Best Practices for AI Governance Structures:
Executive Oversight and Internal Review

Responsible Artificial Intelligence Institute

Best Practices for AI Governance Structures: Executive Oversight and Internal Review

Introduction

Due to the rapid pace of technological and policy change, companies are prioritizing establishing or refining AI governance. The often opaque (black-box) nature of AI systems, their potential to amplify harms, and their ability to scale uniformly create a clear need to connect policy and governance with tooling, training, procurement, and compliance. To drive the change, leaders are adopting a mix of existing and new processes and structures. The composition of this approach depends on an organization's industries, existing organizational structures, needs, culture, and risk appetite.

The substance of AI governance approaches is more important than whether they are incorporated into existing or new processes and structures
Two major elements of Responsible AI governance are:

- **Executive Oversight by a High-Level Board/Steering Committee:** Provide direction, mandates, and resourcing to responsible AI efforts in a timely manner
- **Internal Review by an Operational Committee:** Conduct arm's-length internal reviews of AI systems at various stages, including feasibility and resource allocation, technical scope, and responsible AI considerations

Executive Oversight by a High-Level Board/Steering Committee

Composition: Leaders ensure that the Central AI Board/Steering Committee includes C-suite executives, senior legal and compliance officers, and technology leaders. Some leaders also incorporate or frequently consult external advisors or experts to offer independent insights on complex AI issues.

Scope and Responsibilities: The Central AI Board/Steering Committee's responsibilities generally include making informed decisions about AI strategy, priorities, and resourcing. The Steering committee receives reports from the Operational Committee(s) and resolves escalated high-priority issues. It evaluates AI's impact on the organization, innovation, customer experience, supply chain management, risk reduction, efficiency improvements, and unexplored strategic opportunities. In addition to strategic planning, this body leads efforts to comply with AI-relevant laws and approves related policies and controls. Leaders are also increasingly recognizing the importance of designating specific individuals to lead these Central AI Boards/ Steering Committees. This role can take the shape of a Chief AI & Data Officer, RAI Officer, or AI Champion, depending on the size, structure, and needs of the organization.

Internal Review by an Operational Committee

Composition: Leaders ensure that Internal Review Committees (IRCs) are cross-functional (involving multiple departments), with representation from business lines, technology, and compliance. These committees are tasked with practical implementation and coordination of all responsible AI efforts and workstreams, described in the Appendix.

Scope and Activities: The primary functions of an IRC are the development and implementation of AI policies, governance, tools, and training and the ongoing Internal Review (IR) of AI projects and project-types. IR is a flexible, robust approach to AI governance. It can be leveraged in almost any context and is technology and methodology-agnostic. IR broadly involves convening an interdisciplinary committee of experts who can evaluate a project - an algorithm, a robot, a novel use case of existing tech - for technical, responsible, and resource considerations. IR has a long track record of use in various industries and contexts, including in the biomedical sciences, the pharmaceutical industry and academia. Over the last decade, Big Tech companies have started adopting IR in their governance of the technology lifecycle, demonstrating IR's adaptability and scalability.

Lessons from Leaders on Internal Review

Research suggests that successful IRCs need three key characteristics: impartiality, accountability, and effective consultation.

- **Impartiality:** Organizations must create a strong separation between IRCs and the rest of the organization to mitigate bias and ensure a reliable assessment of risks, impacts, and feasibility.
- **Accountability:** IRCs should be empowered to meaningfully recommend adjustments to projects in balance with strategic business priorities. Embedding IRCs in the technology lifecycle early ensures consultation can be leveraged early to prevent more costly and disruptive adjustments once the build has started.
- **Effective Consultation:** IRCs should have the right mix of expertise to be able to provide a true 360-degree review of a project. Experts should be gathered based on the type of project and underlying technology being discussed and should include a mix of technical, legal, or privacy, business experts. Users or members of the public should be considered where appropriate to provide a non-expert lens on high-impact technology intended for broad use.

An operational IRC's task will be to evaluate a proposed technology project or product along three key review areas: technical, responsible, and resource.

Internal Review Committee
OVERVIEW

How to structure your IR Committee (IRC)

An effective and well embedded IRC should have three **key characteristics:**

Impartiality
Minimize conflicts of interest to ensure evaluation reliability

Accountability
Empower IRC to provide guiding feedback and embed IRC in technology lifecycle

Effective Consultation
Gather the right mix of experience and expertise in the room to provide effective 360°

How to evaluate your project with an IRC

Internal review should evaluate a project based on three **key objectives:**

Technical
Assess the project's technical scope and feasibility

Responsible
Determine the project's ethical impact and alignment to organizational values

Resource
Evaluate the project's financial and staffing needs against current team

- **Technical**: Technical review should assess the technology selection against the stated usage or business goals. Use case scoping, model selection, infrastructure capacity, technical debt, integration needs, and other technical considerations should be considered
- **Responsible**: The project should be evaluated for technical, regulatory, and business risks. The alignment of the project to the organization's stated values and policies should be determined. Downstream ethical impacts from broad-use technology should be anticipated
- **Resource**: The financial and staffing needs of a project should be considered before greenlighting. Do you have all the required expertise in-house or are external resources needed? Does the budget realistically reflect the scope? Change management for business process modernization should be anticipated.

The Internal Review conducted by the IRC should result in a set of clear, actionable recommendations to align a reviewed project with its stated goals, the organization's stated mission, and the broader regulatory environment.

Since AI and technology projects are built iteratively, the Internal Review process should be iterative as well - once a project is reviewed, it should be reassessed if there are major changes. The aim is to optimize the review timeline to minimize the impact on project timelines. This may mean having an IRC meet monthly, every 6 months, or annually - depending on the organization's size and project needs.

Examples of Internal Review in Different Contexts

Big Tech

- Big Tech and many large companies already employ IR or governance structures that resemble IRCs. Well-known IRC committees at Big Tech companies are responsible for creating internal policies around responsible technology delivery, publicizing governance and safety efforts, and ensuring compliance cascades through the organization.
- If yours is a large company, you may have internal committees that can be adapted to perform Internal Reviews on key technology projects. Adapting existing structures reduces change management friction and optimizes embedding in the technology lifecycle.
- Legal and Compliance teams can be augmented or adapted to ensure the necessary expertise for Internal Review is available, similar to their adaptation to privacy practices globalized by the EU General Data Protection Regulation (GDPR).
- Some large organizations establish Chief Ethics Officers to strategically emphasize their efforts to deploy technology responsibly, especially if they operate in highly regulated industries or have high reputational risks.

Healthcare

- Internal Review is a well-established practice in the healthcare space. Academic institutions and pharmaceutical industries are required by law or policy based on their jurisdiction to perform internal reviews of research projects and clinical trials involving human participants.
- The challenge for healthcare organizations will be adapting their current IR practices to be able to evaluate healthtech projects. This will require committee education and training to ensure a baseline understanding of proposed technical projects. It will also require an expanded network of outside technical experts who can be called upon to provide relevant subject matter expertise on proposed projects.

Startups and Small Organizations

- Startups and small organizations will face financial and resource challenges to establish robust Internal Review. If your organization lacks the right mix of expertise for IR, you can consider external consultants for specific project review. If hiring externals is out of reach, there may be responsible tech think tanks or academic collaboration clusters that provide free tools, forums, and events that can be leveraged to start the IR process.
- It is important to note that an organization's small size or scope of project does not necessarily shield it from regulatory obligations - or the financial and legal risks of building a tool with serious unintended consequences. Demonstrating your commitment to safety and responsible deployment early and with the tools you have available is a key step in AI and tech governance.

Supporting You on Your RAI Journey

The Responsible AI Institute is here to support organizations on their AI journeys. Becoming a member enables essential support and direction for making significant progress toward future-proofing RAI governance and implementation.

Click below to learn more about how we help our members achieve their Responsible AI goals:

About Responsible AI Institute (RAI Institute)

Founded in 2016, the Responsible AI Institute (RAI Institute) is a global and member-driven non-profit dedicated to enabling successful responsible AI efforts in organizations. We accelerate and simplify responsible AI adoption by providing our members with AI assessments, benchmarks and certifications that are closely aligned with global standards and emerging regulations.

Where to connect with us:

Appendix: Responsible AI Activities

- **AI Governance Lifecycle, Tools, and Reporting:** Maintain, and promote awareness of, policies, processes, governance gates, and tools for AI governance - such as Responsible AI Policies and AI Impact Assessments.
- **Board Oversight:** Oversee the company's AI strategy in an engaged and informed way.
- **Data Posture:** Maintain relevant and high-quality data, while conforming to best practices in data governance and data management.
- **Risk Management Strategy and Execution:** Consider and action AI's risk, privacy, cybersecurity, compliance, and IP considerations.
- **Resource Planning:** Allocate human and technological resources effectively, including necessary tools for AI and RAI initiatives.
- **Third-Party and Partner Risk Management:** Actively monitor and mitigate risks associated with third parties and partners.
- **Stakeholder Engagement and Diversity, Equity, and Inclusion**: Ensure that AI developer teams and reviewers represent varied backgrounds and incorporate stakeholder perspectives.
- **Collaboration and Feedback Mechanisms:** Foster cross-team collaboration and feedback on AI and RAI projects.
- **Responsible AI Practice Development:** Implement team-level and enterprise-wide AI strategy guidance, ensuring direct reporting lines to the C-suite for transparency and accountability.
- **Upskilling:** Commit to comprehensive training and education on AI and RAI for all employees.
- **RAI Collaboration:** Engage with industry groups, academic institutions, and stakeholders to contribute to and align with industry standards for responsible AI.

Acknowledgments

This document was prepared by RAI Institute staff and Monika Viktorova. Insights in this work stem in part from the multistakeholder and multidisciplinary effort co-led by Ashley Casovan (former Executive Director of the RAI Institute) and Monika Viktorova (RAI Contributor) in 2020 under the name 'Bioethics Working Group'. We thank the participants of the working group for their time and invaluable insights.

If you would like to provide feedback on this resource, utilize this guide as part of your organization's journey or help develop this resource further, please reach out to policy@responsible.ai.

Responsible Artificial Intelligence Institute

Advancing Trusted AI