


BEYOND TECHNICAL LITERACY

AI Governance Without AI Understanding — Building the Discipline

A Practitioner's Methodology for Equipping Boards and Executive Teams to
Govern What They Cannot Yet Comprehend

Giovanni Leonardi

September 2023



Contents

1. Executive Summary

2. The Governance Gap

3. Who This Methodology Is For

4. Stage 1: Governance Architecture

4.1 The Governance Design Principles

4.2 Structural Components

4.3 The Independent Technical Advisor Role

4.4 Outputs

5. Stage 2: Decision Framework

5.1 The Five Questions Framework

5.2 AI Risk Appetite Statement

5.3 Board Information Pack Standards

5.4 Outputs

6. Stage 3: Assurance Structure

6.1 Three Lines of Assurance

6.2 Assurance Activities

6.3 Outputs

7. Stage 4: Maturity Progression

7.1 The Maturity Model

7.2 Progression Activities

7.3 Outputs

8. Practical Considerations

8.1 The Cost of Governance

8.2 Starting Where You Are

8.3 The Regulatory Horizon

9. Conclusion

3

3

4

4

4

5

5

5

6

6

6

7

7

7

8

8

8

8

9

9

9

10

10

10

10

11

AI Governance Without AI Understanding — Building the Discipline

Leadership · No. 105

Originally written September 2023 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The governance gap is not that boards lack technical knowledge — it is that they lack a structured method for exercising judgement over systems whose behaviour they cannot directly inspect.”

1. Executive Summary

Boards are being asked to govern AI programmes they fundamentally do not understand. This is not a criticism — it is a structural reality. The technical complexity of machine learning systems, the opacity of large language models, and the speed at which the field is evolving have created a governance gap that no amount of executive AI training will close in time. The question is not how to make boards technically literate. It is how to build a governance discipline that enables sound oversight despite incomplete technical comprehension.

This methodology provides a complete, operational approach to that problem. It is designed for chairpersons, non-executive directors, governance professionals, and the executive teams who support them. It sets out four stages — Governance Architecture, Decision Framework, Assurance Structure, and Maturity Progression — each with defined roles, artefacts, and practical mechanisms. The goal is not AI expertise at board level. It is AI governance competence: the ability to ask the right questions, interpret the answers, set appropriate boundaries, and hold the executive to account.

2. The Governance Gap

The pattern is now well established. A board receives a presentation on AI strategy. The slides are enthusiastic, the use cases compelling, the competitive narrative urgent. The board approves the investment. Six months later, the same board is asked to approve a deployment that raises questions about data privacy, algorithmic bias, or regulatory compliance — questions the board does not have the vocabulary to interrogate, let alone the frameworks to resolve.

The governance gap is not that boards lack technical knowledge — it is that they lack a structured method for exercising judgement over systems whose behaviour they cannot directly inspect. This is not unprecedented. Boards have governed complex domains

before — nuclear safety, pharmaceutical development, financial derivatives — without requiring every director to be a physicist, chemist, or quant. But in each of those domains, governance matured over decades, supported by regulation, professional standards, and painful lessons. AI governance does not have decades. Organisations are deploying systems now, and boards must govern them now.

What makes AI governance distinctively difficult is the combination of three characteristics:

- **Opacity:** Many AI systems — particularly those built on deep learning architectures — do not produce decision logic that can be read, audited, or explained in the way a traditional business rule can. The board cannot ask to see the reasoning behind a specific output in the same way it might ask to see the calculation behind a financial provision.
- **Emergent behaviour:** AI systems can behave in ways not anticipated by their designers, particularly when deployed on data that differs from their training set. Governance must account for what the system might do, not only what it was designed to do.
- **Velocity:** The technology is evolving faster than governance frameworks, regulatory guidance, or organisational learning can keep pace. A governance approach that is adequate today may be inadequate within months.

3. Who This Methodology Is For

This method addresses three audiences:

- Board members and non-executive directors who need to govern AI effectively without becoming technical experts — and who need a structured approach rather than ad hoc briefings.
- Company secretaries and governance professionals who must design the governance architecture and ensure it functions in practice.
- Executive sponsors and CIOs who need to present AI programmes to their boards in ways that enable genuine governance rather than rubber-stamping.

4. Stage 1: Governance Architecture

Purpose: Design the board-level governance structure for AI oversight, including committee responsibilities, information flows, and the boundary between governance and management.

Duration: 4–8 weeks for initial design; ongoing refinement.

4.1 The Governance Design Principles

Before establishing structures, five principles must be agreed at board level:

1. Governance is not management. The board's role is to set boundaries, approve risk appetite, and hold the executive to account. It is not to approve individual models, review training data, or make technical design decisions. The most common governance failure is boards that oscillate between disengaged approval and anxious micro-management.
2. Comprehension is delegable; accountability is not. The board may — and should — rely on technical advisors, assurance functions, and executive expertise to understand AI systems. But the decision about whether those systems are acceptable risks remains the board's.
3. AI governance is risk governance. It belongs within the existing risk governance framework, not in a separate innovation committee. AI risk is operational risk, reputational risk, regulatory risk, and conduct risk. Treating it as a special category isolates it from the disciplines that would otherwise apply.
4. Proportionality applies. Not every AI application requires board attention. The governance architecture must include a triage mechanism that escalates appropriately — Tier 1 systems to the board, Tier 2 to a delegated committee, Tier 3 to management processes.
5. Governance must evolve. The architecture designed today will need to change as the organisation's AI maturity, the regulatory environment, and the technology itself develop. Build review and adaptation into the structure from the start.

4.2 Structural Components

Component	Purpose	Composition
Board Risk Committee (or equivalent)	Primary board-level oversight of AI risk; receives AI risk reports, approves AI risk appetite, escalation point for Tier 1 decisions	Existing committee membership; no AI expertise required at this level
AI Governance Sub-Committee	Delegated authority for Tier 2 oversight, detailed review of AI risk reports, challenge of executive AI governance	2–3 NEDs, CIO or CTO, Chief Risk Officer, independent technical advisor
Independent Technical Advisor	Provides the board with an independent technical perspective; translates between technical reality and governance language	External appointment: senior technologist with AI/ML experience and governance fluency
Executive AI Governance Forum	Management-level governance: approves Tier 2 and Tier 3 deployments, manages the AI risk register, reports to the sub-committee	CIO/CTO (chair), data and analytics leads, risk, legal, and business unit representatives

4.3 The Independent Technical Advisor Role

This role deserves specific attention because it is the single most important structural innovation in the methodology. Boards governing complex technical domains need a trusted, independent voice that is not the executive asking for investment approval. The Independent Technical Advisor (ITA) serves the board, not management. They attend

sub-committee meetings, review AI risk reports before they reach the board, and provide an independent assessment of whether the information the board is receiving is complete, accurate, and appropriately framed.

The ITA is not a consultant. They are a governance appointment, analogous to the external auditor's relationship with the audit committee. They must have sufficient technical depth to challenge the executive's technical claims and sufficient governance experience to frame their advice in terms the board can act on.

4.4 Outputs

- Board AI governance framework document (structures, authorities, escalation paths)
- Updated terms of reference for the Board Risk Committee and any delegated sub-committee
- Independent Technical Advisor appointment and terms of engagement
- AI risk appetite statement (initial version, to be refined in Stage 2)
- Triage criteria for Tier 1/2/3 classification of AI applications

5. Stage 2: Decision Framework

Purpose: Equip the board and its committees with structured decision-making tools that enable sound AI governance judgements without requiring technical expertise.

Duration: 6–10 weeks, running in parallel with early AI programme activity.

5.1 The Five Questions Framework

At the heart of this stage is a structured questioning framework that board members can apply to any AI proposal, deployment decision, or risk report. The framework is designed to surface the information that matters for governance without requiring the questioner to understand the underlying technology.

1. What problem does this solve, and for whom? This question tests whether the AI application addresses a genuine organisational need or is a technology-push initiative. It also identifies who is affected by the system's outputs — the starting point for any fairness or impact assessment.
2. What happens when it is wrong? Every AI system produces errors. This question forces the presentation of the error profile: how often the system is wrong, what the consequences of errors are, and what mechanisms exist to detect and correct them. It is the single most revealing question a board member can ask.
3. Can we explain it? Not how does it work technically but can we explain to a customer, a regulator, or a journalist why this system made a specific decision? If the answer is no, the governance implications are severe.
4. What could change that would make it unsafe? This question addresses the emergent behaviour risk: shifts in input data, changes in the operating environment, drift in model performance over time, or regulatory changes that alter the compliance landscape.

5. How will we know if something goes wrong after deployment? This tests the monitoring, alerting, and incident response infrastructure. A system that cannot be effectively monitored in production cannot be effectively governed.

These five questions are not a checklist to be completed mechanically. They are a governance discipline — a structured way for non-technical directors to interrogate AI proposals with the same rigour they would apply to a major capital investment or a significant regulatory change.

5.2 AI Risk Appetite Statement

The board must articulate its risk appetite for AI in terms that connect to existing risk categories. The AI risk appetite statement should address:

- Which categories of decision the organisation is willing to automate, and which must retain human oversight regardless of technical capability
- The acceptable error profile for AI-assisted decisions in each business domain
- The organisation's position on transparency — what level of explainability is required for different stakeholder audiences
- The circumstances under which an AI system must be withdrawn from production

5.3 Board Information Pack Standards

A critical output of this stage is defining what information the board receives about AI systems and in what format. The methodology specifies a standard AI Governance Information Pack that accompanies every Tier 1 board submission:

Section	Content	Purpose
Business Case Summary	Problem statement, expected outcomes, affected populations	Grounds the discussion in business value and impact
Risk Profile	Error rates, bias assessment results, data quality summary	Enables risk-based judgement
Explainability Assessment	Can the system's decisions be explained to each relevant audience?	Tests transparency readiness
Monitoring Plan	What is monitored, how often, what triggers escalation	Tests operational governance readiness
ITA Assessment	Independent Technical Advisor's opinion on the above	Provides independent assurance

5.4 Outputs

- Five Questions Framework (documented, with worked examples)
- AI risk appetite statement
- Board AI Governance Information Pack template
- Decision authority matrix (which decisions sit at board, sub-committee, and management level)

6. Stage 3: Assurance Structure

Purpose: Establish the mechanisms that give the board confidence that AI governance is functioning as designed — that the information it receives is accurate, that risk classifications are appropriate, and that deployed systems continue to meet governance requirements.

Duration: 8–12 weeks for initial build; continuous operation thereafter.

6.1 Three Lines of Assurance

The methodology applies the established three-lines model to AI governance:

1. First line — Management controls: The delivery and operations teams are responsible for building and operating AI systems within the governance framework. They produce the monitoring data, conduct the bias testing, and maintain the documentation. First-line assurance means the people doing the work are following the rules.
2. Second line — Risk and compliance oversight: The risk function (supported by the RAI Lead, where the role exists) provides independent challenge to first-line activities. They review risk classifications, validate bias testing methodologies, and assess whether monitoring arrangements are adequate. Second-line assurance means someone independent of delivery is checking the work.
3. Third line — Internal audit: Internal audit provides periodic, independent assurance that the entire governance framework is functioning as designed. They audit not just individual AI systems but the governance process itself: are risk classifications being applied consistently? Is the board receiving accurate information? Are escalation mechanisms working?

The Independent Technical Advisor operates alongside the second and third lines, providing the technical depth that neither the risk function nor internal audit typically possesses.

6.2 Assurance Activities

- Quarterly AI Risk Report: A structured report from the Executive AI Governance Forum to the Board Risk Committee, covering: new AI deployments approved, monitoring alerts and incidents, bias test results and trends, regulatory developments, and the RAI Lead's assessment of overall governance health.
- Annual AI Governance Audit: Internal audit conducts a full review of the AI governance framework against the board's approved design. Findings are reported to the Audit Committee with management actions tracked to closure.
- ITA Annual Assessment: The Independent Technical Advisor provides an annual written assessment to the board sub-committee on the quality of AI governance across the organisation, including specific recommendations for improvement.

6.3 Outputs

- Three-lines assurance model documented and roles assigned
- Quarterly AI risk report template and reporting schedule

- AI governance audit programme (included in internal audit's annual plan)
- ITA annual assessment terms of reference

7. Stage 4: Maturity Progression

Purpose: Build the organisation's AI governance capability over time, moving from reactive oversight of individual systems to proactive, strategic governance of an AI portfolio.

Duration: Ongoing; designed to span 18–36 months from initial implementation to target maturity.

7.1 The Maturity Model

Level	Characterisation	Governance Capability
Level 1 – Reactive	AI deployed without structured governance; board awareness minimal; decisions made ad hoc	Governance structures exist on paper but are not consistently applied
Level 2 – Structured	Governance architecture in place; decision framework operational; assurance beginning	Board receives regular reports; risk appetite articulated; ITA appointed
Level 3 – Integrated	AI governance embedded within enterprise risk management; board capable of effective challenge; assurance mature	Governance operates as business-as-usual; board asks the right questions without prompting; lessons from incidents feed back into governance design
Level 4 – Strategic	Board governs AI as a portfolio; governance informs strategy, not just risk management; organisation contributes to industry governance standards	Board uses AI governance insights to shape strategic direction; governance is a competitive advantage, not a compliance burden

7.2 Progression Activities

1. **Annual Maturity Assessment:** The governance professional (company secretary or equivalent) conducts an annual assessment against the maturity model, identifying the current level and the specific actions needed to progress to the next.
2. **Board Development:** Structured learning for board members — not technical training, but governance-focused development. Case studies of AI governance failures and successes, facilitated discussions with the ITA, and exposure to regulatory thinking. Two sessions per year, integrated into the board calendar.
3. **Peer Review:** Where possible, engage with peer organisations or industry bodies to benchmark AI governance maturity. The emerging work from regulators on AI governance expectations (the FCA's interest in AI explainability, the PRA's focus on model risk management for ML systems) provides a useful external reference point.

4. Governance Framework Review: Every 18 months, conduct a formal review of the entire governance architecture against current best practice, regulatory expectations, and the organisation's own experience. Update structures, terms of reference, and the decision framework as needed.

7.3 Outputs

- Annual maturity assessment report
- Board AI governance development programme
- Governance framework review reports (18-monthly cycle)
- Updated governance architecture (as the organisation matures)

The measure of AI governance maturity is not whether the board can explain how a neural network works. It is whether the board can explain why it approved a specific AI deployment, what risks it accepted, and how it will know if those risks materialise.

8. Practical Considerations

8.1 The Cost of Governance

This methodology requires investment: the ITA appointment, governance professional time, board development, and assurance activities all carry cost. In my experience, the total is modest relative to the AI programme investment it governs — typically between 2 and 5 per cent of the AI programme budget. The cost of governance failure — a biased system deployed without oversight, a regulatory breach discovered by a journalist rather than an auditor, a board unable to explain its own approval decisions — is orders of magnitude greater.

8.2 Starting Where You Are

Most organisations implementing this methodology will begin at Level 1 maturity. This is normal. The methodology is designed to be implemented incrementally: start with the Governance Architecture (Stage 1), add the Decision Framework (Stage 2), and build assurance and maturity progression as the organisation's AI portfolio grows. The critical first step is appointing the Independent Technical Advisor and establishing the Five Questions Framework — together, these two elements provide immediate governance capability even before the full architecture is in place.

8.3 The Regulatory Horizon

As of September 2023, the regulatory environment for AI governance is rapidly crystallising. The EU AI Act, while still in legislative process, will impose specific governance obligations on high-risk AI systems. The NIST AI Risk Management Framework provides a voluntary but increasingly influential reference. Sector regulators are issuing guidance with growing specificity. Boards that build governance capability now will be in a strong position when regulatory requirements harden. Those that wait

will find themselves building under pressure, which invariably produces weaker governance.

9. Conclusion

AI governance without AI understanding is not a contradiction — it is the reality most boards face, and it is a solvable problem. The solution is not to wait until boards become technically literate. It is to build a governance discipline — structures, decision frameworks, assurance mechanisms, and a path to maturity — that enables boards to govern effectively with the comprehension they have, supplemented by the independent technical advice they need. This methodology provides that discipline. It does not require boards to understand AI. It requires them to govern it.