


onix | Google Cloud

# Data/Model Governance in the Age of AI and Machine Learning

# Table of Contents


- The Cost of Biased Data and Flawed Models 04
- Instituting Responsible Data Governance 05
- How Policy Developments Can Incentivize Accountability 06
- Implementing robust Data/Model Governance with Onix 08

# Data/Model Governance in the Age of AI and Machine Learning

Machine learning (ML) and Artificial Intelligence (AI) models are now transforming decision-making across industries, governments, and academic spheres. As big data and computing power increases, AI applications are growing more ubiquitous by the day. Facial recognition programs, predictive healthcare analytics, autonomous vehicle navigation systems, and emotion-detecting algorithms (for personalized educational curricula) represent some of the applications enabled by AI and intelligent automation.

There are self-evident advantages to leveraging the capabilities and efficiencies of AI and ML-powered algorithms, thus enabling solutions outside the limitations of human resources. However, high-stakes decisions guided by artificial neural networks and other modeling techniques also demand heightened diligence around development processes and supporting data infrastructure to ensure fairness, accountability, and transparency.

> History shows that where ethics and economics come in conflict, victory is always with economics. Vested interests have never been known to have willingly divested themselves unless there was sufficient force to compel them.

The rapid adoption rate of responsible and ethical AI represents not an economic divestment but an embracement of governance considerations. Policy and public awareness are compelling a shift primarily due to societal privacy interests as the starting point in ethical considerations over economic interests. However, there is little doubt that there is substantial future-proofing and economic value to the emergence of data governance in AI and machine learning systems.

The net benefits are privacy protections, data, model management, insight, IP protection, and efficiency. The economics, while expensive initially, ultimately reduce costs substantially.

It is equally challenging for companies to quantify the hidden costs. For a start, they can remove the unknowns (as much as possible), and create a more straightforward and quantifiable measure of estimation and authority.



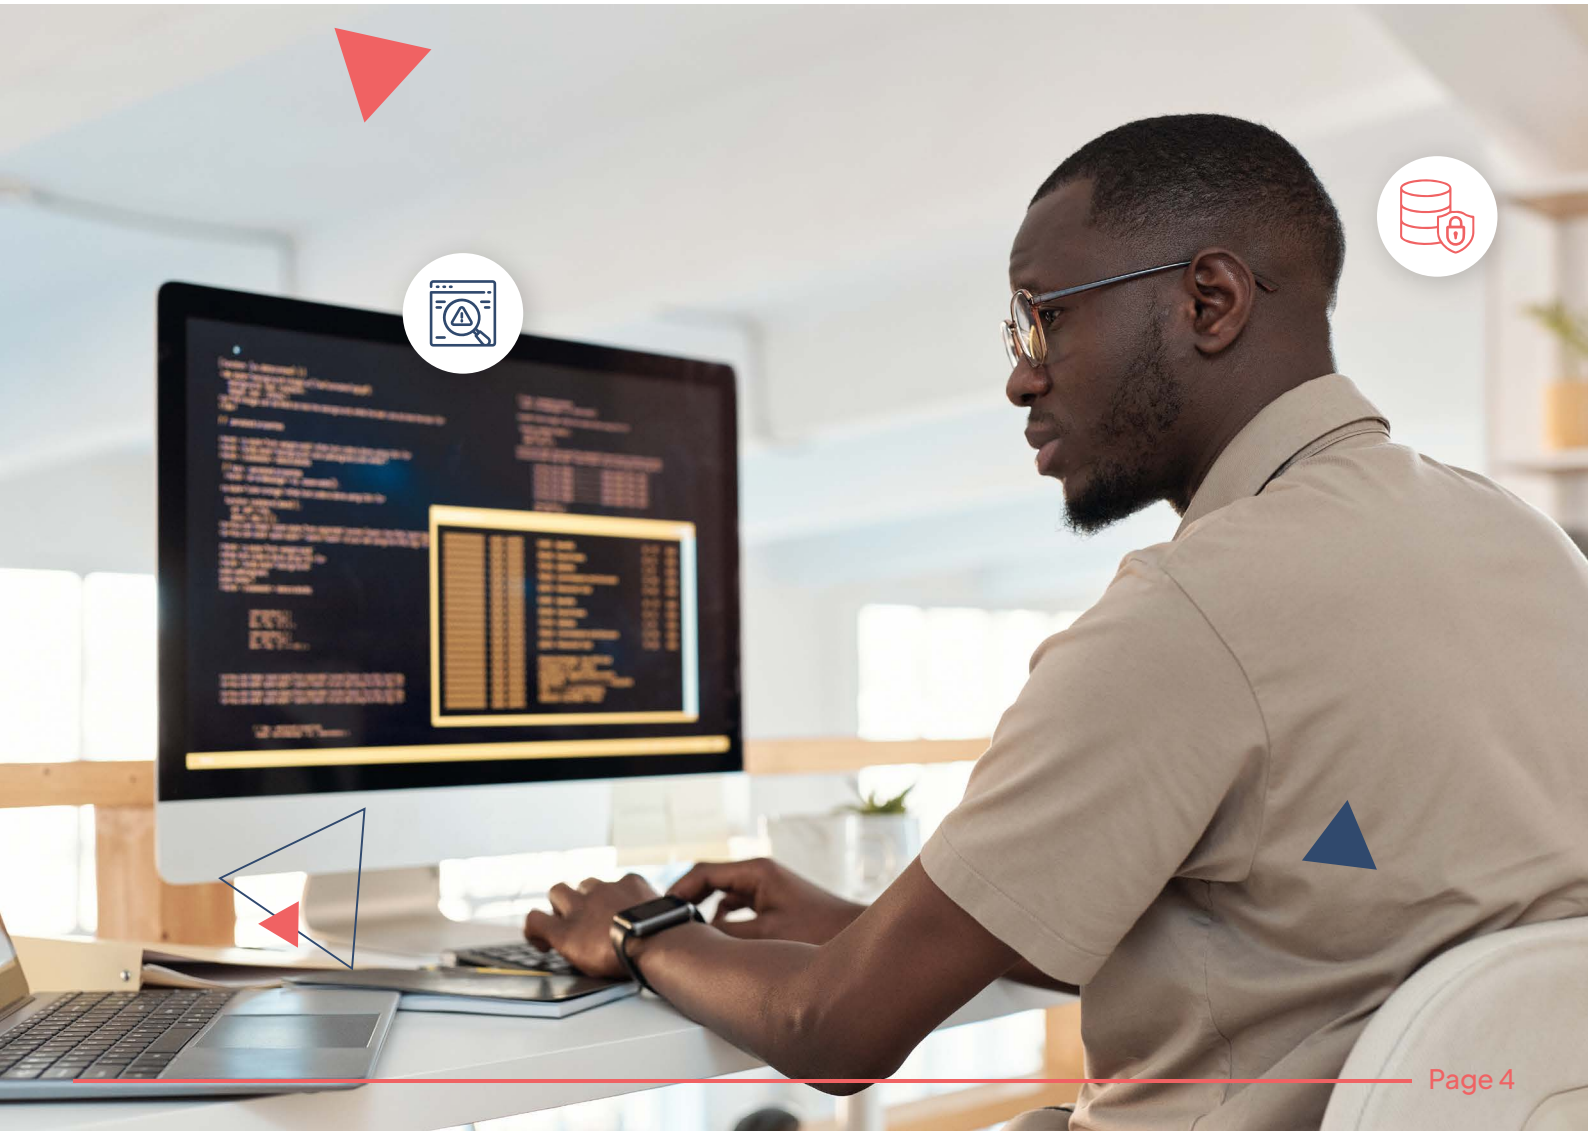
## The Cost of Biased Data and Flawed Models

Recent research shows that inattentive data collection, incomplete categorization considerations, and lack of pipeline oversight propagate hidden biases within sophisticated machine learning architectures. For instance, financial lending risk models have approved straight, white males for credit lines at higher rates than equally qualified minorities.

Similarly, in the healthcare domain ML-powered algorithms designed to spot patients requiring elevated care have, in practice, discriminated against patients based on their race. AI-powered hiring tools designed to standardize and reduce the evaluator's partiality have systematically down ranked female (or older) applicants.

### Other problematic and unethical outcomes continue to persist. This includes:

- The invasion of data privacy through covert data harvesting.
- Opaque governmental surveillance dependent on biometric tracking absent oversight.
- Disproportionate policing of marginalized city zones guided by place-based forecasting.
- The risks of cascading harms to civil rights, equity, access, and justice.



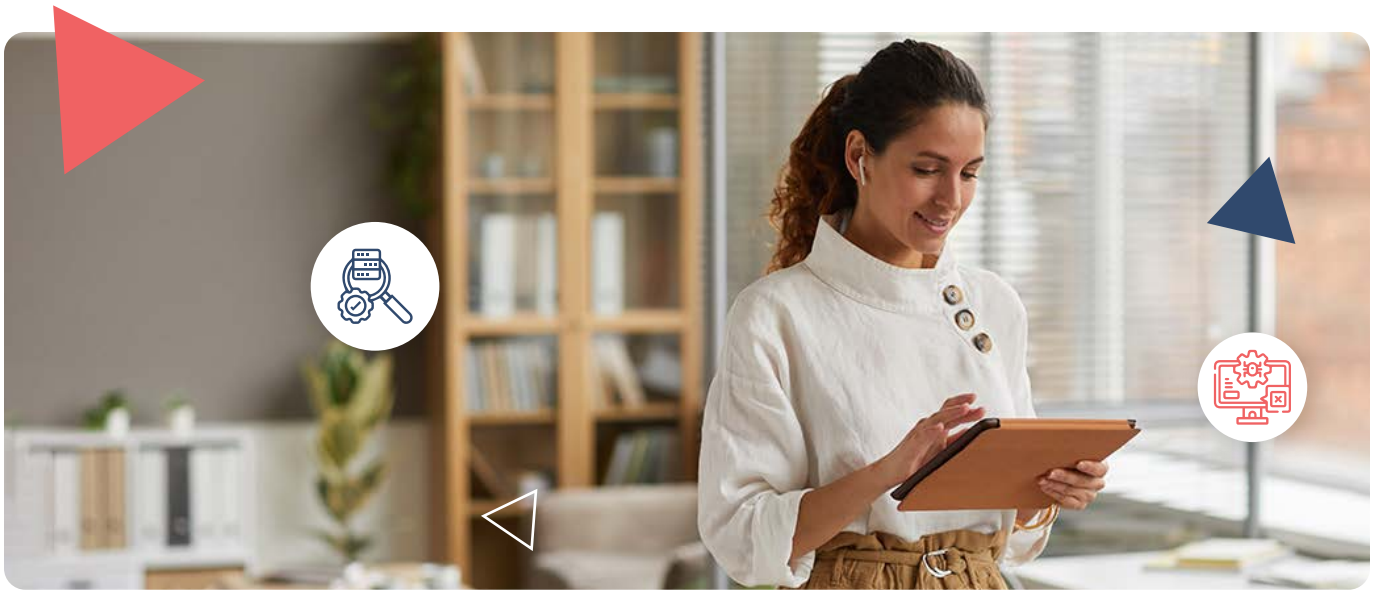
## Instituting Responsible Data Governance

Organizations need to construct equitable and ethically conscious AI models. Comprehensive data governance provides oversight across the integrated pipeline, from initial data collection through encoding, labeling, model development, testing, operationalization, and beyond.

### Key tenets of responsible data governance include:

- Metadata frameworks distinctly log sourcing, ownership permissions, and dataset migrations in detail.
- Statistical testing and auditing procedures to catch skewed samples, incomplete categorization schemas, and disproportionate class representations.
- Validation reports analyzing performance consistency across user subgroups, plus ongoing monitoring for emergent biases.
- Model risk assessment protocols and minimum accuracy/fairness thresholds by application type.
- Explainability features to interpret model behaviors and enable redress.
- Data ethics review boards provide an independent assessment of appropriate usage.
- Access controls, stringent data handling procedures, and security aligned with sensitivity.
- Global (scope) audit trails monitoring and authorizing all data and model access.

Data Governance further entails reconsidering how we conceive data within machine learning contexts, not as raw immutable inputs but as active framing forces imbued with choice around categorization, labeling, and meaning. Such ontological reframing opens space for more constructivist paradigms, more attention to the contingencies and relational nature of data collection, and shifting away from technical remedies toward holistic integration of domain expertise and stakeholder voices that can illuminate unseen pitfalls.t



## How Policy Developments Can Incentivize Accountability

Beyond organizational initiatives, policymakers can guide the burgeoning field of AI and ML algorithmic innovation towards equitable aims, constructive applications, and democratically aligned outcomes. Before productizing models, impact assessment requirements offer one avenue: develop a standardized transparency template reporting around metrics like subgroup performance, data composition, and logging model interventions over time. Some municipalities now necessitate fairness audits of automated decision systems operated by vendors. Establishing centralized bodies providing oversight while enabling data sharing under agreed parameters can potentiate innovation rooted in justice.

Through partnerships across sectors and spheres of expertise, dialogic development of governance structures, anticipatory policy, and communal engagement around the sociotechnical forces shaping lives inch towards integrating ethics into research and design. The potentials await activation.

### Here are some critical fundamentals for robust Data Governance in the age of Machine Learning:

1. **Central data ethics oversight committee:**
   This represents a cross-disciplinary group created to form data governance standards spanning policymakers, ML practitioners, lawyers, and ethicists.
2. **Metadata frameworks and data lineage:**
   These are used to rigorously document all datasets feeding into models, including provenance, content descriptors, access permissions, and modifications.
3. **Algorithmic impact assessments:**
   Review model that influences risks across populations before deployment using statistical and ethical analysis.
4. **Bias detection and mitigation mechanisms:**
   Through audit datasets and by applying human-in-the-loop testing, organizations can uncover and address prejudices by gender, race, and age bracket and mitigate issues.
5. **Model risk severity tiers and monitoring:**
   Organizations can define model accuracy/performance redlines based on application areas - and create systems to enable adjustments for data drift, gaming, and other issues.
6. **Explainability and interpretability requirements:**
   Organizations can enable model understanding and interrogate behaviors by importing key model parameters and descriptions of primary correlations/relationships.
7. **External review processes:**
   These processes perform an independent auditing of core model fairness, benchmarking performance relative to standards by expert panels, and checking alignment with legal duties.
8. **Access controls, permissions, and data ethics training:**
   Access control authorization procedures based on user roles ensure ML practitioners possess grounding in responsible data science.

9. **Issue redress mechanisms:**
   This involves creating transparent channels for affected individuals/groups to submit model-related problems leading to adverse impacts along with processes for investigation and remediation.

10. **Responsible innovation policy partnerships:**
    Government authorities lay the groundwork of regulations and collaboratives with industry/academic leaders focused on furthering innovation rooted in the public good.

*** LLM tools often face data transparency and explainability limitations (LLMOps). Open-source models may only sometimes provide precise citations for the data they use, making it challenging to understand the source and nature of the information they generate. Furthermore, these tools frequently need more explainability, making understanding how they arrived at a particular output or decision challenging.

*** Given these challenges, it is essential to use services that offer indemnification to the customer. Indemnification means that if any issues or liabilities arise from using these tools, such as incorrect or harmful advice or data breaches, the service provider will take responsibility. This protection is crucial in managing the risks associated with adopting LLM-based technologies, particularly in professional or critical applications where accuracy and reliability are critical.



# Implementing Robust Data/Model Governance with Onix



## Cloud Data Loss Prevention (DLP)

This service helps identify and mask sensitive data in your Google Cloud (GCP) environment. It is crucial to ensure that personally identifiable information (PII), financial data, and other sensitive information are adequately protected and managed, essential for compliance with data privacy regulations.



## BigQuery

As Google's enterprise data warehouse, BigQuery allows for the secure and compliant storage and analysis of large datasets. It features fine-grained access controls and comprehensive audit logs, aiding in data governance by ensuring only authorized users can access sensitive data.



## Cloud Identity and Access Management (IAM)

This service enables you to manage access control by defining who (identity) has what access (role) to which resource. It is a fundamental tool for enforcing data governance policies related to user access and permissions.



## Data Catalog

This fully managed and scalable metadata management service helps you discover, understand, and manage your data assets in Google Cloud. It is instrumental in maintaining an organized and easily accessible data environment.



## Cloud Audit Logs

These logs provide a record of actions taken by Google Cloud services and users, offering invaluable insights for monitoring and ensuring compliance with governance policies.



## Cloud Security Command Center (Cloud SCC)

This comprehensive security management and data risk platform for GCP helps you prevent, detect, and respond to threats. It is essential for maintaining the integrity and security of your data in the cloud.



### Cloud Compliance Reports Manager

This tool helps you access, manage, and understand your compliance posture by providing easy access to audit reports, certifications, and other documents that attest to GCP's compliance with various standards and regulations.



### Policy Intelligence

This set of tools, including IAM Recommender and Policy Simulator, helps you understand and improve your cloud resource policies. It helps optimize policy settings in a way that balances security and accessibility.



### Vertex AI

For machine learning specifically, Vertex AI provides a managed service for preparing, building, and deploying machine learning models. It includes features for ensuring the data used for training models is handled securely and competently.



### Cloud Healthcare API

For projects involving health data, this tool ensures that your applications follow healthcare regulations (like HIPAA), offering a bridge between healthcare systems and applications hosted on Google Cloud.

Let us help fast-track your organization's capabilities while centering people first. Why choose between cutting-edge systems or responsible oversight? Through strategy, foresight, and cross-sector collaboration, enterprises can implement ML that accelerates secure and equitable progress.

## Works Cited:


Dr. Bhim Rao Ambedkar - Quote - An Indian jurist, economist, social reformer and political leader who headed the committee drafting the Constitution of India.

Bartlett. "Consumer Lending Discrimination in the FinTech Era." 2021, https://www.sciencedirect.com/science/article/abs/pii/S0304405X20302184

Gebru. "Various Works."

Hanna. "Various works."

Lum, and Isaac. "To Predict and Serve?" 2016, https://rss.onlinelibrary.wiley.com/doi/full/10.1111/j.1740-9713.2016.00960.x)

Doctoral candidate in the Department of Political Science at Michigan State University

Mehrabi. "A Survey on Bias and Fairness in Machine Learning." 2021, https://arxiv.org/abs/1908.09635

Obermeyer. "Dissecting racial bias in an algorithm used to manage the health of populations." 2019, https://www.science.org/doi/10.1126/science.aax2342

Suresh, and Guttag. "A Framework for Understanding Unintended Consequences of Machine Learning." 2021, https://dl.acm.org/doi/10.1145/3461702.3462527



Get in touch

onixnet.com

connect@onixnet.com