Using AI to support governance effectiveness

A practical guide for governance professionals working in further and higher education

Lead author: Rhona Morris, Edge Hill University
Supporting authors: Kath Wass, University of Bradford and Association of Higher Education Professionals & Dan Tinkler, Advance HE

Contents

Preface

This project emerged from a single question posted in an email group asking whether anyone had started to use AI as part of their governance. At that time, widespread use of AI was new and discussion within providers tended to focus on deterring academic malpractice.

From that email, a community of practice emerged, beginning with 12 early explorers and growing to 84 curious colleagues. This group, the Using AI in Governance forum, has proved fertile ground for colleagues to come together to share their interests, doubts and learnings. We have heard from experts sharing insights on topics varying from potential AI uses in governance functions to the risks and flaws built into AI models. However, much of our sessions have been dedicated to learning from one another, and sharing lessons learned from early forays into the technology and templates to support and shape the development of best practice.

It is fair to say that the group represents a wide range of perspectives. When we started, a survey indicated that most were nervous, but curious, about the emerging technology, some were reluctant to pursue a tool which could, in theory, threaten jobs or weaken personal skills and a very small minority were already experimenting seriously. Whilst we all have more practical experience since then, group discussions still bring together a broad range of views.

This diversity was pivotal in realising that a series of Practical Principles would be helpful to governance professionals across the FE and HE sectors. It has been developed to outline good practice in AI use, regardless of confidence, product or experience. This document, actively shaped by discussions held in the Group, is intended to be an aide mémoire, not explaining how to *use* AI but how to *work with* AI in a responsible and sustainable way. We know that the technology will continue to evolve but hope that this guide, along with the underpinning Code of Practice, will prove durable in supporting sound governance practice.

As a governance professional working in HE myself, this is guidance which I am looking forward to applying within my own work. The Group has already helped shift my own practice from naïve experimentation to an approach through which I aim to use AI to grow my skills rather than replace them. Applying the Principles and adhering to the Code of Practice will help me to take this further, giving me confidence to ask questions of myself and others, to act as an ambassador and for this work to be underpinned by ethics and information security.

In producing this Guidance, I am indebted to the skills and collaboration of talented colleagues working in FE and HE across the UK. Particular gratitude is owed to Daniel Tinkler (Advance HE) and Kath Wass (University of Bradford & AHEP), who have been instrumental in keeping the project on track and developing the product we have today. Finally, I want to acknowledge the numerous AI tools which have been used by the Group and in the production of this publication: to experiment, to test thought processes, to build templates, to transcribe meetings and to turn bulletpoints into paragraphs.

Have we used AI in the drafting of this guidance? Absolutely (with human oversight)! But these opening lines have been all my own.

Rhona Morris, Edge Hill University

Chair, Using AI in Governance

Foreword

The governance of higher and further education is undergoing a profound period of change. Institutions are navigating an increasingly complex landscape shaped by rapid technological advancement, heightened regulatory expectations, shifting societal demands, and evolving models of accountability. Within this context, the emergence and adoption of artificial intelligence represents both a significant opportunity and a material governance challenge for governance professionals - one that requires not just new policies, but new ways of thinking and practising governance.

Recent developments in sector guidance have reinforced this shift. The new Committee of University Chairs (CUC) Higher Education Code of Governance (2026) places strong emphasis on effectiveness, adaptability and demonstrable impact - highlighting that good governance is not static, but contingent on behaviours, culture and informed judgement. In parallel, the Advance HE Big Conversation: Shaping the Future of HE Governance has underscored the growing expectation that governance must become more agile, more inclusive and more connected to real-world complexity, rather than relying solely on established frameworks and formal structures.

This guide is situated within that evolving context. It recognises that while codes, principles and institutional strategies are essential, governance effectiveness is enabled and enhanced through the everyday actions of those working within governance systems. As governance professionals and practitioners, we translate intent into practice through challenge, scrutiny, advice and the quality of questions they ask, to design effective structures and systems.

Artificial intelligence (AI) brings this reality into sharp focus. AI systems can influence decisions, shape outcomes, and introduce new forms of risk - often in ways that are not fully visible or easily understood. Traditional governance approaches, if applied passively, risk being outpaced by the speed and complexity of these developments. As such, the governance of AI cannot rely solely on top-down controls or technical frameworks. It must also be grounded in strong professional curiosity, ethical awareness and practical engagement at all levels of the organisation.

What distinguishes this guide is its deliberately practitioner focused perspective. Rather than focusing exclusively on what institutions should do, it explores what individuals can do day-to-day to strengthen governance outcomes. It reflects a broader shift in the sector: from governance as oversight, to governance as active practice; from compliance as process, to compliance as embedded behaviours; and from frameworks as endpoints, to frameworks as enablers of informed judgement.

This approach aligns closely with the direction of travel set out in both the CUC Code and the Advance HE work: that governance must be dynamic, values-driven and rooted in real decision-making. It also recognises that in dynamic areas such as AI - where uncertainty, risk, and innovation intersect - effective governance depends as much on mindset and behaviour as it does on formal structures.

The principles and practices set out in this guide are therefore not intended to replace existing governance frameworks, but to strengthen their application. They have been co-created with governance professionals and are designed to support colleagues supporting governance in HE and FE in becoming more confident, more questioning and more connected to the practical realities of AI use within their institutions. In doing so, they contribute to a wider ambition: to position governance not only as a mechanism of assurance, but as a source of insight, value and trust.

The production of this guidance has involved the active contributions of talented colleagues drawn together from across the sector by a curiosity about AI and the opportunities and risks it brings. This guide invites governance professionals to continue to play an active role in shaping that future.

Introduction

This project began with a simple but pressing question: how can governance professionals adapt to a world increasingly shaped by AI?

Effective AI governance is not defined solely by frameworks, policies or institutional strategies. In practice, it is driven by how individuals exercise judgement, challenge decisions and embed good habits into everyday processes - just as they do across all areas of governance. In higher and further education, where decisions are often made through collective governance structures and are subject to significant public scrutiny, these considerations have particular importance.

This guide reframes seven core AI governance principles through a practitioner-focused lens, with an emphasis on:

+ how governance professionals can adapt their own practice
+ empowering practitioners to explore emerging technologies and pursue professional development
+ taking small, practical actions that strengthen governance over time
+ identifying risks earlier through improved behaviours
+ positioning governance as a value-adding discipline
+ sharing ideas for how AI can be used to support governance professionals in their work whilst still enabling professional growth.

Where possible, we try to illustrate the principles with case studies drawn from members of the Using AI in Governance Group to demonstrate the lessons learned, both positive and negative, to date.

Given the scale and pace of change in this field, clear boundaries have been necessary to ensure that this guidance does not become obsolete immediately. Therefore, this guide does not:

+ focus on any single AI model or function (eg Chat Bots, Large Language Models, meeting pack products), recognising that the technology will continue to evolve and produce new opportunities and risks
+ provide institutional or board-level recommendations, instead concentrating on what is within the control of individual governance professionals
+ offer technical instructions on how to use AI, given that these will likely be specific to the product used and will change as the technology progresses.

By taking this approach, the guide is intended to provide enduring and practical support. The Code of Practice presented at the end is designed as a personal charter to inform and guide future individual professional practice.

1 Commit to ongoing professional development

Effective governance in an AI-enabled environment begins with individual capability. While deep technical expertise is not required, governance professionals must feel confident engaging with AI-driven decisions. In dynamic environments where training is not yet available, or at least not targeted to our specific needs, governance professionals can engage in self-development by seeking webinars, reaching out to colleagues, asking informed questions, recognising limitations and understanding where risks may arise.

Developing this capability and confidence is an ongoing process rather than a one-off exercise. It sits in the habits you build, the curiosity you maintain, and the conversations you initiate within your team. And through individual development, we also build institutional resilience by raising awareness of both opportunities and potential issues.

Embedding this in practice

At its core, building literacy is about developing a questioning mindset. This might include pausing to ask: *"What assumptions is this AI making?"* or *"What is this output based on?"* - particularly where decisions rely on automated or opaque systems.

Practical steps to support this could include:

+ maintaining a simple log of interesting AI use cases encountered in your work – what could you learn from others' mistakes? What would you use again or share with others?
+ setting aside regular, informal learning time (for example, 30 minutes each month)
+ building working knowledge of key distinctions, such as AI versus automation, and the information security implications of different tools
+ considering how AI can be embedded in the institution's systems in less obvious ways eg within student information systems, VLEs or third-party platforms, and how this may need governance consideration
+ stay alert to any ethical considerations and best practices in the use of AI and embed these learnings into your own processes and practice
+ creating space to share insights, questions or lessons learned across your governance team or wider network.
+

Over time, these small actions build confidence and strengthen collective understanding.

Risks to be aware of

Without sufficient literacy, there is a risk that AI outputs are accepted without appropriate challenge, particularly where systems appear authoritative but operate as ‘black boxes’ (ie it is impossible to trace the rationale behind the decision).

A lack of understanding of how AI tools operate and their potential flaws can also mean early warning signs are missed, limiting the ability of governance professionals to identify and respond to emerging risks.

Opportunities this creates

As literacy develops, so too does the ability to contribute more meaningfully and position governance as a strategic enabler to unlock AI benefits. Governance professionals can move from passive review to trusted advisors, helping to educate others on the proper use of emerging technologies.

AI literacy brings additional benefits for the governance professional’s own work. Harnessing the opportunities presented by AI tools could enable practitioners to enhance the depth of their understanding of complex issues, support committees to improve the quality of scrutiny and add value to decision-making processes.

Considerations

Building this capability requires time and sustained attention. It can feel complex at first, and the pace of technological change means knowledge will continue to evolve. However, the long-term benefits: stronger judgement, increased credibility and more proactive governance, make this a foundational investment.

2 Connect AI to real decisions

AI governance becomes most effective when it is grounded in the reality of organisational decision-making, including the work of boards, governing bodies and committees. When AI is treated as abstract, or discussed in general terms without a clear link to outcomes, it risks becoming disconnected from the very issues governance is intended to oversee. As its use is becoming more mainstream, there is also the potential that AI practice will become 'scattergun', with colleagues and trustees creating uses to match the tools, rather than to add value to pre-existing aims.

For governance professionals, the critical task is to consistently connect AI use to real decisions, priorities and risks, such as those relating to student outcomes, financial sustainability, estates planning or widening participation. This ensures that governance remains relevant, proportionate and focused on impact rather than process.

Embedding this in practice

A practical starting point is to bring discussions of AI back to fundamentals. When reviewing proposals or initiatives, this may involve asking simple but clarifying questions such as:

+ what decision is this AI influencing?
+ what outcome is this intended to improve?
+ how does this support our statutory duties and regulatory obligations?

Supporting this approach could include:

+ linking each AI use recommendation to a clearly defined business outcome, regulatory requirement, strategic aim or identified risk
+ challenging instances where AI is being introduced without a clear purpose or rationale
+ introduce simple mechanisms to weigh the potential risks and benefits of AI adoption for informed decision making
+ diarise time periodically, as an individual or team, to look back and evaluate AI use. It is often easier to see whether value has been realised once you are at a distance from the event.

These actions help ensure that AI is understood not as a standalone innovation, but as part of broader organisational activity.

Risks to be aware of

Where AI is not clearly connected to decision-making, there is a risk that it is deployed without a defined purpose, leading to inefficiencies or unintended consequences. This can also result in misaligned priorities across teams, or governance activity that is disconnected from where real impact is occurring.

There is also a more subtle risk of overconfidence in AI outputs. Where AI is perceived as authoritative, the 'common sense' judgement and professional challenge typically offered by governance professionals may be overlooked, or not voiced at all.

Maintaining appropriate human involvement at each stage of governance processes is therefore critical, both to ensure meaningful challenge and to guard against complacency.

Opportunities this creates

By maintaining a clear line of sight between AI and decision-making, governance professionals can strengthen transparency and accountability. This approach supports more focused use of resources, improves organisational coherence, and ensures that governance activity is seen as practical and enabling rather than abstract.

In doing so, governance becomes more clearly positioned as a driver of value, supporting better decisions rather than simply overseeing process.

Considerations

This approach requires persistence. It may involve challenging proposals that are unclear or underdeveloped and at times may be perceived as slowing innovation. Its effectiveness is also influenced by the quality of underlying business cases and organisational clarity of purpose.

However, by consistently anchoring AI in real decisions and outcomes, governance professionals can help ensure that innovation is meaningful, aligned and sustainable.

3 Make compliance practical

As we all know, compliance is most effective when it is integrated into day-to-day activity rather than treated as a separate or retrospective exercise. When embedded within routine workflows, it becomes part of how decisions are made. This in turn supports consistency, transparency and confidence across governance processes.

When working within such a changeable environment, traditional compliance processes may no longer be sufficient. For governance professionals, this means moving away from periodic or reactive checks, and towards a model where compliance is continuously considered and evidenced as AI work progresses.

Embedding this in practice

In practical terms, this begins with making compliance visible and manageable within everyday activity. This might involve using simple prompts or tools to guide thinking, rather than relying on complex or highly formalised processes.

Examples include:

+ keeping abreast of relevant laws, regulations and sector-wide developments relative to your use of AI and embedding these into your processes (eg data protection legislation, equality duties)
+ using a straightforward checklist when reviewing AI-related initiatives, covering areas such as data usage, bias, and explainability. This should include consideration of obligations under the Public Sector Equality Duty and any student-facing impacts
+ asking clarifying questions including whether any external assurance or reporting may be required
+ encouraging teams to document decisions as they are made, rather than retrospectively, to support transparency, understanding and assurance
+ sharing examples of both strong and weaker compliance practices to build collective awareness and learning
+ modelling good practice across your institution
+ being alert to emerging risks in practice (eg uninvited AI agents joining meetings) and developing appropriate mitigations
+ following your institutional policy or, where this is not yet available, developing a local policy to guide appropriate use.

Taken together, these approaches help to normalise compliance as part of everyday governance behaviour.

Risks to be aware of

The risks in AI governance and compliance are the same as in all other aspects of our work, meaning that it can be approached in much the same way. Where compliance is treated as a formality, there is a risk that it becomes a ‘box-ticking’ exercise without meaningful oversight or input. This can lead to inconsistent application of requirements, or in some cases, regulatory obligations being overlooked entirely.

A lack of embedded practice may also mean that issues only emerge late in the process, increasing the likelihood of rework or exposure.

Opportunities this creates

Embedding compliance within daily workflows can significantly strengthen confidence in AI governance processes. It supports earlier identification of issues, reduces the need for last-minute interventions and enables organisations to demonstrate clear and consistent control to regulators and stakeholders. This is particularly critical when working with technology which is not universally understood or trusted or with potentially sensitive information.

Over time, this approach also contributes to a more mature governance culture, one where compliance is understood as enabling effective decision-making, rather than constraining it.

Considerations

This approach requires consistency and discipline. When you’re using a tool to make your work more efficient, it can feel counterproductive to add in compliance processes. Some elements may feel repetitive, and there is a risk of over-standardisation if processes become too rigid. However, when applied proportionately, embedding compliance in this way can make governance more efficient, reduce potential regulatory exposure and support more reliable outcomes.

4 Apply practical ethical challenge

Ethical AI governance is fundamental. To ensure that it is effective, ethics must be actively discussed and modelled, rather than simply documented in policies. While formal frameworks provide a foundation, it is through everyday conversations, questioning, and reflection that ethical considerations are truly tested and applied.

For governance professionals, this involves creating and sustaining space for deliberate challenge, ensuring that ethical implications are considered in process design alongside operational and strategic factors, not after the fact. Where mistakes do occur, we must ensure that lessons are identified and remedied to prevent this becoming commonplace.

Embedding this in practice

In practice, ethical awareness often begins with asking the right questions at the right time. Simple prompts can be highly effective in surfacing issues that might otherwise go unexamined. For example:

+ who could be negatively impacted by this?" (eg student groups, applicants, or staff cohorts?)
+ would we be comfortable explaining this decision publicly?
+ could this disadvantage any groups in ways that conflict our widening participation or equality commitments?

Alongside this, governance professionals can support ethical practice by:

+ promoting transparency in your use of AI
+ considering ethical practice when designing processes and when selecting AI products
+ raising awareness of the potential conflict between increased adoption of AI tools and attainment of strategic environmental sustainability targets
+ as far as possible, ensuring that AI tools are fair, accountable and have mechanisms to reduce bias as far as possible
+ raising concerns early, even where information is incomplete, to ensure they are part of the conversation from the outset
+ maintaining professional scepticism about the information provided by an AI tool – be prepared to verify data before using it
+ encouraging open discussion in a way that feels constructive rather than escalatory
+ using real examples and case studies to ground discussions and make abstract considerations more tangible

+ raising awareness of ethical issues inherent in AI within your team, department and/or institution
+ encouraging the development of an AI Policy to guide practice in your institution
+ making sure there is human involvement in every process involving AI support.

These approaches help normalise ethical reflection as part of routine governance activity.

Risks to be aware of

Without active engagement, ethical considerations can be overlooked, particularly in time-pressured environments where the focus is on delivery. This can allow issues such as bias to go unchallenged and become embedded, increasing the risk of unintended consequences or reputational harm.

There is also a risk that, in the absence of open discussion and in an environment of widespread experimentation, concerns remain unspoken, limiting the effectiveness of governance oversight.

Opportunities this creates

Embedding ethical awareness within day-to-day practice supports a culture of responsible and informed decision-making. It strengthens both internal confidence and external trust, demonstrating that decisions are not only compliant, but considered and accountable.

Over time, this contributes to more resilient governance, where risks are identified earlier and decisions are better aligned with organisational values.

Considerations

Ethical questions are not always clear-cut when technology is still evolving, and discussion may introduce tension or differing perspectives as an institutional position is determined. Creating space for this kind of dialogue and enabling subjective views to be explored and balanced, requires confidence and facilitation, and it can take time to embed culturally.

However, by consistently encouraging critical thinking and thoughtful challenge, governance professionals can improve the quality of decision-making and reinforce the role of governance as a source of insight, not simply oversight.

5 Stakeholder engagement

We cannot deliver effective governance alone. Much of our role as governance professionals is to translate and integrate sound governance process into the realities of academic, professional service and student-facing practice. Where oversight relies solely on reports or outputs, there is a risk that important context is lost. By actively engaging with those using AI in operational settings, governance professionals can develop a more complete and grounded understanding of both risks and opportunities.

This principle is about maintaining visibility, not through additional layers of process, but through meaningful connection with people and practice.

Embedding this in practice

In practical terms, this involves creating opportunities for direct engagement beyond formal governance structures. This might include:

+ speaking directly with colleagues and operational teams who are using AI tools, to understand how they are applied in practice (eg academic staff, registry, IT, student service or data teams)
+ asking open questions such as: ‘How is AI affecting student experience or staff workload in practice?’ to surface issues that may not be formally reported
+ listening to developing views and be pragmatic in suggesting a proportionate governance response which can adapt to a digital world
+ sharing governance perspectives in clear, accessible language to support mutual understanding
+ building informal networks and relationships that complement formal reporting lines.

These approaches help bridge the gap between governance and delivery, ensuring that oversight is informed by lived experience.

Risks to be aware of

Where governance operates at a distance, there is a risk of becoming isolated from real-world issues. This can limit visibility of emerging risks, reduce the relevance of governance activity and create a perception that governance is disconnected from practice. Working directly with colleagues and stakeholders, co-creating guidance and operating with pragmaticism as well as robust process is vital.

A lack of engagement may also contribute to resistance from stakeholders, particularly where governance is seen as imposed rather than informed.

Opportunities this creates

By strengthening connections with operational teams, governance professionals can identify issues earlier and respond more effectively. This approach supports greater trust and cooperation, improving not only oversight but also the success of implementation.

It also reinforces the role of governance as informed, collaborative, and engaged rather than remote or purely procedural.

Considerations

Active engagement requires time and consistent effort, as well as strong interpersonal skills. Informal approaches, while valuable, may lack the structure of formal processes and need to be balanced appropriately.

However, by investing in these connections, governance professionals can significantly enhance transparency, deepen their understanding of practice and strengthen the overall effectiveness of governance.

6 Strong data governance

AI governance is only as strong as the data on which it is built and the prompts which we provide. While outputs may appear authoritative or precise, their reliability is fundamentally shaped by the quality, completeness and assumptions of the underlying data and any bias built into the algorithm.

For governance professionals, this means looking beyond outputs to consider what sits beneath them. We need to develop an informed curiosity about how data is sourced, interpreted, and applied in practice, then build measures to ensure that data is used and processed properly and with appropriate safeguards where necessary.

Embedding this in practice

In many cases, strengthening data awareness begins with asking simple but probing questions. For example:

+ where does this data come from?
+ how reliable is it? Can I verify it?
+ would I be happy if someone shared this information about me?

Alongside this, governance professionals can support better outcomes by:

+ challenging assumptions about data quality, rather than accepting it at face value
+ understanding that obligations relating to privacy, information security and professional integrity (eg UK GDPR, data protection legislation, and sector expectations around student and staff data use) continue to apply when working with an AI tool
+ raising concerns where data appears incomplete, inconsistent or potentially misleading, even where issues seem minor
+ seeking education on the potential risks involved in the AI products adopted
+ building relationships with data owners and specialists, such as with registry, planning teams or data protection officers, to improve understanding and visibility
+ anonymising confidential or sensitive material before inputting data into an open-source AI tool
+ adhering to institutional policies, where these are available
+ prioritising information security whenever working with AI models.

These actions help ensure that data is treated as an active area of governance consideration, rather than a fixed or unquestioned input.

Risks to be aware of

The risks involved in data governance when using AI are significant, and can incur real damage to institutional reputation, commercial interests and the privacy and dignity of our stakeholders. Where data is not subject to appropriate scrutiny or privacy, there is a risk that decisions are shaped by flawed or incomplete information. This may include hidden bias within datasets, gaps in coverage, or inaccuracies that go unnoticed.

Overconfidence in data, particularly when presented through AI-generated outputs, can further compound these risks, creating a false sense of certainty.

Information security is critical, and it is important to remember that AI tools are not extensions of our own minds. They are products, developed by third parties, who may have interests in our data that we are not aware of. Whilst a conversation with a large language model can be useful, there is a risk that the data we are providing could be used or sold on multiple times.

Opportunities this creates

By strengthening focus on data quality and assumptions, governance professionals can significantly improve the reliability and effectiveness of AI-supported decision-making. This enables more robust insights, enhances credibility and enables governance to engage more confidently with complex or data-driven proposals.

It also positions governance as a critical partner in ensuring that organisational data is not only used effectively, but understood appropriately.

Considerations

Data-related issues are not always easy to resolve and may reflect wider systemic challenges. Addressing them can require persistence, collaboration and, at times, difficult conversations about data quality or ownership.

However, by maintaining a consistent focus on what underpins AI outputs, governance professionals can reduce hidden risks, strengthen decision-making and support more effective and accountable use of AI.

7 Ongoing performance and evaluation

Governance effectiveness is not achieved through one-off review, but through sustained curiosity and consistent follow-through. In the context of AI, where systems evolve over time and conditions can change, initial assurance is only the starting point.

For governance professionals, this means maintaining an active interest in how AI systems, and our ways of working with them, continue to perform after they have been approved, ensuring that oversight remains relevant and responsive.

Embedding this in practice

In practice, continuous oversight is often built on simple, repeatable behaviours rather than complex structures. This might include:

+ regularly asking: ‘is this still working as expected?’
+ seeking clear, proportionate performance information, that can be meaningfully interpreted by governing bodies and committees, rather than relying on complex or overly technical reporting
+ following up on known issues to ensure they are addressed, rather than allowing them to drift over time
+ maintaining a visible record, aligned with existing risk registers or assurance frameworks if available, of identified risks and agreed actions to support accountability and continuity
+ being open to adapting and enhancing processes and approaches based on feedback, results and new information
+ ensuring AI-related risks and performance are appropriately reflected in committee reporting cycles.

These actions help ensure that governance remains active and engaged throughout the lifecycle of AI use.

Risks to be aware of

Without ongoing attention, there is a risk that AI performance degrades over time, particularly where models are sensitive to changing data or conditions. Issues that are identified early may be overlooked or deprioritised, while initial approvals can create a false sense of confidence that reduces scrutiny.

This can result in risks emerging gradually, without clear ownership or timely intervention.

Opportunities this creates

By embedding continuous oversight, governance professionals can identify issues earlier and support more effective, long-term outcomes. This approach strengthens accountability, reinforces the relevance of governance activity and ensures that decisions remain under appropriate review as circumstances evolve.

It also supports a more adaptive approach to governance, one that recognises that assurance is not static, but ongoing.

Considerations

Sustaining this level of attention requires consistency and persistence. Some activities may feel repetitive, and responsibility for follow-up may not always be clearly defined within existing structures.

However, by maintaining a focus on curiosity and follow-through, governance professionals can ensure that AI governance remains responsive, effective and aligned with the pace of change.

8 AI Group Case studies

Developing an AI governance framework in a university context

As artificial intelligence tools began to proliferate across the University - from ChatGPT to embedded functions in Microsoft 365 - it became clear that we needed a structured institutional response. The lack of clear ownership, coupled with growing concerns around risk, ethics, and practical use, highlighted the need for a Framework to guide staff, students and leadership.

The Framework was developed through our AI Working Group, which reports to the University Leadership Team (ULT) and brings together representatives from governance, digital services, academic practice, and professional services. With no single function owning AI governance, the process was collaborative and focused on building shared understanding.

To inform our approach, we ran an internal AI Usage Survey to better understand how staff were engaging with AI tools. The responses gave valuable insight into current practice, concerns, and areas of opportunity. These findings directly shaped the Framework and highlighted the differing needs of academic and professional services staff.

The result is a short, high-level Framework that sets out key principles to guide the University's approach. It is not a detailed policy or rulebook - instead, it provides a foundation for future decisions, clarity of direction and a common language for discussing AI use. It is expected to evolve as the technology changes.

ULT formally endorsed the Framework as a live document. The initial phase of implementation focuses on awareness-raising, aligning it with existing digital initiatives, and identifying governance gaps. The AI Working Group continues to meet, now turning its attention to how best to build down from the Framework - translating principles into more detailed guidance, tools, and support for staff.

A key lesson: in complex, fast-moving domains like AI, it is better to start small, work with what you have and iterate based on real-world needs and feedback.

From personal innovation to sector leadership in AI governance

Initial curiosity about how AI could support workload management led me, as a governance professional in further education, to explore generative AI tools to streamline policy reviews, enhance board communications and summarise key sector updates; followed then by introducing AI to streamline board meeting management. This hands-on approach offered efficiencies, but also raised important questions about ethics, accountability and transparency - particularly in how institutions govern AI use. One key lesson learned during this process was the importance of data protection, particularly regarding access controls in AI platforms. This led to the implementation of clearer usage protocols, data risk assessments and alignment with GDPR.

Recognising the strategic implications, a focus on AI governance at board level was introduced. To drive this forward, the college hosted a national AI in Governance Conference, bringing together sector leaders and board members from across the country. This platform initiated key conversations about board responsibility in overseeing AI risks, legal compliance, and ethical consideration. At institutional level, AI is now embedded into strategic reporting cycles, with AI oversight featuring in risk registers and board agendas. Draft AI governance principles are being developed to guide policy, procurement and decision-making, ensuring AI supports, not undermines, institutional integrity and public trust.

To maintain momentum, an FE AI Governance Working Group is also established, who collaborate on defining board-level responsibilities, assurance reporting structures, and ethical AI principles, aligning their work with national guidance from JISC and leading FE lawyers.

This journey from personal innovation to sector-wide leadership demonstrates how governance professionals can play a pivotal role in shaping ethical, accountable AI adoption across education and something I am grateful to champion.

Strengthening AI governance following an unintended AI chatbot disclosure

During a routine farewell gathering for a departing staff member, held via Microsoft Teams, a lesson emerged that underlined the critical importance of robust AI protocols. Unbeknown to organisers, an AI chatbot linked to an attendee's personal email account had joined the meeting. The bot was designed to automatically capture, summarise and email key points from meetings. Without human review or consent, it recorded the entire event, including an informal speech by the line manager in which a joke, intended as light-hearted, could be misconstrued as offensive.

The AI then automatically generated a transcript and summary and emailed it to the full meeting invite list, to over 100 individuals, amplifying the risk of reputational damage, personal distress and a possible breach of data.

Immediately, a review of meeting protocols and institutional AI governance was initiated. AI chatbots are now explicitly covered under the AI policy, with clear restrictions on their use during meetings unless pre-approved. Staff are briefed on the presence and risks of third-party AI tools and protocols have been introduced for vetting all digital attendees, including CAPTCHA-style identity verification for external participants. Meeting organisers must now confirm whether AI transcription tools are active and ensure appropriate content management and consent procedures are in place.

This incident highlighted the need for AI oversight beyond technical systems and highlighted our professional role in setting policies, raising awareness and embedding ethical safeguards. By applying these lessons, the institution reinforced its commitment to responsible AI use, balancing innovation with trust, transparency and protection of its people.

Using AI in meetings

During discussion with colleagues, as AI started to become more widely available, it became clear that there were some behaviours taking place which raised concerns. Probably the biggest potential impact was people sending AI Assistants to meetings. While this was not for the 'top tier' committees it was concerning as they were being sent in place of the person attending the meeting. This did not support either good decision making or good governance.

It also quickly became clear that these assistants represented an information security risk with the potential for anything shared with them being made widely available on the internet. Concerns were also raised that the use of such applications increased potential for a cyber security breach. Given the risks and the unethical use of AI Assistants, two actions were taken. Firstly, it was confirmed that all such assistants could be removed from the meeting should they appear (basically giving secretaries permission to do something they were already doing) and IT blocks were put in place to prevent their use. This does seem to have worked.

The discussions held around all of this also allowed some other concerns to be raised about the use of AI and its wider impact including amount of storage needed and the environmental impact. The IT storage issues are being supported by a review of document retention and disposal policy, the environmental impact is very much up for wider debate as measurements do vary.

It has been a search for good practice that has ultimately led the desire to support the production of these principles for the use of AI in HE Governance. It may not solve all the dilemmas but should provide a solid foundation from which we can build.

Normalising the use of AI in governance

Introducing AI into governance practice initially proved challenging. When I first tried to encourage exploration within my institution, there was clear hesitation and I was told that pursuing this technology could risk my colleagues' jobs. Many colleagues were unsure whether they were 'allowed' to experiment, and even starting a conversation about potential applications for staff felt like pushing against a closed door. For our institution, AI conversations were limited to academic malpractice. Individual curiosity amongst staff was growing, but confidence and permission were not.

Everything changed almost overnight once tools like ChatGPT became mainstream. Suddenly, senior colleagues were experimenting enthusiastically, often without much technical understanding and with a degree of naivety about risk. I saw practices that made me wince, full CVs, confidential business cases and sensitive internal documents pasted straight into public AI tools. The link between third-party large language models, data protection responsibilities and GDPR simply wasn't front of mind, and the gap between excitement and responsible use became increasingly obvious.

Trying to shift the conversation from the bottom up was difficult, especially when risky behaviours were being modelled from the top. I wanted to encourage experimentation but also needed people to understand the realities of the risks inherent in the tools: hallucinations, data leakage and over-reliance. Within governance, I was also keen to move beyond basic note-taking and explore more strategic, value-adding uses.

On a personal level, AI has become a genuinely helpful part of my working routine. I use it to plan and prioritise my day, run quality checks on reports, and pull out actions and decisions from meeting notes. To avoid losing core skills, I focus on prompts that strengthen my thinking rather than letting the tool do the work for me.

As the technology evolves, I'm committed to using it ethically and thoughtfully. Normalising AI in governance, for me, is about balancing innovation with accountability, embracing the benefits while staying clear-eyed about the risks.

9 Conclusion: from everyday practice to personal commitment

Across UK FE and HE, governance is already shaped by complexity, balancing student outcomes, academic integrity, financial sustainability and regulatory accountability. The emergence of AI does not change these core responsibilities, but it does change how they are exercised. It challenges governance professionals to look more closely, ask more probing questions and engage more actively with how decisions are informed, shaped and delivered in practice.

The principles set out here are not intended as a new layer of governance, nor as a framework to sit alongside existing requirements. Instead, they reflect a shift in emphasis, from formal structures to the everyday behaviours that give governance its meaning and impact. Effective AI governance is not created by policy alone; it is built through the consistent application of professional judgement, curiosity and constructive challenge in the moments that matter.

This shift can be understood in simple terms. It is a movement:

+ from policy to practice
+ from oversight to active challenge
+ from frameworks to real decisions and outcomes
+ from top-down control to bottom-up effectiveness.

In this context, the risks we face are not entirely new. However, they may become more pronounced. Governance can become passive, with outputs accepted at face value rather than interrogated. Processes can be followed without reflection, giving reassurance without necessarily providing insight. And governance itself can drift away from operational reality, limiting its ability to identify emerging issues or add meaningful value.

Yet the opportunity is equally clear. Where governance professionals engage actively by connecting with practice, questioning thoughtfully, and maintaining a clear line of sight to real-world impact they can strengthen decision-making across their institution. Governance becomes not simply a mechanism for oversight, but a source of insight, assurance and improvement. It builds confidence internally and trust externally, demonstrating that emerging technologies are being used responsibly, transparently and in line with institutional values.

At its heart, this is about how we show up in our roles. The most effective governance professionals are not defined solely by the frameworks they apply, but by the habits they build and the behaviours they model. They ask better questions. They remain curious, even when answers seem clear. They engage widely, listen carefully, and challenge constructively. And they maintain a consistent focus on the real-world impact of decisions particularly for students, staff, and the wider communities our institutions serve.

This is why improving AI governance does not start with new structures. It starts with individuals.

The next step, therefore, is to move from shared principles to personal commitment. A code of practice provides a practical way to do this setting out clear, accessible expectations that governance professionals can adopt in their own work. Not as an abstract standard, but as a deliberate choice about how they will act, question, and contribute in an AI-enabled environment.

What follows is intended as that commitment: a code grounded in daily practice, designed to support professional judgement, and shaped by the realities of governance within UK FE and HE.

Appendix 1: AI Code of Practice for governance professionals

The principles set out in this guide are intended to support governance professionals in navigating the opportunities and challenges presented by artificial intelligence in a practical and informed way. However, as with all areas of effective governance, impact is not determined by principles alone, but by how they are applied in day-to-day practice.

The Code of Practice that follows is designed to translate these principles into a clear and personal commitment. It recognises that governance effectiveness is shaped through individual behaviours, including the questions we ask, the judgements we make and the standards we choose to uphold through our work. In this sense, this Code is not an additional requirement, but a practical tool to support consistency, confidence and professional integrity in an AI-enabled environment.

AI introduces new considerations, but it does not alter the fundamental expectations placed on governance professionals. Accountability, transparency, ethical awareness and sound judgement remain central. What changes is the context in which these are exercised, often at greater pace, with less visibility and with unpredictable forms of risk and complexity. The Code therefore provides a structured way to remain anchored to core expectations, while adapting practice to reflect emerging technologies.

Importantly, this Code is intended to be proportionate and usable. It is not a prescriptive or exhaustive framework, but a guide that can be applied flexibly across different institutional contexts, levels of experience and types of AI use. It is equally relevant whether engaging in early experimentation or supporting more established applications. Its value lies in supporting reflective practice, encouraging consistency and reinforcing the role of governance professionals as active contributors to responsible AI use.

By adopting this Code, governance professionals are making an explicit commitment: to engage thoughtfully with AI, to remain professionally curious, to apply appropriate challenge and to act with integrity in all circumstances. In doing so, they help ensure that AI is used not only effectively, but responsibly, strengthening both governance practice and institutional trust.

Purpose

I will use AI to support efficiency, insight, and quality in governance work while upholding the principles of accountability, integrity and transparency.

Professional judgement

+ I remain accountable for all outputs I use or share.
+ I will not rely on AI without applying critical thinking and professional expertise.
+ I will challenge, sense-check, and validate AI outputs before use.
+ I will ensure decisions are made by people, not tools.

Data responsibility

+ I will not input confidential, personal or sensitive information into AI tools unless operating within institutional policies.
+ I will follow all organisational data protection and information governance requirements.
+ I will minimise data exposure wherever possible.

Accuracy & quality

+ I will verify outputs against trusted sources, policies and known facts.
+ I will recognise the limitations of AI (including errors and omissions).
+ I will ensure outputs are appropriate for their intended purpose.

Ethics & integrity

+ I will use AI in a way that aligns with the values of good governance: openness, fairness and accountability.
+ I will be transparent about AI use.
+ I will take responsibility for content where I have used AI. I will not present AI-generated work as wholly my own.
+ I will not present AI-generated work as wholly my own without review and ownership.

Ethical judgement

+ I will consider whether AI use could mislead, bias or reduce trust.
+ I will use AI to support, but not replace, scrutiny or professional challenge.
+ I will reflect on whether I could justify both the outcome *and* the method used.

Proportional use

+ I will use AI appropriately based on the level of risk.
+ I will avoid using AI in high-risk or sensitive decisions without proper oversight and/or expert advice.
+ I will seek support to ensure that benefits outweigh the risks where there is uncertainty.

Transparency & audit

+ I will maintain clarity on how outputs were generated where needed.
+ I will ensure I could explain and defend my use of AI if challenged.
+ I will keep appropriate records where AI use influences governance decision making.

Continuous improvement

+ I will build my understanding of AI capabilities and limitations.
+ I will learn from experience and share good practice with colleagues.
+ I will stay aware of emerging risks, guidance and standards.

Personal commitment

+ I will support, model and advocate for professional and ethical use of AI within my institution.
+ I will share my AI literacy and experience to support the identification of best practice.
+ Before using AI outputs, I will ask:

– does this make sense?
– is this safe and appropriate?
– would I be comfortable defending this?

If not, I will stop and reassess.

Contact us

All enquiries

Email: enquiries@advance-he.ac.uk

Advance HE helps HE institutions be the best they can be, by unlocking the potential of their people.

We are a member-led, sector-owned charity that works with institutions and higher education across the world to improve higher education for staff, students and society. We are experts in higher education with a particular focus on enhancing teaching and learning, effective governance, leadership development and tackling inequalities through our equality, diversity and inclusion (EDI) work.

Our strategic goals to enhance confidence and trust in HE, address inequalities, promote inclusion and advance education to meet the evolving needs of students and society, support the work of our members and the HE sector.

We deliver our support through professional development programmes and events, Fellowships, awards, student surveys and research, providing strategic change and consultancy services and through membership (including accreditation of teaching and learning, equality charters, knowledge and resources).

Advance HE is a company limited by guarantee registered in England and Wales no. 04931031. Company limited by guarantee registered in Ireland no. 703150. Registered as a charity in England and Wales no. 1101607. Registered as a charity in Scotland no. SC043946. Registered Office: Advance HE, Innovation Way, York Science Park, Heslington, York, YO10 5BR, United Kingdom.

To request copies of this report in large print or in a different format, please contact Advance HE: enquiries@advance-he.ac.uk