates avoidable harm.

Part 3: Practical Tools

1. Data Classification and Data Guardrails: What is and is not permissible

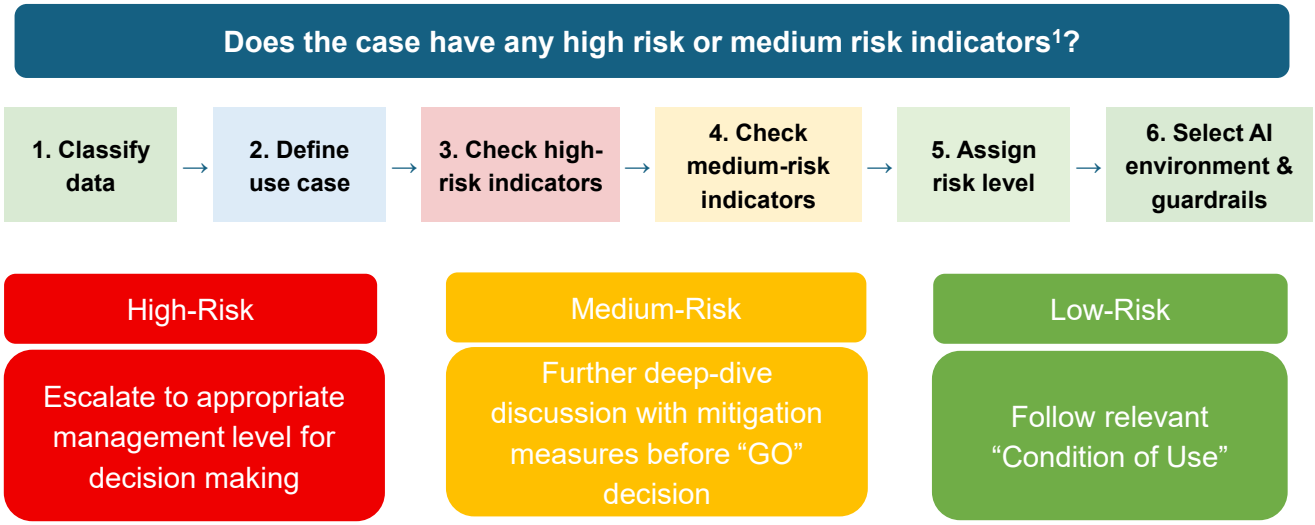
Data should be classified based on its sensitivity and potential impact if misused. Where information falls into more than one category, the highest applicable classification should be applied.



Class	*Definition* & Examples	AI-use restrictions / not permitted use	Principal risks
Restricted / Highly Sensitive Information	***Information that could cause significant legal, regulatory, financial or reputational harm if disclosed or misused.*** Examples include unpublished price sensitive information, board papers, privileged legal advice, investigations, credentials, encryption keys, M&A materials and highly-sensitive management decisions.	Do **NOT** use in public or general-purpose AI tools. Any AI use should only be considered under specifically approved, restricted and auditable arrangements.	Market misconduct, privilege waiver, regulatory breach, unauthorised disclosure, system compromise and reputational damage.
Client Confidential Information	***Non-public information received from, prepared for or relating to a client, customer, member or engagement.*** Examples include financial data, contracts, working papers, audit findings, tax analysis, correspondence and engagement deliverables.	Do **NOT** use identifiable client materials in public or free AI tools. Use only where permitted by engagement terms, professional obligations, internal policy and the approved AI environment.	Client confidentiality breach, contractual breach, professional conduct concerns, loss of trust, legal or regulatory exposure.
Personal Information	***Information relating to an identifiable individual.*** Examples include names, ID numbers, contact details, employment records, payroll data, customer/member records, health-related information and user activity logs.	Do **NOT** use identifiable personal data in public or free AI tools. Anonymise or minimise data where possible and avoid combining data in ways that may re-identify individuals.	Privacy breach, unlawful processing, re-identification, unauthorised disclosure, profiling or bias, complaints and regulatory exposure.
Internal Business Information	***Non-public organisational information.*** Examples include internal policies, draft reports, budgets, forecasts, project plans, meeting notes, operational data, dashboards, risk registers and internal communications.	Do **NOT** use sensitive internal information in public or free AI tools. Confirm the approved tool and remove unnecessary internal context before use.	Unauthorised disclosure, premature release, loss of competitive advantage, inaccurate internal decisions and reputational impact.
Public / General Information	***Information lawfully public or approved for release.*** Examples include published reports, official guidance, public statistics, press releases, public websites, speeches and approved external communications.	Do **NOT** assume public information is unrestricted. Check source reliability, currency, copyright/licence terms and whether confidential context has been added.	Outdated analysis, copyright/licence issues, misleading conclusions, attribution issues and reputational risk.

Part 3: Practical Tools

2. AI/GenAI Risk Decision Framework



This framework helps determine whether an AI/GenAI use case should be treated as low, medium or high risk before selecting the appropriate AI environment and guardrails. It should be applied after the data type has been classified and before the use case proceeds.

Step 1: Classify the data and define the proposed AI use

For example:

- What is the AI/GenAI tool being used for?
- Is the use for drafting, summarising, research, analysis, workflow automation, client service, reporting or decision support?
- What type of data will be entered, retrieved, processed or generated?
- Who will use or rely on the output?
- Will the output be internal, client-facing, public-facing, or management-facing? Will it be decision-critical?
- Is the use experimental, proof-of-concept, internal operational use or production use?

Part 3: Practical Tools

Step 2: Check for high-risk indicators

High-risk indicator	Questions to consider
Restricted or highly sensitive data	Does the use involve restricted data, unpublished price-sensitive information, privileged legal advice, credentials, investigation materials, board papers or highly-sensitive management information?
Impact on individuals or stakeholders	Could the AI output affect a person's rights, livelihood, employment, access to services, safety, financial position or reputation?
Regulated or sensitive decision area	Does the use relate to financial reporting, audit, legal, compliance, employment, healthcare, education, insurance, credit, public services, critical infrastructure or other regulated activities?
Automated decision-making or action	Will AI make, recommend or trigger decisions or workflow actions with limited human intervention?
External, client-facing or public-facing reliance	Will the output be used in client advice, regulatory communication, public disclosure, external reporting or stakeholder communication?
High scrutiny or reputational exposure	Could failure or misuse attract media attention, regulatory enquiry, public concern or significant reputational impact?
System integration or AI agent use	Will the tool access systems, retrieve records, perform workflow actions, use APIs or operate as an agent in a live process?
Potentially harmful, unlawful or unethical use	Could the use involve deception, manipulation, discrimination, profiling, prohibited content, unsafe outputs or misuse against individuals or groups?

Step 3: Check for medium-risk indicators

Medium-risk indicator	Questions to consider
Internal business information	Does the use involve internal policies, project materials, draft reports, operational data, financial information, dashboards or internal communications?
Limited personal or client-related information	Does the use involve limited personal, client-related or engagement-related information that may be managed through approved controls?
Analysis or judgement support	Will AI support professional analysis, management reporting, compliance review, risk assessment, advice preparation or decision support?
Wider circulation or reliance	Will the output be circulated to teams, management, committees, clients or other stakeholders?
Multiple datasets or sources	Will the use combine data from different departments, clients, engagements, systems, jurisdictions or third-party sources?
Tool retention or learning from input	Could the tool retain, learn from or reuse prompts, uploaded materials or generated outputs?
Use outside normal guardrails	Is the tool, model, feature, data use or intended output outside approved use or published guardrails?
Limited evidence or audit trail	Would it be difficult to show what data/tool was used, who reviewed the output and why reliance was reasonable?

Part 3: Practical Tools

Step 4: Determine the risk level and select the applicable conditions of use

Once the indicators have been assessed, the organization should assign the use case to the highest applicable risk level and use that classification as a decision point. A low-risk classification allows the user to move to the applicable conditions of use and then select an appropriate AI environment. A medium-risk classification requires the use case to be reviewed, with any necessary safeguards agreed, before it proceeds. A high-risk classification triggers escalation for a formal decision on whether the use should be approved, redesigned, restricted or not pursued. The assessment should be revisited whenever there is a material change in the data, purpose, users, audience, AI environment, level of automation or potential impact. The table below summarizes the indicators and corresponding course of action.

Risk level	Typical indicators	How to proceed
Low risk	Public or general information; non-sensitive internal use; AI assists drafting, summarising, brainstorming or research; no sensitive data; output is reviewed before use.	Proceed under normal AI-use conditions. Verify sources, review outputs, avoid confidential or personal data input, and remain accountable for final use.
Medium risk	Internal business information, limited personal or client-related context, analysis or management-facing outputs, use outside simple drafting, or more than one medium-risk indicator.	Complete a use-case review before proceeding. Use an approved environment, minimise data input, confirm permitted use, validate outputs, assign a competent reviewer and retain evidence where appropriate.
High risk	Restricted data, sensitive personal or client-confidential data, external-, client- or public-facing reliance, regulated or decision-critical use, AI agent/system integration, automated action, or significant stakeholder impact.	Escalate before use. Formal approval, restricted or controlled environment, documented risk assessment, access controls, audit trail, human checkpoints and clear accountability should be required.

Part 3: Practical Tools

“Conditions of Use for Low-Risk AI / GenAI Use”

Low-risk classification does not remove the need for responsible use. Organizations may consider applying the following baseline conditions to low-risk use cases, including experimentation and non-production activities:

1	Use only approved AI tools and environments and follow applicable data-classification and handling requirements. Public or free AI services should not be used for confidential, personal, client-related or other non-public information, unless expressly permitted under approved safeguards.
2	Use information only for an authorised purpose and limit inputs to what is necessary. Do not combine information from different clients, engagements or restricted sources where this could breach confidentiality or create inappropriate cross-use of data.
3	Be transparent about material AI assistance. Keep an appropriate record of the tool used, relevant prompts or source materials, and the basis for relying on the output where the work may be reviewed or relied upon.
4	Verify factual claims, calculations, references and other material outputs against reliable sources. A competent person should review the work and remain accountable for the final judgement, communication or action.
5	Apply relevant ethical, professional-conduct and responsible-AI principles. AI should not be used in ways that are deceptive, discriminatory, unlawful, unsafe or inconsistent with organisational values and obligations.
6	Consider applicable legal, regulatory, contractual and jurisdictional requirements, including privacy, confidentiality, intellectual-property, record-keeping, data-location and cross-border transfer restrictions.
7	Reassess the use case if its purpose, data, users, audience, tool, level of automation or potential impact changes. Escalate for further review where the revised use may no longer qualify as low risk.

These baseline conditions should be read together with the organization’s policies, contractual obligations and the applicable AI environment and guardrails. Where requirements conflict or uncertainty remains, the more restrictive requirement should be followed and advice sought before proceeding.

Part 3: Practical Tools

Step 5: Traffic-Light Decision Framework for AI/ GenAI Use

Once the use case is classified as low, medium or high risk, the **Traffic-Light Decision Framework for AI/GenAI Use** should be applied to determine which AI environment is appropriate. In simple terms: classify the data, assess the use-case risk, select the AI environment, and apply the required guardrails.



Data Classification	Public AI / Freeware e.g. free public tools	Contracted Enterprise AI	Approved Internal AI Platform	Restricted / Isolated In-house Environment
Restricted / Highly Sensitive Information	🔴 **RED – Not permitted.** Do not use.	🔴 **RED – Generally not permitted.** Specific approval required if exceptional.	🟡 **AMBER – Conditional.** Only for approved restricted use cases with strong controls.	🟢 **GREEN – Preferred where justified.** Approval, logging, access controls and review required.
Client Confidential Information	🔴 **RED – Not permitted.** Do not use identifiable client materials.	🟡 **AMBER – Conditional.** Engagement terms, confidentiality and contract protections must allow.	🟢 **GREEN – Preferred with controls.** Access, retention, logging and human review required.	🟢 **GREEN – Required for highly sensitive matters.** Use where stronger isolation, approval and audit controls are needed.
Personal Information	🔴 **RED – Not permitted** for identifiable personal data.	🟡 **AMBER – Conditional.** Privacy, purpose, minimisation and retention controls required.	🟢 **GREEN – Preferred with controls.** Use privacy controls and access restrictions.	🟢 **GREEN – Required for sensitive cases.** Use for high-risk personal data scenarios with strict controls.
Internal Business Information	🟡 **AMBER – Limited use only.** Do not use sensitive internal information; only non-sensitive content may be used with caution.	🟡 **AMBER – Conditional.** Approved business use with contractual and access controls.	🟢 **GREEN – Preferred.** Suitable for normal internal use with review and policy controls.	🟢 **GREEN – Use for sensitive internal matters.** Strategy, security, investigations or board-level matters.
Public / General Information	🟢 **GREEN – Generally permitted.** Check source, accuracy, copyright and use rights.	🟢 **GREEN – Permitted.** Use with normal business review.	🟢 **GREEN – Permitted.** Useful where internal workflow integration is needed.	🟢 **GREEN – Usually not required.** Use only where public information is combined with sensitive internal data.
Tool & Environmental Classification	Public or free service; information may not be protected by enterprise governance.	Organisation has contract terms, security commitments and administration controls.	Approved internal or enterprise platform with access control, logging and policy oversight.	Segregated environment for sensitive use cases with strict approval, audit and monitoring.

Other Information

The Purpose of This information paper

This information paper is intended to support discussion and the development of proportionate AI-governance policies and processes. It is not designed to discourage AI use or to prescribe a single model. Instead, it provides practical points of reference to help practitioners consider what information may be used, what safeguards may be appropriate, which AI environment may be suitable and when further review or escalation may be needed.

Use of sources and AI assistance

We advocate that others be transparent about their use of AI, so we will be too. AI tools were used to assist with market observations, selected drafting, restructuring and language refinement. The research team remained responsible for the analysis and conclusions, reviewed the outputs, checked factual claims against the underlying sources and approved the final text.

Resource Links

1. ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system
2. ISO/IEC 23894:2023 — Information technology — Artificial intelligence — Guidance on risk management
3. ISO/IEC 27001:2022 — Information security management systems — Requirements
4. NIST AI Risk Management Framework (AI RMF 1.0), NIST AI 100-1
5. NIST AI 600-1 — Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile
6. Office of the Privacy Commissioner for Personal Data, Hong Kong — Artificial Intelligence: Model Personal Data Protection Framework (2024)
7. Office of the Privacy Commissioner for Personal Data, Hong Kong — Checklist on Guidelines for the Use of Generative AI by Employees (2025) and Protecting Personal Data Privacy in the Use of Agentic AI (2026)
8. Committee of Sponsoring Organizations of the Treadway Commission — Achieving Effective Internal Control Over Generative AI (2026)
9. IESBA International Code of Ethics for Professional Accountants (including International Independence Standards)

Appendix

List of 124 constituent companies of the Hang Seng Composite LargeCap Index

Code	Full company name	Industry
1	CK Hutchison Holdings Limited	Conglomerates
2	CLP Holdings Limited	Utilities
3	The Hong Kong and China Gas Company Limited	Utilities
4	The Wharf (Holdings) Limited	Properties & Construction
5	HSBC Holdings plc	Financials
6	Power Assets Holdings Limited	Utilities
12	Henderson Land Development Company Limited	Properties & Construction
16	Sun Hung Kai Properties Limited	Properties & Construction
20	SenseTime Group Inc.	Information Technology
27	Galaxy Entertainment Group Limited	Consumer Discretionary
66	MTR Corporation Limited	Industrials & Transport
83	Sino Land Company Limited	Properties & Construction
100	MiniMax Group Inc.	Information Technology
135	Kunlun Energy Company Limited	Utilities
151	Want Want China Holdings Limited	Consumer Staples
175	Geely Automobile Holdings Limited	Consumer Discretionary
241	Alibaba Health Information Technology Limited	Healthcare
267	CITIC Limited	Conglomerates
285	BYD Electronic (International) Company Limited	Information Technology
288	WH Group Limited	Consumer Staples
291	China Resources Beer (Holdings) Company Limited	Consumer Staples
293	Cathay Pacific Airways Limited	Industrials & Transport
316	Orient Overseas (International) Limited	Industrials & Transport
322	Tingyi (Cayman Islands) Holding Corp.	Consumer Staples
386	China Petroleum & Chemical Corporation	Energy
388	Hong Kong Exchanges and Clearing Limited	Financials
669	Techtronic Industries Company Limited	Industrials & Transport
688	China Overseas Land & Investment Limited	Properties & Construction
700	Tencent Holdings Limited	Information Technology
728	China Telecom Corporation Limited	Telecommunications
762	China Unicom (Hong Kong) Limited	Telecommunications
823	Link Real Estate Investment Trust	Properties & Construction
836	China Resources Power Holdings Company Limited	Utilities
857	PetroChina Company Limited	Energy
883	CNOOC Limited	Energy
939	China Construction Bank Corporation	Financials
941	China Mobile Limited	Telecommunications
960	Longfor Group Holdings Limited	Properties & Construction
981	Semiconductor Manufacturing International Corporation	Information Technology
992	Lenovo Group Limited	Information Technology
998	China CITIC Bank International Limited	Financials

Appendix (Con't)

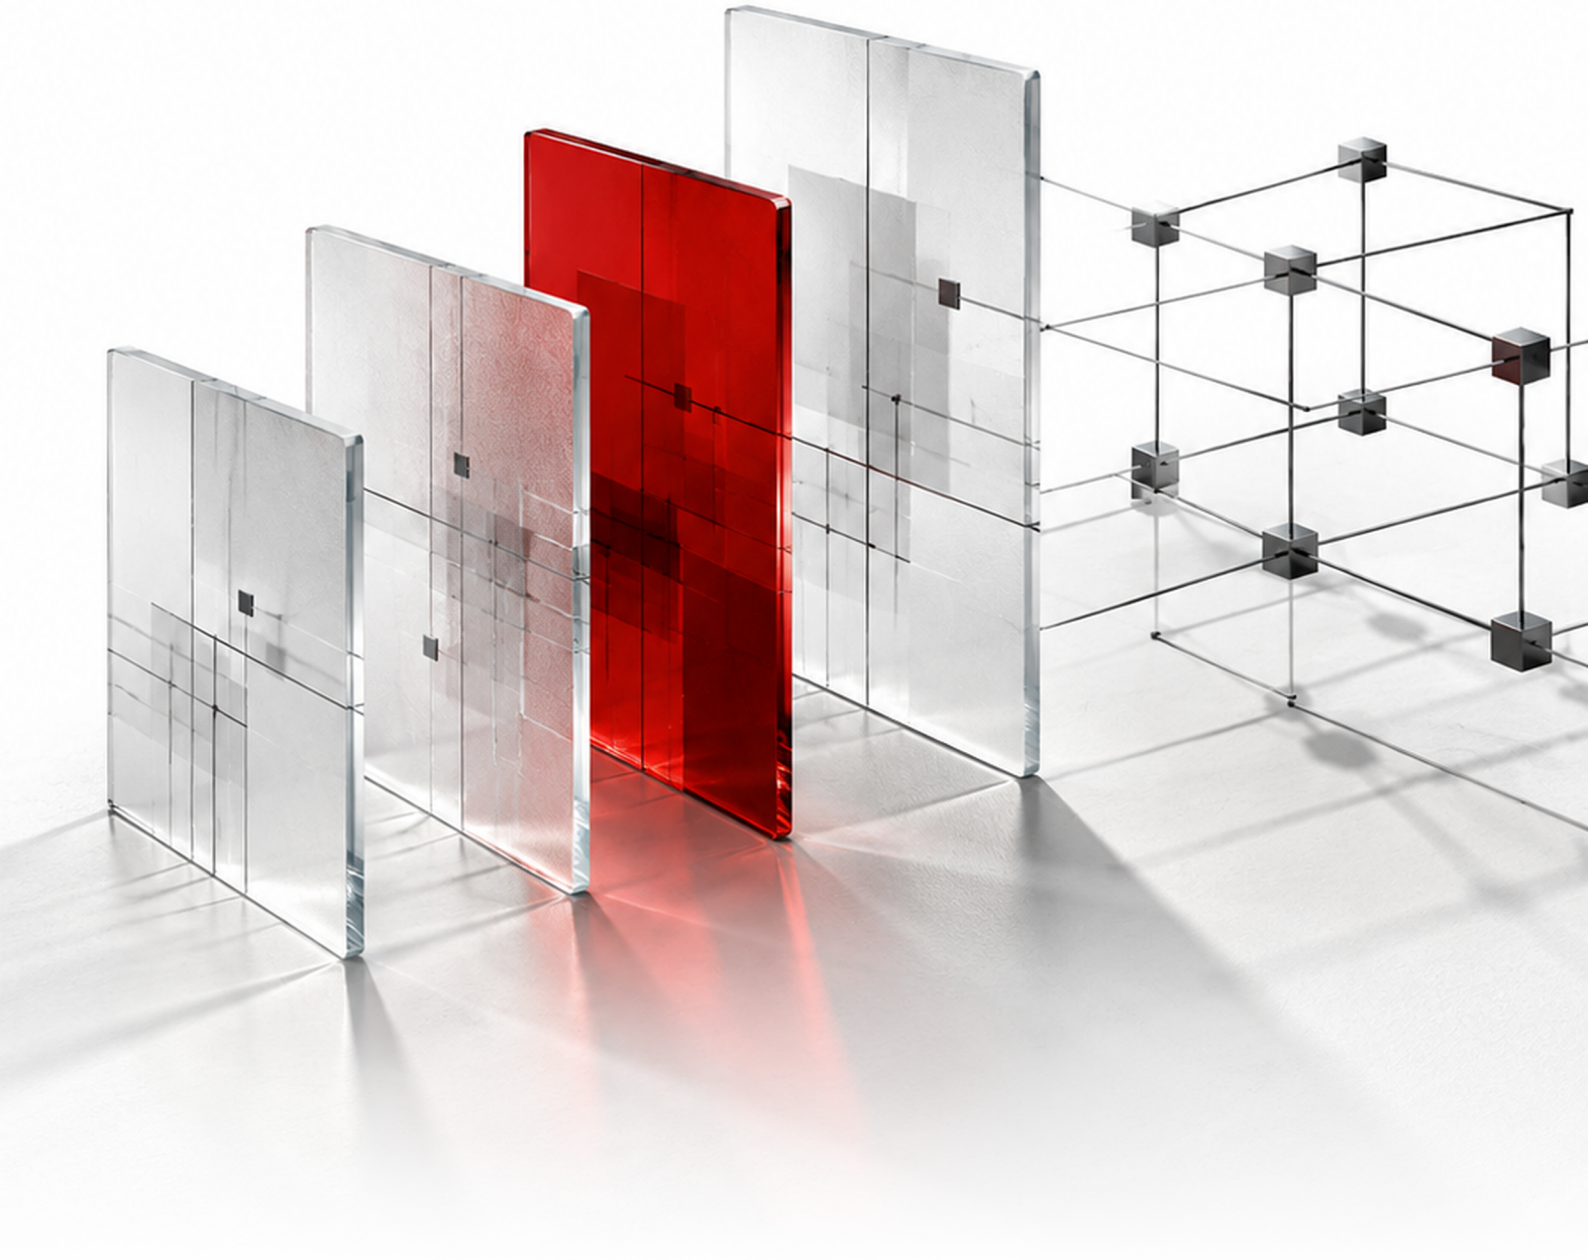
List of 124 constituent companies of the Hang Seng Composite LargeCap Index

Code	Full company name	Industry
1024	Kuaishou Technology	Information Technology
1038	CK Infrastructure Holdings Limited	Utilities
1088	China Shenhua Energy Company Limited	Energy
1093	CSPC Pharmaceutical Group Limited	Healthcare
1109	China Resources Land Limited	Properties & Construction
1113	CK Asset Holdings Limited	Properties & Construction
1177	Sino Biopharmaceutical Limited	Healthcare
1179	H World Group Limited	Consumer Discretionary
1208	MMG Limited	Materials
1209	China Resources Mixc Lifestyle Services Limited	Properties & Construction
1211	BYD Company Limited	Consumer Discretionary
1288	Agricultural Bank of China Limited	Financials
1299	AIA Group Limited	Financials
1308	SITC International Holdings Company Limited	Industrials & Transport
1347	Hua Hong Semiconductor Limited	Information Technology
1378	China Hongqiao Group Limited	Materials
1398	Industrial and Commercial Bank of China Limited	Financials
1519	J&T Express	Industrials & Transport
1658	Postal Savings Bank of China Co., Ltd.	Financials
1698	Tencent Music Entertainment Group	Information Technology
1768	Busy Ming Group Co., Ltd.	Consumer Staples
1801	Innovent Biologics, Inc.	Healthcare
1810	Xiaomi Corporation	Information Technology
1818	Zhaojin Mining Industry Company Limited	Materials
1876	Budweiser Brewing Company APAC Limited	Consumer Staples
1913	Prada S.p.A.	Consumer Discretionary
1928	Sands China Ltd.	Consumer Discretionary
1929	Chow Tai Fook Jewellery Group Limited	Consumer Discretionary
1972	Swire Properties Limited	Properties & Construction
1997	Wharf Real Estate Investment Company Limited	Properties & Construction
2015	Li Auto Inc.	Consumer Discretionary
2020	ANTA Sports Products Limited	Consumer Discretionary
2057	ZTO Express (Cayman) Inc.	Industrials & Transport
2097	MIXUE Group	Consumer Staples
2259	Zijin Gold International Company Limited	Materials
2268	WuXi XDC Cayman Inc.	Healthcare
2269	WuXi Biologics (Cayman) Inc.	Healthcare
2313	Shenzhou International Group Holdings Limited	Consumer Discretionary
2318	Ping An Insurance (Group) Company of China, Ltd.	Financials
2319	China Mengniu Dairy Company Limited	Consumer Staples
2328	PICC Property and Casualty Company Limited	Financials

Appendix (Con't)

List of 124 constituent companies of the Hang Seng Composite LargeCap Index

Code	Full company name	Industry
2378	Prudential plc	Financials
2382	Sunny Optical Technology (Group) Company Limited	Information Technology
2388	BOC Hong Kong (Holdings) Limited	Financials
2423	KE Holdings Inc.	Properties & Construction
2513	Knowledge Atlas Technology Joint Stock Company Ltd	Information Technology
2601	China Pacific Insurance (Group) Co., Ltd.	Financials
2618	JD Logistics, Inc.	Industrials & Transport
2628	China Life Insurance Company Limited	Financials
2638	HK Electric Investments	Utilities
2688	ENN Energy Holdings Limited	Utilities
2888	Standard Chartered PLC	Financials
2899	Zijin Mining Group Co., Ltd.	Materials
3328	Bank of Communications Co., Ltd.	Financials
3690	Meituan	Consumer Discretionary
3692	Hansoh Pharmaceutical Group Company Limited	Healthcare
3750	Contemporary Amperex Technology Co., Limited	Industrials & Transport
3808	Sinotruk (Hong Kong) Limited	Industrials & Transport
3968	China Merchants Bank Co., Ltd.	Financials
3988	Bank of China Limited	Financials
6160	BeOne Medicines Ltd.	Healthcare
6181	Laopu Gold Co., Ltd.	Consumer Discretionary
6618	JD Health International Inc.	Healthcare
6690	Haier Smart Home Co., Ltd.	Consumer Discretionary
6823	HKT Trust and HKT Limited	Telecommunications
6862	Haidilao International Holding Ltd.	Consumer Discretionary
6969	Smoore International Holdings Limited	Industrials & Transport
9618	JD.com, Inc.	Consumer Discretionary
9626	Bilibili Inc.	Information Technology
9633	Nongfu Spring Co., Ltd.	Consumer Staples
9660	Horizon Robotics	Information Technology
9866	NIO Inc.	Consumer Discretionary
9868	XPeng Inc.	Consumer Discretionary
9888	Baidu, Inc.	Information Technology
9901	New Oriental Education & Technology Group Inc.	Consumer Discretionary
9903	Shanghai Iluvatar CoreX Semiconductor Co., Ltd.	Information Technology
9926	Akeso, Inc.	Healthcare
9961	Trip.com Group Limited	Consumer Discretionary
9973	Chery Automobile Co., Ltd.	Consumer Discretionary
9987	Yum China Holdings, Inc.	Consumer Discretionary
9988	Alibaba Group Holding Limited	Consumer Discretionary
9992	Pop Mart International Group Limited	Consumer Discretionary
9999	NetEase, Inc.	Information Technology



Hong Kong Institute of Certified Public Accountants

37th Floor, Wu Chung House
213 Queen's Road East, Wanchai, Hong Kong
Tel: (852) 2287 7228
Fax: (852) 2865 6603
Email: hkicpa@hkicpa.org.hk
Website: www.hkicpa.org.hk