AI Governance Maturity Models 101: Assessing and Improving AI Governance Over Time

TL;DR

AI Governance Maturity Models measure your progress in implementing best practices for AI governance. Conduct assessments using maturity models to chart a clear course towards more stable and reliable AI risk management.

Introduction

AI systems are a powerful new wave of technologies that present a plethora of business opportunities. But with new opportunities come new risks. While there are now several comprehensive and widely adopted frameworks for Responsible AI (RAI), Artificial Intelligence Governance and AI Risk Management, such as the EU AI Act and the NIST AI Risk Management Framework, how can companies assess their own policies and practices in light of these broad frameworks to lower their risks while harnessing the capabilities of AI systems?

AI Governance Maturity Models (or Responsible AI Governance Maturity Models) are designed to answer this question. A maturity model is a measurement tool for assessing how developed an organisation's capabilities and practices are within a given business function. For example, there are industry-standard maturity models in areas like cybersecurity and HR. Naturally, an AI Governance Maturity Model applies this kind of framework in the field of AI governance.

These models are important tools for evaluating how effectively a business is implementing industry-standard best practices and regulations. For example, the maturity model based on the NIST AI Risk Management Framework gives a detailed questionnaire on all facets of AI governance, such as risk measurement, documentation and monitoring. It also includes a scoring procedure to get a concrete sense of which areas of AI governance within a business need to be improved and how to do this.

Key Takeaways

1. Reliably track your organisation's AI governance progress by consistently conducting structured assessments using maturity models.

2. Thoroughly document the assessment process and results, including evidence supporting assessment criteria verdicts. Involve knowledge experts from across the organisation throughout the process.

3. Use the results of your assessments to develop clear and detailed improvement plans to increase your organisation's AI governance maturity.

Understanding AI Governance Maturity Models

As described, AI Governance Maturity Models are measuring devices for assessing an organisation's progress in implementing consensus AI governance guidelines and recommendations. While different models take on different structures, some common components include the following.

Assessment Criteria

The assessment criteria describe the dimensions along which AI governance maturity is assessed. They may take the form of questions that need to be answered, statements to evaluate for degree of accuracy (such as "Completely Accurate" or "Somewhat Accurate") or rubric descriptions that are placed within tiers (such as "Optimised" or "Initial Stages").

The NIST-based maturity model, for example, takes the approach of giving statements and sub-statements about various areas of AI governance, which are then scored on a scale of 1 - 5 for the degree of accuracy. One such statement in AI transparency, for instance, states, "We document the system risk controls, including in third-party components."

The Data Ethics Maturity Model, on the other hand, gives rubrics for different areas of data ethics containing detailed overall evaluations of company policies and procedures within those areas. The evaluator then chooses which description most closely fits the company being evaluated on a scale from "Initial" to "Optimised".

Scoring and Aggregation

The evaluations on the individual assessment criteria are aggregated and scored, with many maturity models grouping the final scores into tiers or levels of maturity. The exact scoring procedure differs between maturity models. The NIST-based maturity model includes methods for aggregating along the NIST framework's "Responsibility Dimensions," which include such values as fairness, privacy and human oversight, or along the "NIST Pillars," which are the AI governance tasks "MAP," "MEASURE," "MANAGE" and "GOVERN."
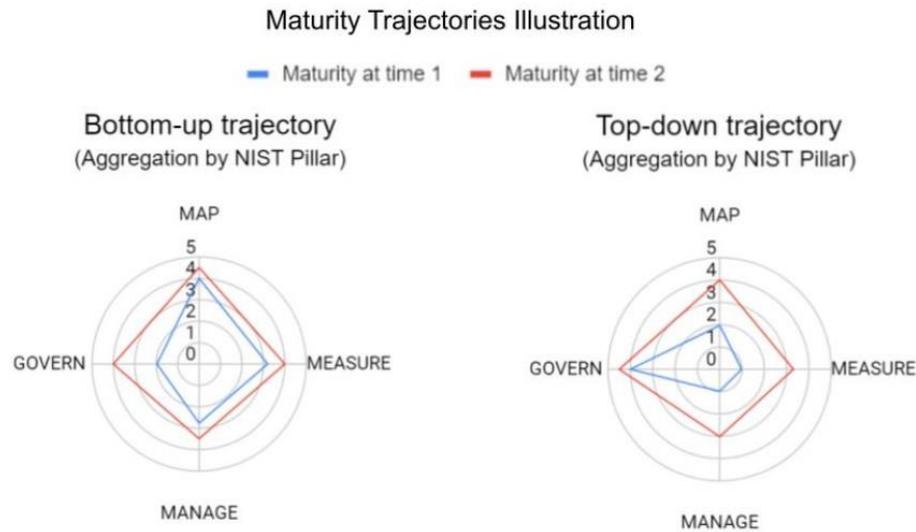
Figure: Maturity progress trajectories. Source: NIST AI RMF Maturity Model (Dotan et al.)

Improvement Pathways

While all maturity models can help improve AI governance by pointing out areas for improvement, some maturity models also offer specific suggestions for implementing improvements. For example, the AI Ethics Maturity Continuum gives an "Action for Improvement" within each ethical value, including different actions depending on the level of value maturity and business stage.

The Importance of AI Governance Maturity Models

The goal of an AI Governance Maturity Model is to help mitigate an organisation's AI risks through effective governance. The following are three specific ways in which these models achieve this goal.

Structured Assessment

It's obvious that assessing AI governance practices is key to managing AI risks. Adopting a structured approach to assessment by using maturity models offers various advantages over a more ad-hoc method of assessment. With a comprehensive maturity model, you are less likely to overlook any aspects or areas of AI governance. Moreover, a structured approach is documented and repeatable, allowing progress in AI governance to be reliably tracked over time.

Continuous Improvement

Maturity models identify areas of weakness in AI governance and risk management, highlighting improvement pathways and enabling businesses to take actions to address these vulnerabilities. With structured assessments being performed on a consistent basis, progress towards AI

governance maturity is measured reliably and which policy changes are most effective becomes transparent.

Benchmarking and Comparison

With the wider adoption of AI governance maturity models, businesses will have a standard measure to compare their AI governance approach with that of comparable industry peers. This incentivizes less mature organisations to accelerate the implementation of best practices and provides evidence for more mature organisations of the effectiveness of their approach to AI governance.

Levels of AI Governance Maturity

AI Governance Maturity Models often define tiers, or levels, of AI governance maturity and readiness. While various models define the levels differently, a useful example comes from the Data Ethics Maturity Model, which defines five levels of maturity. In order of increasing maturity, these are **Initial**, **Repeatable**, **Defined**, **Managed** and **Optimising**.

1. **Initial**: Relevant governance practices are either nonexistent or completely ad-hoc and informal, with no documentation or oversight.
2. **Repeatable**: Relevant governance practices exist but are determined individually by distinct teams and business units, with no organisation-wide standards.
3. **Defined**: Relevant governance practices are documented and standardised company-wide, but may not be fully implemented or adopted within all areas of the organisation.
4. **Managed**: Relevant governance practices are documented, fully implemented and monitored to measure effectiveness and compliance.
5. **Optimising**: Relevant governance practices are documented, fully implemented, monitored and measured, and are continuously improved, updated and adapted to align with strategic initiatives and changing regulatory frameworks.

Using AI Governance Maturity Models

AI Governance Maturity Models are effective tools for improving overall AI governance posture when used properly. The following describes the different uses of these models and the best practices for each use.

Conducting Assessments

The main function of an AI Governance Maturity Model is conducting assessments of organisations' AI governance maturity. Here are tips for evaluators to do this effectively:

- **Thoroughly document the assessment process and results.** Leave a sufficient paper trail so that the process can be repeated consistently and the results can be understood in their proper context. If possible, make this documentation public to increase transparency around AI governance and allow industry-wide benchmarking and comparison.

- **Use and document evidence when completing assessment criteria.** The final verdicts on assessment criteria, such as whether a given facet of AI governance falls within the "Managed" tier of maturity, do matter, but it's also important to document what evidence was used in making these assessments. This increases trust in the assessment results for both internal and external stakeholders. It also gives vital details that can be highly important when the results of the assessment are used to improve governance practices.
- **Involve members of the organisation who are knowledgeable on the relevant practices when conducting the assessment.** A wide range of organisation members should be interviewed or otherwise contacted to get reliable and evidence-based information for the assessment. While assessments are often spearheaded by a single risk management- or governance-focused business unit within an organisation, they should involve all business units that implement or are affected by internal AI systems and policies.

Identifying Gaps and Opportunities

The verdicts on individual assessment criteria and aggregate scores for risk areas both help to identify weaknesses in current AI governance practices and opportunities for improvement. Maturity models can uncover a gap in metrics for assessing bias or a lack of documentation concerning data collection practices, for example. Steps can then be taken to address these gaps by implementing bias-related metrics in evaluating AI outputs and developing documentation concerning internal or external data collection.

Developing Improvement Plans

Effective improvement plans fall out of assessments using maturity models once gaps and weaknesses are clearly identified. This is especially true when assessments are conducted effectively by documenting evidence for verdicts and involving a wide range of business units affected by AI governance practices. With specific evidence in hand once the assessment is completed and documented, the evaluators have a clear roadmap for improving AI governance and the organisational knowledge of who can implement each aspect of that roadmap.

Best Practices for Improving AI Governance Maturity

Regardless of the particular weaknesses identified by using an AI Governance Maturity Model, there are some general best practices that help improve overall AI governance effectiveness for any organisation across all facets of AI governance.

Stakeholder Engagement

AI governance policies affect people and organisations both internal and external to your company. It's important when developing and improving AI governance practices to get input and feedback from a diverse body of stakeholders that are, or will be, affected by your practices. Stakeholder engagement can reveal overlooked considerations and bring important voices to the table throughout the governance process.

Regular Reviews and Updates

Consistently performing assessments of your practices using AI Governance Maturity Models means reliable tracking of progress towards governance goals. It also means that governance practices will be responsive to any changes in business strategy, technological developments and regulatory updates in a timely manner.

Training and Education

Regular training and education is necessary both to inform stakeholders of updates to governance practices and to give employees the tools to implement these practices. Evaluators should also be trained on effectively conducting AI governance audits using maturity models. Education helps foster a culture in which AI governance is understood and taken seriously across the organisation.

Challenges in Assessing and Improving AI Governance

Improving your AI governance posture requires knowing the challenges that you are likely to confront and possible solutions. The following are some of the most common.

- **Organisational resistance to change**: AI governance best practices include increasing documentation, measurement and transparency. These activities may be perceived as unnecessary distractions or hindrances in some contexts. To combat organisational resistance, offer training and education that clearly explains the value of governance practices and fosters a culture of continuous improvement.
- **Lack of accurate measures of governance effectiveness**: Many current AI governance frameworks and maturity models emphasize metrics related to the implementation of policies and processes. However, it can be difficult to assess whether these are effective in mitigating risk. Organisations might consider supplementing existing models with assessment criteria related to incidence rates and using statistical measures of fairness and bias in AI systems. External expertise can be helpful in this developing area.

Final Thoughts

Achieving AI governance maturity allows you to harness the exciting upsides of AI technologies while lowering their inevitable risks. AI Governance Maturity Models are a powerful tool to help you get there. A detailed and comprehensive model gives you a structured assessment that can be consistently used to identify gaps and develop clear improvement pathways. With effective use of AI Governance Maturity Models, you will be ready for the unexpected changes and developments AI brings.