AI governance: embracing the opportunities and challenges posed by democratisation and agentic AI

Emma Keeling, Senior Knowledge Lawyer and Steve Wood, Special Advisor on Data Protection, A&O Shearman

explore the practical reality of AI implementation for organisations

One clear theme emerging from 2025 is the transformative opportunities and competitive and market pressures for organisations to rapidly implement AI. The technology evolves apace across a wide range of business functions. AI systems are now widely accessible and this democratisation, as well as a proliferation of use cases, imposes increasing demands on legal and compliance functions to facilitate responsible AI adoption quickly.

We are also now moving into an "agentic era", where AI agents, executing specific tasks traditionally performed by humans, can increasingly work together in a sophisticated way to pursue higher level objectives with greater autonomy and complexity. This provides great commercial opportunities, but may also amplify risk.

In this article, we explore the practical reality of AI implementation for organisations and how governance is evolving to support compliance in this environment, as well as the delivery of trustworthy AI that can enable business growth.

The legislative and regulatory landscape for AI

The legislative and regulatory landscape for AI globally continues to be increasingly fragmented and complex, with different approaches to legislating based on different national strategic objectives. In the US alone, at a federal level the new administration retracted the Biden AI Executive Order, but US states are taking the lead on a raft of different AI related legislative developments, with hundreds of state-level AI measures being enacted or proposed over the last year. This month we see a new Executive Order from the White House entitled 'Ensuring a National Policy Framework for Artificial Intelligence'. It seeks to ensure a unified federal approach to AI regulation and innovation, with the stated policy objective of sustaining and enhancing the US' global AI dominance through a minimally burdensome national standard. The evolution of the US AI regulatory landscape therefore remains one to watch.

Regulators are beginning to recognise the challenge of fragmented and complex regulation. They are also coming under political pressure to more actively support growth and innovation and to smooth the way to AI adoption. As such, there is something of a trend towards enablement rather than pure censure.

The EU Artificial Intelligence Act ('the EU AI Act') is in force and attention is turning to effective application, with guidance being issued and support structures being established. In tandem, we now also see the EU considering a programme of wider regulatory simplification (if not relaxation) in the form of the Digital Omnibus. The proposals span data governance, including GDPR and EU Data Act obligations, for example. With regards to the EU AI Act, amongst many other things, the European Commission has proposed a delay in the application of Chapter III regarding high-risk AI system obligations until there are standards or other support tools available (such as common specifications, harmonised standards and Commission guidelines), with a longstop date for application of 2nd December 2027 (rather than August 2026). The proposals will now work their way through the lengthy EU legislative process.

The UK is holding fire on horizontal AI regulation, continuing to adopt a 'wait and see' approach (though regulators may of course apply existing technology-agnostic regulations to AI use cases within their remits). Whilst the Data (Use and Access) Act 2025 addressed automated decision-making reforms, it did not include any substantive provisions regarding AI.

However, we do see initiatives such as the proposed AI Growth Lab (announced October 2025) looking at ways to free AI innovation from unnecessary or inappro-

priate regulatory burden. This cross-economy approach to a sandbox is looking to tackle the realities of AI development and deployment in the context of existing law and regulation.

Very targeted AI related legislation is also possible. For example, in the Crime and Policing Bill, we see proposals to empower designated organisations to scrutinise AI models and ensure that safeguards are in place to prevent them generating or proliferating child sexual abuse material.

Overlapping, conflicting or simply unclear regulation produced without AI in mind, particularly requirements such as risk and impact assessments, makes navigating the legal requirements an ongoing challenge for organisations.

However, regulators are increasingly engaging with industry and those at the front line to ensure that they stay abreast of developments, and minimise the gap between technological capability and regulatory understanding.

It is recognised that we are likely only part way through the evolution of data and other law to accommodate and appropriately regulate fast moving and transformative technology such as agentic AI. In that context and to address this complexity across multinational structures, organisations are often developing their own global AI principles and frameworks, blending compliance with the EU AI Act with standards such as the Organisation for Economic Co-operation and Development principles, National Institute of Standards and Technology risk management framework or ISO 42001.

Governance structures

Against this backdrop, to exploit the opportunities afforded by AI, effective governance will require collaboration between a wide range of business functions, including technology, data science, legal, risk, ethics, compliance and security. There will never be a single solution to the ongoing challenge of data governance and digital regulation. Organisations will need to tailor their approach to their risk profile, business and operating model, wider compliance governance, size and level of maturity, and sector-based regulatory challenges.

> ***"There will never be a single solution to the ongoing challenge of data governance and digital regulation. Organisations will need to tailor their approach to their risk profile, business and operating model, wider compliance governance, size and level of maturity, and sector-based regulatory challenges."***

Who should be responsible for AI governance?

Alongside clear links to senior management, responsibility for AI could sit with: each existing business function; the Chief Privacy Officer ('CPO'); the data protection team; or a new AI department.

In the short-term, we're observing that a centralised driver or coordinating function (which may be the CPO or privacy team) is particularly important to ensure that AI risks are being considered at each stage of the lifecycle, at each level of the business and by each relevant team. This encourages engagement in AI both horizontally and vertically across the business. Many CPO roles are being transformed into wider roles such as 'Privacy, Data Responsibility Officer' or 'Chief Privacy and Trust Officer'. The role of the coordinating function can be to drive a standard approach to AI risk assessment frameworks. For example, when AI is being deployed, the business accesses a set of standard questions covering the inherent business risks of compliance, privacy, cyber, ethics, etc. This will also help streamline initiatives, minimise parallel workflows and tackle compliance fatigue.

Some organisations may use a governance board that ensures alignment of AI use cases with the organisation's core values. It may operate in an advisory capacity but also have the ability to determine whether a project has the go-ahead. Others may look to an ethics forum that considers whether a particular action or approach is something that should be undertaken, even if feasible from a regulatory perspective.

What is clear is that siloed structures for data governance will create significant inefficiencies and risks of inconsistency. These will reduce opportunities for innovation and collaboration in finding solutions.

Decentralisation and a risk-based approach

It is becoming increasingly straightforward for business teams and employees to access, create and deploy AI systems and agents for their needs without centralised engagement, whether as stand-alone products or bolt on functionality for existing tools. Many organisations recognise the value of agile deployment in supporting innovation. Many AI agents will be self-created at the business level, for example.

Decentralisation of AI oversight is necessary to manage this democratisation as well as the broader pace of deployment. Organisations are exploring how the role of the coordinating function can dissolve into each business function, with the aim of integrating AI risk into the first line. Similarly, a risk-based, tiered approach to governance can help to prioritise and focus risk assessment against legislative requirements and the organisation's values and policies.

"Test and learn" can be another practical way to allocating, often limited resources, to mitigate AI risk. However, care is required to ensure that AI

(Continued on page 14)

(Continued from page 13)

risk categorisation does not prevent consideration of risk in different dimensions such as IP leakage. System-based controls and technical and organisational measures can also be essential to help manage risk, for example, through central management of model context protocol ('MCP') connections or limits on the ability of AI agents to access certain data stores.

Shadow AI

Democratisation of the technology inevitably risks "shadow AI", i.e. employee use of AI tools for work without organisational approval or oversight. This raises concerns about unknown unknowns.

Effective governance and record keeping helps to avoid the risks associated with shadow AI. This can include house-keeping processes, looking to remove certain dormant AI agents that have been created by colleagues for instance.

Additionally, up-to-date record keeping and clear tracking of tools and assets may offer opportunities for effective deployment, rather than simply acting as a risk mitigator. Awareness of AI usage across the organisation can allow central functions to break down silos and share that information with others. Understanding what other tools exist within the organisation, for example AI agents that could be brought together to form part of an agentic system, may support deployment and the development of additional use cases.

Separately, the risk of shadow AI could, in fact, be a driver to encourage organisations to implement official, effective tools, for use. If organisations implement the best systems for use by their people, it may help to avoid colleagues looking outside the estate for functionality.

Democratisation and vendor engagement

The democratisation of AI is recognised in the context of vendors and suppliers too. Organisations are not necessarily imposing the same stringent limits on vendor use of AI as may previously have been the case. Given the pervasive nature of the technology, many organisations' supplier restrictions are now tending to relate only to high risk or prohibited AI system use.

Standardised documentation in a democratised environment can also help triage AI systems and decentralise governance. AI related playbooks, for example the one published by the UK government in 2025, are also increasingly used and assist employees with understanding AI and the best practices to use when deploying and procuring.

Ongoing engagement

The need to address the risks of point-in-time oversight is important. If an AI tool initial falls below a threshold for in-depth engagement, governance teams should continue to ascertain whether this is true, ensuring that any further development or evolution in us has changed the risk status.

There are clear benefits to monitoring and observation, particularly given the potential for scope creep, concept drift and bias amplification (for instance) in the context of AI agents. Monitoring compliance with governance frameworks means that controls can be introduced or adapted, perhaps with governance board review. Understanding challenges arising during deployment in one space may also help address other root cause issues.

Scalability

The need for a strategy and a common framework that is scalable is also key. Expanding upon existing approaches to data governance and accounting for existing standards and compliance requirements may support effective AI governance methods. Integrating governance workflows that follow uniform approaches can also enable innovation at scale.

Systemising governance and building in data capture opportunities to existing platforms and systems is a way to gather records and information about AI agents and tools from the outset. This can support audit, analysis and ultimately accountability.

What does human oversight look like in the context of AI agents?

Whilst "human in the loop" is understood as an approach to enable AI oversight, manage hallucination risk, and support data protection compliance, amongst other things, it sits at odds with the benefits offered by AI agents. However, if anything, the potential for enhanced risks in the context of AI agents (e.g. the risk of goal misalignment, compounding of errors, amplification of bias, loss of data and confidentiality, unauthorised access, explainability and transparency concerns amongst others) means that it is necessary to consider how human oversight and governance could be introduced to an AI agent scenario, whether that is through a human on the loop, human before the loop or human after the loop approach.

Trigger points and conditions for action may need to be introduced to the AI agent cycle to enable adequate checks and balances. For example, if a particular category of individual was impacted by the actions of an AI agent, the AI agent may bring in a human decision phase. If the AI agent requires access to additional databases, human authorisation may be necessary first. It would be impractical to build in human engagement at every decision point and so clear, specific parameters defining action will also be necessary.

As automated decision making becomes more prevalent, for example in the context of recruitment, the needs of human oversight are more complex. The use of regulatory sandboxes such as that of the UK Information Commissioner's Office could be helpful in these circumstances.

Key Performance Indicators ('KPIs')

"Speed to deployment" is recognised as a KPI measure of interest for leadership and one that particularly impacts the approach taken by legal and compliance teams, incentivising a

more pragmatic approach to risk.

Other more value-based ROI measures can be seen as being harder to assess, often requiring business team input. That said, KPIs measuring AI adoption rate, decision-making effectiveness and customer satisfaction for example are possible and may support investment in governance processes.

It is worth noting that KPIs are perhaps a threshold indicator of the success or otherwise of a process, and an ongoing process of assessment and monitoring of AI governance programmes is likely to be a more informative means of measuring efficacy.

Conclusion

To take advantage of the transformative opportunities offered by AI, including the use of AI agents, organisations are continuing to evolve their governance structure, controls, policies and processes.

While centralised frameworks remain crucial, a risk-based approach is vital to enable agility across organisations and avoid unnecessary steps that can slow down experimentation and use of AI tools in the day to day. Many of the risks posed by agentic AI are not new but the step change in autonomous actions across supply chains will significantly exacerbate the nature of the risks and how they can impact on people, and ultimately compliance and trust in the digital services and products offered.

Organisations will need to strategically understand how their AI risk profile is changing and how their AI governance flexes to address the key risks.

Emma Keeling and Steve Wood
A&O Shearman
emma.keeling@aoshearman.com
steve.wood@aoshearman.com