gitops **cicd** **continuous-delivery** **git**

AI Governance

AI governance is about ensuring greater transparency across the AI lifecycle and the model itself. IBM recently announced watsonx.governance, a next generation enterprise toolkit which is designed to automate and accelerate workloads **across the AI lifecycle** while providing **risk management** and facilitating **regulatory compliance.** Use IBM AI Governance services to accelerate responsible, transparent, and explainable AI workflows with an AI governance solution that provides end-to-end monitoring for machine learning. Monitor machine learning assets from request to production. Collect facts about models that are built with IBM tools or third-party providers in a single dashboard to aid in meeting compliance and governance goals. Implement an approval workflow to meet complaince goals.

These Cloud Pak for data services can be used individually or together as part of your governance and MLOps plan.

AI Factsheets allows automated collection and documentation of model metadata at all stages, from model idea to production

Model and process metadata is captured in a central meta store. Having all model facts in central place is important both to increase the productivity of the MLOps process (model facts are immediately visible to all parties involved in the lifecycle of an AI model) and to comply with regulatory requirements. Data scientists benefit from assistance and automation of the documentation process. Transparency of model metadata supports audits and brings more clarity to stakeholder or customer requests. Metadata captured in AI factsheets includes model

details, training information, metrics, input and output schemas, or details about the models used, such as quality metrics, fairness or drift details.

AI Factsheets - Official Documentation AI Governance - Factsheets - Official Documentation

OpenPages: Govern models through the complete AI workflow considering policies and regulations

The next generation governance-toolkit provides a range of capabilities to identify, manage, monitor, and report on risk and compliance. It accelerates the creation of models at scale, from use case idea (model candidates) to production deployment, by incorporating approvals in the workflow-based approach. Full transparency of any type of model (e.g., task specific data science artefacts or foundation models) is ensured and made visible in customisable risk monitoring dashboards. Additionally in Open Pages corporate policies and regulations can be assigned to models, e.g., annual bias review (required for EU AI ACT) to ensure that models are fair, transparent, and compliant [4].

OpenPages - Official Documentation

OpenScale allows to monitor, explain, and benchmark your model

Model monitoring is an ongoing task to track models and to drive transparency. This includes the monitoring of the general model performance (e.g., accuracy) and more specifically monitoring of fairness or model and data consistency over time (i.e. drift). Open Pages supports threshold definitions for model performance metrics and combines those with automated detection of threshold violations to trigger model retraining. It implements explainability by supporting explanations how the model arrived at certain predictions. Model benchmarking is supported – it is common practice to compare and

See OpenScale in this Handbook.

For more information check out the official documentation or the example Notebooks.

Last update: April 8, 2024

Authors: Dominik Kreuzberger, Julianne Forgo