



INTERNATIONAL
INSTITUTE FOR
ANALYTICS™

RESEARCH BRIEF
RESEARCH & ADVISORY NETWORK



CLIENT ONLY

A Comprehensive Guide to Responsible Analytics Governance

GIOVANNI LEONI, GLOBAL ENTERPRISE LEADER, RESPONSIBLE ANALYTICS

FEBRUARY 2023

Today, many global enterprises are on a journey to become more efficient, sustainable, and relevant in an omnichannel reality. In this omnichannel reality, digital technology is widely used in business operations. This requires a new mindset to ensure that your organization stays true to its corporate values in the digital age. Therefore, digital ethics becomes an important and fundamental element of the digital journey.

Responsible analytics is about delivering data insights in an ethical, transparent, and reliable way. Responsible analytics is also one of the key capabilities needed to build and achieve data-enabled value chains. Responsible analytics governance guides the capability development and supports co-workers to make values-based decisions around data and technology. Global organizations need to be prepared to respond to new customer demands and regulatory requirements across the world.

This research brief lays the foundation for operationalizing analytics governance across all business functions in the enterprise. This guide will help you:

- Frame analytics governance within the context of your business model and values
- Build out an analytics program that is a natural part of digital ethics
- Consider the dimensions of process, people, and technology when creating your analytics governance program

Analytics Governance Within the Context of Company Values

To stay competitive in today's landscape, most organizations are not just concerned with the bottom line, but rather trying to find better ways to get things done and to bring out the best in themselves and others. Digital ethics is the next logical step in this journey: doing the right thing for people, the environment and society, and finding new and better ways to navigate data and digital technologies. Think about it: Every digital or data-driven interaction or decision represents how employees stretching across an organization build the company together. Consequently, you need to think about how you want your digital and data-driven interactions to represent your company's values.

At its core, it would be ill-advised to go any further with your analytics program until it is driven by your company's values. This exercise forces your team and senior leadership to address *why* your organization needs an analytics program in the first place – and aspirational goals for such an endeavour.

For example, does your organization believe the digital world should be designed responsibly to benefit people first and foremost? Does it believe in honest and genuine transparency for the good of society? Does it want to be a brand leader in digital trust and respect of customers and co-workers? These are big questions that can no longer be avoided given the increasing amount of scrutiny on how companies use data and digital technologies. Further, it is now the norm for customers to align with organizations that respect their data and digital interactions. Negative incidents through misuse, underuse, and abuse of digital technologies have also created a push toward greater regulation.

So, the task at hand for data and analytics leaders is thinking – and acting – beyond regulatory requirements. In other words, instead of taking a “defensive posture” to digital ethics, the spirit of this guide is rooted in taking proactive steps to protect your brand, your customers, and to *promote* the values of your company.

Establishing a Policy on Digital Ethics

Once you’ve framed analytics governance through the context of company values, the next logical step is establishing a policy that guides your team in living these values in a digital age. Digital ethics adds the element of moral and conscious decision-making to a company’s activities with data and digital technologies. It provides a perspective that helps to guide your business decisions to choose what is good, morally right and ethically justifiable when developing digital products or services, automating processes, acquiring and using data, and designing and deploying algorithms or robotics. For practices regarding data and digital technologies to be trustworthy, they must be legal, ethical, secure and designed in accordance with your company’s values. The best governance polices are:

1. Driven by clear commitments
2. Give high-level direction and guidance to make conscious decisions around data and digital technologies
3. Enable digital ethics by design

Uniting around your company values and having clear and common principles creates the foundation on where your enterprise can act together with a common agenda, and a common reference to what is “good.”

Together, you define the meaning of what it means to be “ethical” through a clear governance framework that covers people, process, technology, and data.

Digital ethics principles are unique to the company, and it is up to you to define these principles in a collaborative manner. Regardless of the exact language or length, “handling data ethically” should be positioned as a core aspect of your data and analytics teams’ frameworks and considered a core capability of your business model.

Integrating Ethics into Data Frameworks and Business Models

The ideal state of your analytics governance is that the principle of “handling data ethically” spans across the complete data framework of your organization and enables your company’s transition to a data-enabled value chain.

Sound data ethics principles will:

- Protect, secure and develop the enterprise long-term
- Enable company co-workers to make data-driven decisions
- Connect and drive insights from data to adapt to the ever-changing world around us

To truly integrate ethics into all corners of the business, you want to map out how data ethics principles and practices inform both fundamental business capabilities and higher-level interdependencies. For example, is ethics integrated into information management, information security and data privacy, data and analytics, and innovation hubs? For higher tier capabilities, consider how ethics drives policy management, design, auditing, and responsible analytics and AI.

Key questions to consider:

- Do you manage policies on data and algorithm ethics based on business needs, ethical principles, and socially good outcomes?
- Do you assess and advise business initiatives and projects on their ethical data and algorithm lifecycle management including acquisition, development, and usage and monitoring?
- Are you able to execute audits of algorithmic processes and meet up with legislative demands at low cost?
- Can you govern the lifecycle and procurement of analytics and AI models and are you developing practices/tooling in a way that aligns with ethical principles?

No one will have all these aspects of digital ethics mastered. The point is building maturity in these capabilities and creating an interdependence between your company values, data, and business needs.

Responsible Analytics Governance: The Algorithmic Lifecycle

Responsible analytics governance strives to enable the business to use algorithmic processes in adherence with the digital ethics principles of your enterprise. It provides the support, structure, and guidance around processes, people, and technology and data involved in algorithmic processes. It supports business teams to take ownership of the ethical aspects of algorithmic processes and provides guidance for taking concrete actions and decisions in daily operation, both tactical and strategic. Responsible analytics governance is enabled through:

Process integration into everyday business

processes. Responsible analytics governance means ethical considerations should be part of the business process by design. Through checkpoints and steppingstones to ensure the algorithmic processes meet your company's standards, you ensure the trustworthiness of data and algorithmic processes in the most resource-efficient manner. These steppingstones are incorporated into the existing businesses processes, both those designed, developed and deployed internally and for all data, products and solutions you procure.

People sharing responsibility and making a

difference. You are building a distributed model of responsible analytics that enables everyone to work responsibly. Every employee is an important contributor to responsible analytics, both contributing themselves and to proactively seek solutions with others. Everybody has a part to play and must take ownership to act with integrity and care. This is the most flexible, effective and efficient way to do responsible analytics – fitting it into your company's values.

Technology augments process, people and data.

You want responsible analytics to govern all algorithmic processes across the enterprise. To make this effort scalable, affordable, flexible and possible, technology is an enabler to bring it all together. Technology will come in both the form of a platform for responsible analytics governance and tools that support responsible analytics processes. The responsible analytics governance platform support to tangibly enable governance on operational level through registry of algorithmic processes, impact and risk assessment, collaboration, communication, auditing and reporting. The tools could either be part of the platform, stand-alone solutions or integrated in products or solutions that have algorithmic processes.

The following framework gives structure to and guides the processes, people, technology and data involved in algorithmic processes.

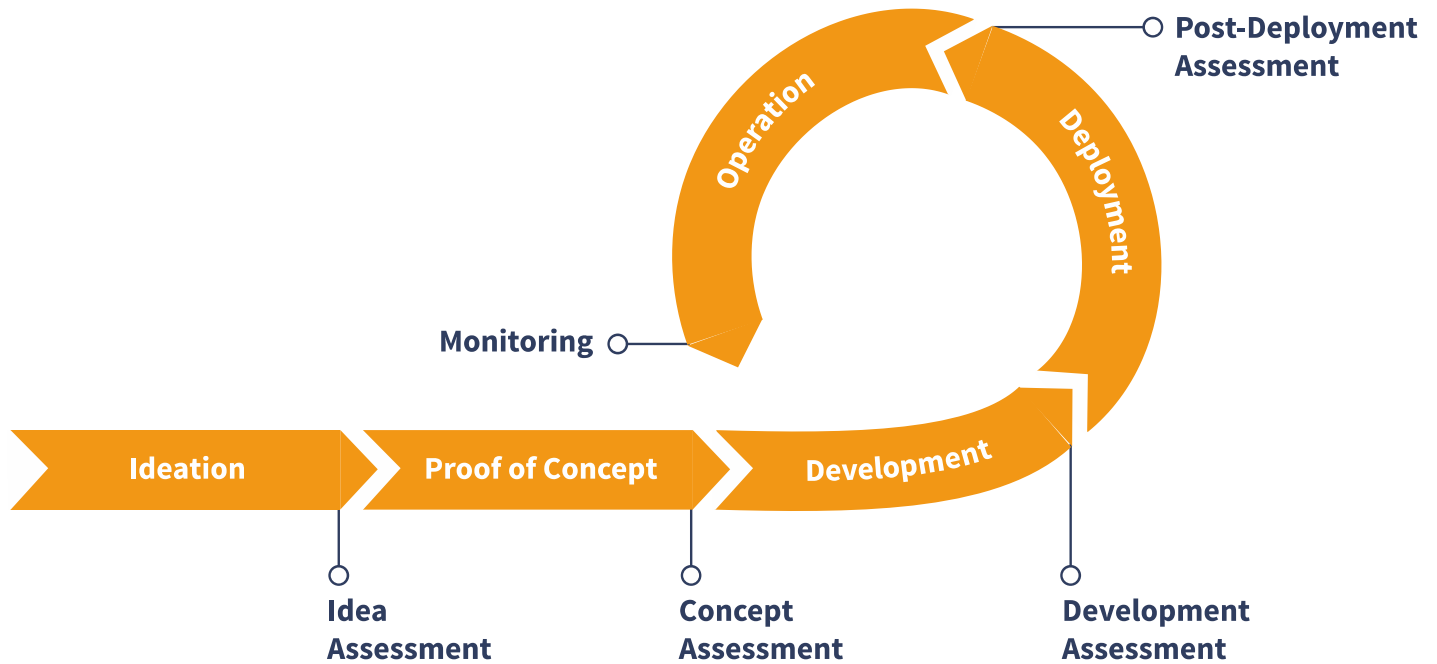
These tools could be very specific for the processes it supports, as identifying and addressing the bias and discrimination with fair machine learning technology and practices, all serving the purpose to support the outcome of responsible analytics.

Data is the highly valued asset that makes everything work. Responsible analytics governance promotes the idea of accountability of data through the capabilities of data ownership, data quality, data security and data privacy. Securing data as an asset and enabling data literacy, data and insight life cycle management and in the end analytics, is critical to partner it with the approach of responsible analytics. There is an interdependency between the capabilities of securing data as an asset and responsible analytics, where both contribute to deliver to common interests.

The following framework gives structure to and guides the processes, people, technology and data involved in algorithmic processes. It avoids detailing exactly how it should be applied in every algorithmic and business process.

The reason for this is that the rate of change is so high for any business environment that being prescriptive would deem the framework irrelevant before it saw the light of day. Rather, this framework is designed to give clear direction for responsible analytics governance of all algorithmic processes, and to then be integrated into its respective business process on a case-by-case basis.

We will start with the governance **process** which is integrated into the development lifecycle for both internally and externally developed algorithmic processes. Thereafter, we discuss the **people** perspective with roles and responsibilities to execute the processes. Finally, we identify the enabling **technology** to support the process and people in a scalable and efficient way. Data is present across all domains but is not focused upon as the data is primarily a precondition for successful responsible analytics, treated as an asset at all times.



PROCESS

Responsible analytics governance requires a close integration of digital ethics considerations and actions at all stages of the algorithmic processes, from ideation to development and operations. These stages are encompassed in a governance framework that supports the registering, assessing, monitoring, auditing, and reporting of the algorithmic processes, while taking actions that strengthen the governance framework over time.

Responsible analytics governance covers the entire lifecycle of internally developed or procured algorithmic processes. At a high level, algorithmic processes can be divided into five stages: 1) Ideation, 2) Proof of Concept, 3) Development, 4) Deployment, and 5) Operation. There are a couple of important considerations to consider when using this simplified model of algorithmic processes and responsible analytics governance integration.

- **Encompassing all algorithmic processes:** The process should cover both those developed within the enterprise and those products, solutions or data with algorithmic processes that have been procured. Furthermore, it includes a broad range of algorithmic processes, from the most rule-based to advanced use of machine learning as well as administrative and operational algorithmic processes.
- **Full integration with the algorithmic process lifecycle:** Responsible analytics governance becomes a natural part of decision-making along the lifecycle by offering tangible guidance as an ever-present dimension.
- **Catering to an iterative process:** Requirements often emerge during operation and can be developed as new features or improvements of existing algorithmic processes. This may require new and alternative digital ethical impacts to be assessed.

- **Entering the responsible analytics governance framework at any given stage in the lifecycle:**
Today there may be hundreds of existing algorithmic processes running at different lifecycle stages and in different parts of an organization. Therefore, it is important to create a low threshold for these algorithmic processes to enter into the responsible analytics governance framework at any point in their lifecycle. You can support this through clear expectations for innovation at every stage of iteration.

Assessments at Different Stages of the Algorithmic Lifecycle

Over time, responsible analytics should become natural when the analytics organization initiates a new development or procurement of an algorithmic process. Until that day, responsible analytics governance needs to be applied in three different scenarios.

Scenario #1: Support the application of analytics governance at the stage of ideation. The goal here is to integrate governance into the business process as tightly as possible and illuminate the added value of having responsible analytics present at every step of algorithmic development.

Scenario #2: Support the application of analytics governance at any stage of development – from ideation to deployment to operations. In this scenario, the algorithmic process needs to enter the framework, regardless of where you are in the process.

Scenario #3: Support the mindset of responsible analytics in everyday decision-making. All told, this is the most important application of analytics governance, especially at the early maturity stage of responsible analytics implementation. This scenario is

where you will find the most impact – at the level of every employee driven by the spirit of your organization’s principles and values around digital ethics.

To achieve these goals, the algorithm process lifecycle should pass through **five assessments**.

- **Idea Assessment:** An early-stage evaluation of the identified need, anticipated use case, and envisioned technological solutions and needed data.
- **Concept Assessment:** An evaluation of the refined design concept, anticipated use case and benefits, clear technology requirements and specifications and identified and defined performance and digital ethics metrics that the final solution should be developed to meet.
- **Development Assessment:** An evaluation of the final solution, its final design and performance results, capabilities and limitations, as well as intended use case and anticipated benefits, prior to deployment.
- **Post-Deployment Assessment:** An evaluation of the final solution and its impact on the real world, its integration into the business process and teams, its performance and use in practice.
- **Monitoring:** Continuous and regularly iterative review of the performance, use and impact of the algorithmic solution after its deployment.

These assessments are carried out in different stages of the algorithmic process lifecycle. In other words, each assessment can gradually encompass more information and be more specific.

Each assessment covers **four governance review dimensions** and some of them can extend to a fifth optional dimension:

A proportionate algorithmic process should bring more positive than negative outcomes across the dimensions of purpose, impact, performance and risk.

- **Purpose & Impact:** This dimension looks at the purpose behind the innovation, the expected benefit, and considers the impact, both positive and negative, of the algorithmic process.
- **Performance:** This dimension focuses on data and technology. What are the preconditions and requirements for the solution, as planned at that stage, to function as intended and as desired from a technical perspective? It also covers the necessary steps of technical work, engagement with stakeholders, validating results and plans to ensure that the final solution functions and is used as intended.
- **Risk:** This dimension seeks to assess the potential risks that may arise from the design, procurement, and use of the solution, as it is planned. It covers risks arising from the performance of data and technology, as well as from features of the business and human context in which the solution will operate. Where risks are identified, this dimension will ask you to identify mitigation techniques to eliminate, transfer or mitigate the risks.
- **Proportionality:** This dimension is a step where the anticipated benefits and risks of the solution, as it is planned at this stage, are balanced. A proportionate algorithmic process should bring more positive than negative outcomes across the dimensions of purpose, impact, performance and risk. In addition, some risks may be considered unacceptable, no matter how great the benefits of the algorithmic process.
- **Procurement (Optional):** This dimension is optional and should be included in the proof of concept or development/deployment assessments whenever the algorithmic process that is assessed is procured from an external vendor. The governance review dimension of procurement seeks to provide insights into vendors' compliance with your company's values and principles.

Below we illustrate the five assessments in greater detail.

IDEA ASSESSMENT



Governance Review Dimensions

Purpose & Impact	Understand the business purpose of the algorithmic process, how it aligns with your company values, digital ethics principles, and its potential impact on users and society.
Performance	Identify the prerequisites for developing the algorithmic process, including but not limited to: necessary data sourcing, data quality, technical capabilities, and skills.
Risk	Identify potential digital ethical risks of the algorithmic process and its intended use. Assess the risk level and propose risk mitigation accordingly.
Proportionality	Analyze the proportionality of the idea of the algorithmic process through the lens of purpose, performance complexity, risk level and impact.

Outcomes

- Registration of high-level description of the use-case and envisioned solution, including preferred data, technical performance requirements, and capabilities, and limited documentation of the governance review dimensions.
- Ethical considerations, including impact and risk review, risk mitigation steps, and proportionality review documented to support future steps in the algorithmic process lifecycle.

At the stage of ideation there is an identification of opportunities with new data, products or algorithmic solutions with the presence of algorithmic processes. The design specifications, requirements for the algorithmic process and its intended use and impact on the business start to develop at a high level. The purpose of the idea assessment at this stage is to pre-empt and identify potential positive and negative impact from a holistic perspective (from individual users to society at large). If necessary, changes can then be made to adapt the idea or the intended use case. The idea assessment covers four governance review dimensions and has two outcomes.

CONCEPT ASSESSMENT



Governance Review Dimensions

Purpose & Impact	Understand the business purpose of the algorithmic process, how it aligns with company values and digital ethics principles and its anticipated impact on people, environment and society.
Performance	Identify the feasibility of developing the algorithmic process, including, but not limited to, data sourcing, data quality, technical capabilities and skills necessary. Define and test a limited set of metrics to measure the algorithmic process performance from a digital ethics perspective.
Risk	Identify potential digital ethical risks of the algorithmic process and its intended use, assess the risk level and propose risk mitigation accordingly.
Proportionality	Analyse the proportionality of the algorithmic process, keeping in mind the aspects of purpose, performance complexity, risk level and impact on people, planet and society.
Procurement if externally sourced	If sourcing data, product or solution that is based on algorithmic processes externally, then the vendor will be evaluated on their alignment to company values, digital ethics principles, and compliance to procurement standards.

At the stage of proof of concept there is an exploration of opportunities with new data, products or solutions with the presence of algorithmic processes. Design specifications or requirements for the algorithmic process are formulated, although limited in scope or time, and with defined planned impact on the business

in its context. The purpose of engaging in a concept assessment at this stage is to identify potential positive and negative impacts of the proposed algorithmic process on people, planet and society and adapt the idea and conceptual design of the algorithmic process if necessary. If an idea assessment has been done already, then the concept assessment only complements the idea assessment with more details and revisits the previous assessment if the planned algorithmic process has substantially changed. At this stage there will be tangible aspects to assess, such as data and initial pilot algorithms, but it is still early to have complete understanding of impact, proportionality, risk and performance of the final solution. The concept assessment covers the four governance review dimensions, adding the procurement review if the data, product or solution is based on an algorithmic process that is procured. The assessment has two outcomes.

Outcomes

- Registration of a description of the use case and envisioned solution, including needed data, technical performance requirements, capabilities, and performance metrics, relevant to digital ethics and documentation of the governance review dimensions, adding procurement review if applicable.
- Ethical considerations, including impact and risk assessment, risk mitigation steps, proportionality review, procurement assessment, documented to support future steps in the algorithmic process lifecycle.

DEVELOPMENT ASSESSMENT



Governance Review Dimensions

Purpose & Impact	Understand the business purpose of the algorithmic process, how it aligns with your company’s values and digital ethics principles and its anticipated impact.
Performance	Identify the feasibility of developing the algorithmic process, including but not limited to: data sourcing, data quality, technical capabilities, and skills necessary. Define and test meaningful metrics to measure the algorithmic process performance from a digital ethics perspective. Iteratively test the solution against set performance targets, adapt design, and assure the desired level of quality.
Risk	Identify potential digital ethical risks of the algorithmic process and its intended use, assess the risk level and propose risk mitigation accordingly.
Proportionality	Analyze the proportionality of the algorithmic process through the lens of purpose, performance complexity, risk level and impact.
Procurement if externally sourced	If the data, product or solution that is based on algorithmic processes is externally sourced, then the vendor will be evaluated on their alignment to company values, digital ethics principles, and compliance to procurement standards.

At the stage of development many opportunities are explored for using new data, products or solutions with the presence of algorithmic processes. There are clear design specifications and requirements for the algorithmic process and with defined impact on the

business in its context. The purpose of engaging in a development assessment at this stage is to identify potential positive and negative impacts of the algorithmic process and adapt the design and development decisions if necessary. If a concept assessment has been done already, then the development assessment only complements the concept assessment with more details and revisits the previous assessment if the algorithmic process has substantially changed. There will be tangible aspects to assess, as data and algorithms, to understand the proportionality, risk, or performance of the algorithmic process. The development assessment covers four governance review dimensions, adding the procurement review if the data, product or solution is based on an algorithmic process that is procured. The assessment has two outcomes.

Outcomes

- Registration of description of the use case and algorithmic solution – underlying logic, architecture, intended use, performance results along different metrics, capabilities and limitations, and documentation of the governance review dimensions, adding procurement review if applicable.
- Ethical considerations, including impact and risk assessment, risk mitigation steps, proportionality review, procurement review, documented to support future steps in the algorithmic process lifecycle.

POST-DEPLOYMENT ASSESSMENT



Governance Review Dimensions

Purpose & Impact	Understand the business purpose of the algorithmic process, how it aligns with company values and digital ethics principles and its <u>actual</u> impact.
Performance	Identify and test meaningful metrics to measure the algorithmic process performance from a digital ethics perspective and relevant level of quality assurance. Review and assess the performance of the algorithmic process in practice. Set up periodic reviews to ensure the performance remains as desired long-term and the data inputs and outputs remain as expected. Evaluate feedback from users and affected stakeholders of the solution, if appropriate.
Risk	Identify potential digital ethical risks of the algorithmic process and its intended use. Assess the risk level and propose risk mitigation accordingly.
Proportionality	Analyze the proportionality of the algorithmic process through the lens of purpose, performance complexity, risk level and actual impact.
Procurement if externally sourced	If the data, product or solution that is based on algorithmic processes is externally sourced, then the vendor will be evaluated on their alignment to company values, digital ethics principles, and compliance to procurement standards.

At the stage of deployment, algorithmic processes are embedded in business processes and begin impacting the business and people daily. There are clear design specifications of the algorithmic process and its input and output data. The purpose of engaging in a post-deployment assessment at this stage is to identify potential positive and negative impacts of the algorithmic process at a holistic level – from individual users to society. If a development assessment has been done already, then the post-deployment assessment only complements the development assessment with more details and revisits previous assessments if the algorithmic process has substantially changed. There will be tangible aspects to assess, such as data and algorithms, framework of roles and procedures around the use of the algorithmic process in practice, and some results and logs, which allow to understand the impact, proportionality, risk, and performance of the algorithmic process on the business.

The development assessment covers four governance review dimensions. If an externally procured algorithmic process has undergone the development assessment to ensure its quality and performance meets company requirements, it can be considered as “vetted” and the same post-deployment assessment can be applied to internally developed and externally procured solutions at this stage. However, if a procured solution is already deployed and operating within the business, and has not yet undergone any prior assessment, a complete procurement review should be added to the post-deployment assessment. The assessment has three outcomes.

Outcomes

- Registration of description of the use case, solution and the context of its deployment and use, roles and responsibilities, channels for human oversight and interference, logs and mechanisms for recording performance and impact long-term, and documentation of the governance review dimensions, adding procurement review if applicable.

- Clearly define the key points of reference in each review dimension to be monitored and how often they should be reviewed.
- Ethical considerations, including impact and risk assessment, risk mitigation steps, proportionality review, procurement assessment, documented to support future steps in the algorithmic process lifecycle.

MONITORING



At the stage of operations, algorithmic processes are a part of running business processes and impact the business and people daily. To ensure human autonomy over the algorithmic processes, there needs to be continuous monitoring of key points of reference in each review dimension as defined in the post-deployment assessment (purpose, impact, performance, risk and proportionality).

Some points of reference can be automated. These will be automatic triggers, for example, if the performance of the algorithmic process deviates from expectations, as in the case of data drift. In addition to automatic monitoring, there will be manual ad-hoc deviation handling if the data inputs, outputs, or if the algorithmic process substantially changes, creating a need for a new post-deployment assessment of purpose, impact, risk or proportionality. This also applies to vendors if they change their approach after post-deployment assessment analysis.

Outcomes

- Dynamic and up-to-date responsible analytics governance status of all algorithmic processes running across the enterprise.
- Pre-conditions and requirements for keeping records to ensure the auditability of individual algorithmic processes, their registration and assessments.
- Pre-conditions for being able to report, internally and externally, on responsible analytics governance and the individual algorithmic processes the enterprise has registered and assessed.
- Recommendation for renewal of post-deployment assessment if deviation occurs.
- Periodically review and update post-deployment assessment, independent of deviations.

Additional Steps in the Algorithmic Lifecycle

There are two activities in the responsible analytics governance process that could be relevant to engage in, depending on the context. If the data, product or solution is externally procured, then a procurement review needs to take place. Furthermore, if monitoring key points of reference can be integrated into the responsible analytics governance platform, then this monitoring should be applied.

PROCUREMENT REVIEW



If externally sourced data, products, or solutions are based on algorithmic processes, then the vendor will be evaluated on the alignment to company values, digital ethics principles, and compliance to procurement standards.

The alignment to values is not about being a perfect match, but rather that both parties can agree on important viewpoints without limiting the nature of the agreement or actions on both sides. It is about finding common ground and building a trustful collaboration beyond the specifics of the contractual agreement. This would also encompass a tangible review of the vendor's digital ethics point of view, using your company's digital ethics policy and responsible analytics approach as a reference for what "good" means. It is not about approaching it as a compliance exercise, but rather agreeing on high-level values and principles and identifying digital ethics risks that come with the partnership.

Regarding governance review dimensions (purpose, impact, and so on), these are considered with the same scrutiny as if you would have developed the algorithmic process internally. Having the same expectations, without having the same natural access to data and algorithms, creates a higher demand for

creating the right contractual pre-conditions for transparency, robustness, and data and algorithmic portability. As business processes depend on the vendor's external data source, product, or solution, the contractual agreement needs to prioritize the integrity of the business process, regardless of the status of vendor partnership.

Outcomes

- Procurement review documented.
- Key points of reference to be monitored clearly defined.
- Ethical considerations documented to support future steps in the algorithmic process lifecycle.
- Communication and collaboration with legal, procurement, and compliance functions.

RESPONSIBLE ANALYTICS GOVERNANCE

PLATFORM INTEGRATION

Ideation

Proof of Concept

Development

Deployment

Operation

For the sake of cost-consciousness, simplicity and efficiency, we should always strive to automate what we can and integrate algorithmic processes into a responsible analytics governance platform. The purpose of the integration is not to move data from the systems with algorithmic processes, but rather to focus on metadata that triggers actions along the responsible analytics governance process. The integrations should support four main activities: 1) Registering, 2) Monitoring, 3) Reporting, and 4) Collaboration.

Registering an algorithmic process can be more easily supported when the underlying system with algorithmic processes indicates whenever a new algorithmic process is being introduced. By triggering this indication to the responsible analytics governance platform, it also triggers subsequent actions within the responsible analytics governance process.

Monitoring algorithmic processes is best supported when deviations are identified automatically and triggers indications to the responsible analytics governance platform. What will be triggered is dependent on each use case and determined by key points of reference during the post-deployment assessment and then monitored for deviation. This could be anything from scope of data used, expected level of performance of the algorithms, or other changes that might create a deviation of assessed impact, purpose, risk, or proportionality.

Reporting should be carried out ad-hoc or periodically. Integrations between the responsible analytics governance platform and dependent systems should eliminate manual administration and promote automated processes. The dependent systems could either be internal (e.g. risk compliance systems) or external (e.g. national or multinational governance bodies).

Collaboration is best done when the trigger of change supports an integrated way of working across the enterprise, providing the right information on the right topic is delivered to the right person in an automated way. By integrating the responsible analytics governance platform into data and analytics monitoring tools (and project management tools), your analytics governance is engaging with the workers as opposed to the workers engaging with analytics governance.

Outcomes

- Evaluate and design the integration between the responsible analytics governance platform and the dependent system.
- Develop, deploy, test, and monitor integration.
- Deliver the right information to the right person at the right time.

The RACI Model Applied to Responsible Analytics Governance

Operationalizing the responsible analytics governance process is first performed by people, and then supported with technology. We do not want to be too prescriptive in telling people what they should do and when, but we still should give direction to main activities and responsibilities for certain roles that engage in the responsible analytics governance process. Detailed process descriptions are designed alongside the implementation of responsible analytics across the organization and dependent on the demands of each algorithmic process.

ROLES AND RESPONSIBILITIES

Everyone across the enterprise has an individual responsibility of applying the mindset of responsible analytics on everyday decision-making. Literacy, training, good processes and technology supports this effort, but it all starts with this mindset: We want the outcome of our actions to be aligned with the digital ethics principles of the company. That said, there are key roles that are more engaged in the responsible analytics governance process and here we will cover their main responsibilities.

The **business owner** takes the overall ownership and accountability of the algorithmic processes applied in their business domain. This ownership is most natural with a domain-based design driving your data framework. Most importantly, the business owner plays a key role in understanding and communicating business needs, the actual use case and context in which the intended algorithmic solution is to be used, and represents the needs of the intended end-users and business teams. Finally, the business owner guides long-term oversight to ensure that algorithmic processes are applied in a responsible way.

The **data and analytics manager** is responsible for ensuring that their team is equipped with the right level of literacy and training around responsible analytics, supporting responsible analytics governance processes, and enabling technology. The data and analytics manager must be informed of the algorithmic process and assessments related to purpose, impact, performance, risk, and proportionality – and support action if necessary.

The **product manager** sees that algorithmic processes are developed and executed according to the requirements of responsible analytics. This includes engagement in assessments along the different stages of the algorithmic process lifecycle and making sure action is taken to meet desired outcomes.

The **data scientist** has a key responsibility in designing analytics solutions and machine learning models which fulfils the business requirements and digital ethics principles, such as fairness, explainability, interpretability and robustness. The data scientist is also responsible for documenting and communicating the technical information.

The **data engineer** is key to identifying data quality, documenting and communicating deviations, or shifts in the way data has been recorded over time which undermines its utility.

The **digital ethics lead** is accountable for the management of responsible analytics governance through the development and implementation of governance processes and tools, supporting people in their roles, and facilitating the use of technology in the form of a responsible analytics governance platform. Furthermore, the digital ethics lead is responsible for having an active role as a partner along the responsible analytics governance process, engaging in active dialogue in assessment from end-to-end, and recommending courses of action.

Users give a reality check to how the algorithmic process could impact people. This could be anything from how they perceive the algorithmic process in general to how they want to use the technology, their degree of understanding of the tools, and how the solution would affect their work or life in practice.

Important RACI Considerations:

- **Accountability always falls on the business owner:** This holds true if the outcome is augmented by algorithmic processes and this ownership does not deviate regardless of level of detail of the action or who is responsible for taking that action. This requires the need to create mature processes that enable transparency, insight, and conscious leadership.
- **Who needs to be informed is case-by-case:** Who needs to be informed of algorithmic processes and their status depends on the business context. Assigning who is “informed” in the RACI model cannot be prescriptive, but rather evaluated on a case-by-case basis. Regardless of business context, the outcome should be the same: Keeping relevant stakeholders informed at any given time.
- **Responsibilities are assessment-dependent:** During the responsible analytics governance process there are key actions that need to be taken by different roles. These responsibilities flex at different stages of assessment. For instance, data engineers and data scientists contribute deep analytics competence in the ideation stage and assess how data and technology can meet business needs in the development stage. To fully leverage the power of RACI as it relates to analytics governance, explicitly map out responsibilities against current roles and context of your enterprise.

Responsible Analytics Governance and Supporting Technologies

To make the responsible analytics governance process scalable and efficient, you need the right technologies in place. We can break down these technologies into two types:

1. The overall **governance platform** to support the analytics governance process end-to-end. For example:
 - a. Automated interfaces for assigning, submitting, and reviewing assessments
 - b. Intelligent risk detection for newly added data or changes to models
 - c. Visualization and reporting of the overall responsible analytics governance status
2. Specific **tools** to develop features in use cases. For example:
 - a. Detecting bias in the data and assessing the fairness of a model
 - b. Desensitization of personal information in data inputs and outputs
 - c. Enhancing the explainability and interpretability of an analytics solution

These technologies work together to operationalize the responsible analytics governance process. Furthermore, the responsible analytics governance platform and tools should be integrated into your data and analytics platform and other relevant systems with algorithmic processes.

For example, the responsible analytics governance platform should be able to:

- Fetch metadata from data catalogue and lineage tools to manage potential risks in the dataset.
- Connect to the model repository and lifecycle management tools to track the model's performance against digital ethics features such as fairness, robustness, and reliability.
- Connect with other project tools to manage responsible analytics actions and outcomes.

From these technological assumptions, we can lay out the key features of the responsible analytics governance platform and tools, and state how these key features will support the responsible analytics governance process.

RESPONSIBLE ANALYTICS GOVERNANCE PLATFORM

The responsible analytics governance platform is built to support the operationalization of responsible analytics governance. Think of the platform as a common workplace for business owners, product managers, the data and analytics team, and the digital ethics lead. For optimal performance, you want the platform to perform the following activities in relation to governance.

1. **Registry** of algorithmic processes. The first – and most important – step of responsible analytics governance is to register algorithmic processes in the platform. This will provide visibility into which algorithmic processes are currently being explored, developed, and deployed across the enterprise.
2. **Assessments** throughout the algorithmic process lifecycle. The responsible analytics governance platform should provide a user-friendly interface and guidance for assessments at different stages of the algorithmic process lifecycle. The platform serves as the repository of status updates for each algorithmic process, and supports communication, collaboration, and execution of next steps. The platform should intelligently trigger new assessments or request an update when there is new data or when a new model has been added to an algorithmic process. Finally, the governance platform should continuously track the impact of algorithmic processes in business operation.
3. **Risk management** of algorithmic processes. The responsible analytics governance platform should support risk management in three different ways. First, the platform recommends risk categorization when an algorithmic process is registered. This recommendation is based on the specific business scenario, the type of data, and company digital ethics standards and principles. Second, the governance platform detects risks during development, deployment, and operation of an algorithmic process when new data and models are added. Third, the governance platform recommends risk mitigation methods for an analytics solution, based on the business context and detected risks.
4. **Reporting** of algorithmic processes against the responsible analytics governance framework. There is a need for reporting around the responsible analytics governance status of algorithmic processes, both internal and external (when applicable), ad-hoc or periodically. Internal reporting could take the shape of periodic reports, management meetings and governance risk compliance systems. External reporting usually

involves national or multinational governance bodies. Reporting supports responsible analytics governance through transparency, awareness, and actionable insights.

5. **Communication and collaboration** for all relevant stakeholders. Action items should be clearly assigned to responsible persons via project management tools (allowing for full integration into the product team's backlog). Ultimately, progress and outcomes should be shared, reviewed, and approved in a seamless manner to improve the efficiency and scalability of the overall governance process.

RESPONSIBLE ANALYTICS GOVERNANCE TOOLS

You also need relevant tools to achieve digital ethics principles (e.g. fairness, explainability, robustness, and privacy) for specific use-cases. Ideally, the team leverages open source tools while creating an environment of innovation internally. Here are a few examples:

- **Fair** machine learning tools. The key to this type of method is to identify sensitive features (e.g. gender or age) and mitigate systematic bias that is usually introduced in the training data. It is also important to design or optimize a model based on fair objectives.
- **Privacy** preserved data processing and federated machine learning. Sometimes sensitive information is embedded in the data and you want to analyze the statistical patterns without leaking sensitive information. Also, you likely face the challenge of data from distributed systems not being centralized for training models. Differential privacy is a technique to query and analyze patterns from data with sensitive information (e.g. personal information) while protecting the individual's privacy.

- **Explainable** and **interpretable** machine learning. You should always strive for explainability and interpretability and this can be achieved in several ways. Understanding the data's statistical property is an important step before moving into model selection training. During model training, explainable models are typically superior to over-complicated ones. Finally, help business users understand the model's output by providing statistical evidence and uncertainty measurements.

Responsible Analytics Governance Management

Responsible analytics governance is a process and processes fail if not properly managed. Important activities that need to be managed to build necessary capabilities include:

AUDITING

All algorithmic processes within the scope of the responsible analytics governance framework need to be available for auditing. There are few examples of specific requirements for the auditing of AI today, but this is changing given recent developments in national and international legislation around data and privacy. Still, the act of auditing something is based on fundamentals that we can take advantage of and prepare for future regulations. First, there should be clear and ongoing documentation of decisions, evaluations, and considerations. Second, transparency is very important. Third parties (i.e. the auditor) will want to understand what has been done and why. Lastly, there should be clearly defined roles and responsibilities outlining who is accountable, who is responsible, who should be informed, and so on.

REPORTING

The monitoring process and available information in the responsible analytics governance platform will create the pre-conditions for relevant reporting, both internal and external.

LITERACY

This guide is written primarily with digital non-native organizations in mind who are going through a digital transformation. The ambition is that everyone across such an enterprise can apply the mindset of responsible analytics in everyday decision-making. This means there's a high demand to build awareness across the enterprise for what responsible analytics means and how it impacts everyone within the company. Literacy should be built into the communication of all relevant channels, co-worker engagement and, most importantly, through active partnership with the business in applying responsible analytics to algorithmic processes.

TRAINING

Beyond building literacy, you must offer education and skills training to employees across the company. Training and education programs should range from basic literacy/awareness to more advanced analytics governance practices. You need to meet people where they're at and their role in the algorithmic process. The end-goal is to integrate literacy and education in such way that responsible analytics becomes second nature to everyone at the enterprise.

EXTERNAL ECOSYSTEM ENGAGEMENT

Responsible analytics is a nascent field in a corporate context. The high rate of technological and societal change means it is critical to connect with and learn from peers, the industry, private enterprise best practices, governmental bodies and academia. Again, responsible analytics governance is not just about compliance. It's about driving innovation and best practices and contributing to the global narrative of why responsible analytics is important for all companies and society at large.

GOVERNANCE DEVELOPMENT

Algorithmic processes and the principles driving responsible analytics governance is an extremely complex dance, one that is never perfect. In fact, "perfect governance" should not be the point, as advancements in technology and society's relationship to technological progress is ever-evolving. Rather, your goal should be creating a framework that supports business relevance and the responsible development and use of algorithmic processes. You are truly *enabling* all parties involved – from business owner to data scientist to vendor to end-user – to realize technological innovations and digital ethics principles. Core activities include:

- Continuous alignment with the overall strategic direction of the enterprise and digital ethics policies and principles.
- Continuous development of the responsible analytics governance process.
- Continuous development of responsible analytics governance technologies (platform and tools).
- Continuous development of assessments and increasing relevance over time.



RESPONSIBLE ANALYTICS GOVERNANCE INTERDEPENDENCIES

For responsible analytics governance to work it must work with key functions across the enterprise. These interdependencies are critical, as successful governance of any kind is not a siloed practice. There must be a conscious and deliberate effort to engage, collaborate, and build on each other's strengths. To fully develop your responsible analytics governance program, consider interdependencies with data and analytics, IT, supply chain systems, procurement, and risk management and compliance functions, to name a few areas that span a global corporate ecosystem. Connecting the dots within and between your company's ecosystem (internal and external) is essential to integrating responsible analytics governance into every step of the algorithmic lifecycle.

Final Note on the Way Forward

The goal of this guide is to capture some of the main thoughts and ideas on the topic of responsible analytics governance in a global enterprise context. The field is evolving and will develop as rapidly as AI technology. What we know for certain is that enterprises will benefit from resilient and agile processes, a dynamic approach to the changes in technology and business context, while always striving for efficiency. Responsible analytics governance has a role model – AI technology itself – as we aspire to innovate governance and create value alongside the advancements of artificial intelligence in business. The most exciting part of it all is this is just the beginning – the future is bright on the foundation of responsible governance for artificial intelligence.



GIOVANNI LEONI

Giovanni Leoni is the Global Head of Algorithm and AI Ethics at Inter IKEA Group. Leoni is responsible for developing, implementing, and leading the operationalization of responsible analytics across the Inter IKEA Group of companies. This concerns all algorithmic processes, including administrative processes, operational processes, analytics and AI. He is an active contributor to the global AI ethics community, engaging in legislative industry consultations and various forums and bodies, among them being on the Board of Directors for Ethical Artificial Intelligence Governance Group. For more than 20 years he has been active across multiple sectors, private and public, driving development of user experience, sustainable change management, and clarity in impacting factors for better business and wrapping everything in numbers and decision-models. He has worked in diverse areas, including public sector processes, sales, procurement of services, digital and physical products, supply chain, automatization and business analytics. In his current assignment it all comes together, being a partner to create better business with the use of data through values-led decision-making.

IIANALYTICS.COM

Copyright © 2023 International Institute for Analytics. Proprietary to subscribers. IIA research is intended for IIA members only and should not be distributed without permission from IIA. All inquiries should be directed to membership@iianalytics.com.