LTIMindtree

POV

Responsible AI: An Integral part of ModelOps

Prepared by:

Deep Sharma

Sr. Specialist, AI & Data Engineering,
Data & Analytics practice, LTIMindtree

Abstract Line:

Responsible Artificial Intelligence (AI) has become integral to AI/Machine Learning (ML). As organizations experiment with new-age use cases around them, the need is to have a holistic framework to standardize, streamline, and de-risk the journey of piloting to operationalizing the models and adherence to Responsible AI parameters.

Table of Contents

Background

The advent of digital transformation has resulted in troves of data being generated and stored with ease leveraging the power of cloud data platforms. The types of data being handled across the organization have evolved from simply being a system of records and transactions to systems of interaction and observation, resulting in multiple types of data being available within organizational data systems.

As organizations move from managing the data to monetizing the large data troves by leveraging exponential technologies, such as Artificial Intelligence (AI) and Machine Learning (ML), empowering the business teams with the insights generated to enrich the overall customer experience becomes the primary focus. Using innovative products can help businesses approve loans in minutes, diagnose cancer, select candidate profiles, approve claims, and implement many more such use cases. This will further help businesses generate new revenue streams, drive cost efficiency, and deliver new-age products.

With AI/ML being embedded across products and processes by organizations irrespective of the revenue size and domain, managing the operations side of the model is critical to this value chain. This means that a major chunk of key business decisions will have AI fabric, resulting in millions of dollar amounts being dependent on it. Moreover, most AI/ML models currently in the pilot phase will have to move into full-scale operation. ***A recent Gartner report predicts that more than 75% of organizations will shift from piloting AI technologies to operationalizing them by the end of 2024.***

Democratization of AI/ML with Operations = ModelOps

The scope of AI/ML products/applications is expanding every day, resulting in multiple technology-embedded touch points in our daily lives. This has expanded the scope of AI/ML products from big technology players to Fortune 500 and Global 1000 companies with 100s & 1000s of models helping us deliver simplified experiences. This has made AI applications and products a necessity.

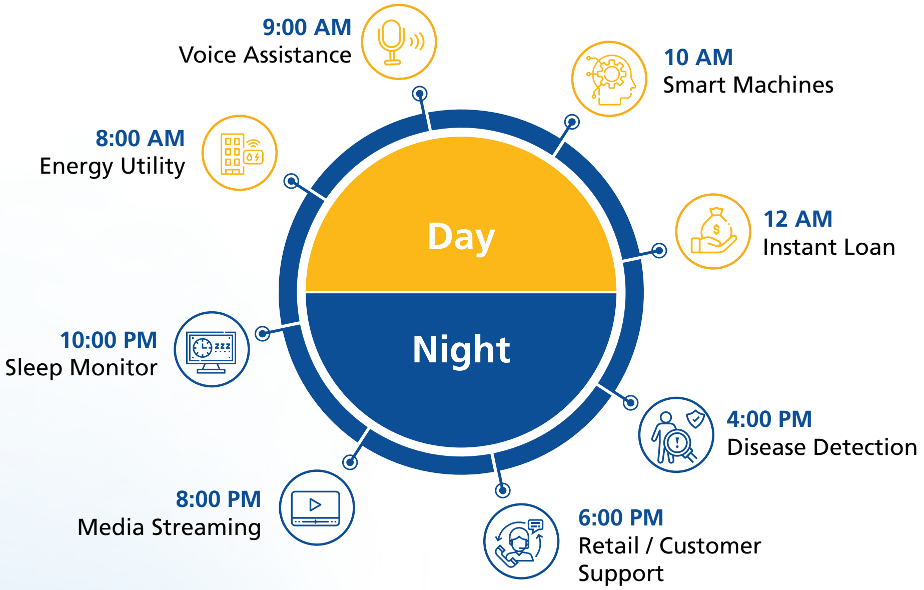

Figure 1 Day in a life Scenario with AI

The big question which comes is, what if an AI/ML model dysfunctions, and what will be the ripple effect of it in terms of revenue, business, and reputational loss along with unsatisfied customer experience? An organization that has billions or millions of dollars in business revenue riding on AI systems has to ensure the models are bias-free, the output can be explained easily, and models, when deployed into production, can raise a flag in case something goes wrong. In short, how the operational side is managed becomes a critical aspect in ensuring the success of models.

To help amplify the adoption of AI/ML across the business, certain key aspects need to be dealt with critically. First, how do we ensure that the predictions made by AI/ML systems don't pose any business or reputational risk? Second, how can we ensure a frictionless operationalization of AI/ML models at scale?

Rise of Responsible AI with ModelOps

As technology-related touch points rise in our day-to-day lives, the decisions made by AI/ML products/systems can have a massive impact. There have been multiple instances in public life where AI systems have directly interfered with human rights, resulting in providing partial model insights having a bias toward a particular class.

Across the spectrum, many well-documented and publicized instances of AI systems have left many companies red-faced due to alleged discrimination in the insights drawn. These instances range from discrimination based on class (gender, age, sex, income group, ethnicity), invasion of privacy, spreading distrust in society, or cases covering regulatory, financial, and legal threats.

As AI systems expand their horizons and scale, if something goes wrong, it has a ripple-down effect across all people involved in its value chain. Hence, the need is to ensure a holistic framework for Model Operations (Model Ops), culminating in responsible AI + ML operation integrated across all the AI systems. The key elements of the frameworks should ensure that the right models are pushed into production adhering to all the key rules and have a mechanism to ensure continuous integration, continuous deployment, continuous training/re-training, and continuous monitoring of the models.

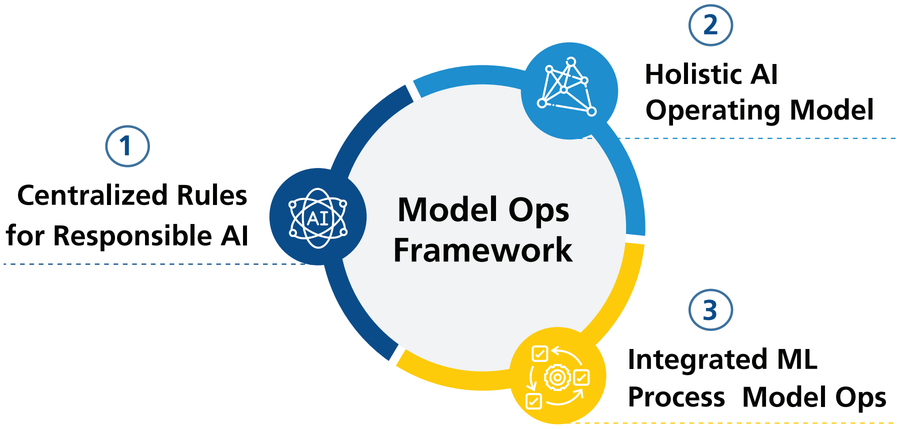
1
Centralized Rules for Responsible AI
AI
Model Ops Framework
2
Holistic AI Operating Model
3
Integrated ML Process Model Ops

Figure 2 Key Elements of ModelOps Framework

Model Ops Framework for Integrated Responsible AI

1. Define Responsible AI Rules

The key to helping businesses prepare and safeguard their AI initiatives is to have a centralized pillar of responsible AI. The key elements of the pillar will help create the guiding principles for organizations to completely manage and govern models across their AI projects which are fair, ethical, and responsible for the insights it derives.

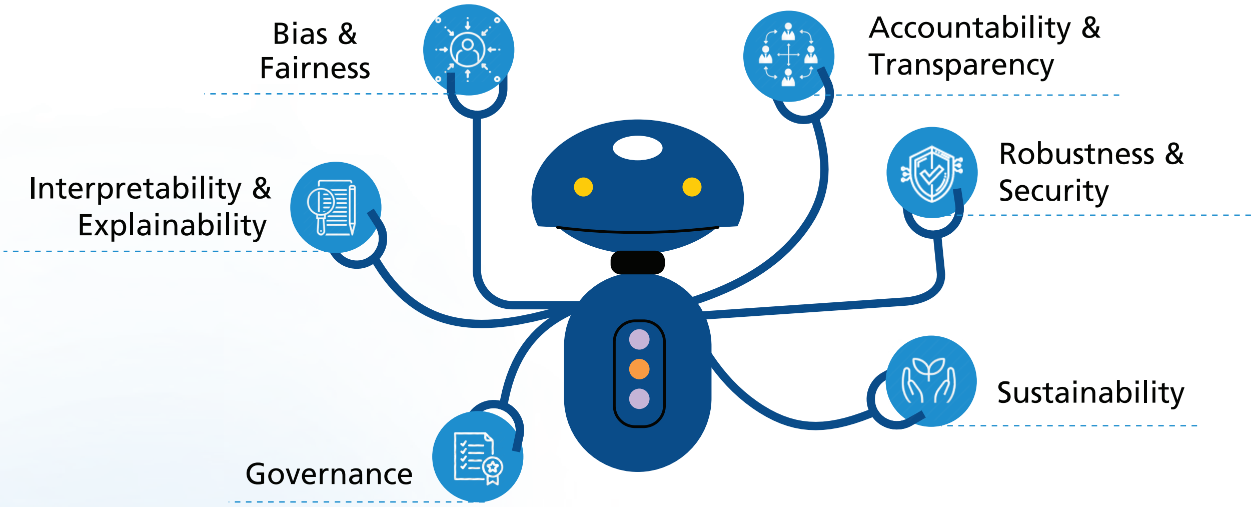
Bias & Fairness
Accountability & Transparency
Interpretability & Explainability
Robustness & Security
Sustainability
Governance

Figure 3 Key Pillars of Responsible AI Policy

Governance

Determine standard governance policy covering key considerations across data, model, and regulatory compliance to ensure end-to-end governance of all the models in production. Have a holistic view across the data and ML lifecycle to ensure appropriate adherence to governance parameters such as:

- Data governance: Pre-defined Data Quality rules, data stewards the custodian of data, registry of the data, establish golden data record to enable ground truth evaluation.
- Model governance: Model registry and versioning of the models, model risk classification mechanism based on usage and criticality, model metadata and lineage to ensure audibility, pre-defined ML process, and steps to ensure adherence.

- Security: Ensure role-based access control with hierarchy-wise controls, enabling cloud platform-based security measures, and update key user groups list at a set time interval.
- Regulatory compliance: Adherence to key data privacy rules like the General Data Protection Regulation (GDPR), Central Consumer Protection Authority (CCPA), etc., to ensure the safeguarding of Publicly Identifiable Information (PII) data along with adherence to key financial, healthcare, and other AI-related compliances and rules

Interpretability and Explainability

Have a mechanism to help the teams understand the variables inside the model and the output we receive in the form of insights. The key is to ensure that the process of what goes between the input and output can easily be deciphered and there remains no black box for models consumed by various teams. This will help both teams, the one that builds the model and the one that consumes or operates the model, to understand and interpret the models. Key elements for this include:

- Feature importance: Hierarchical ranking of key features which played an essential role in determining the output of the models.
- What-if analysis: Provide the ability to determine the impact of the model output if there is any change in the given parameters or features. This helps the team to plan scenarios and easily test the models.
- Model metadata: Enable storing of the underneath layer (key metadata) of the model covering all the critical information of the serialized model.
- Also, the need is to have a central hub to store all the key model-related key artifacts like exploratory data analysis, experimentation results, model validation results, and QA reports in a centralized place.

Bias and Fairness

The output of the models or the way a model behaves is critical for its widespread acceptance in the consumer community. But to ensure a larger acceptance of AI systems, the need is to ensure that the model provides fair insights and doesn’t possess any biases towards a particular class (gender, age, income group, religion, ethnicity, etc.). To ensure the business from reputational risks, it is essential to ensure models are fair and free from bias.

- Pre-defined checks: Ability to perform pre-defined checks on the data used to mitigate biases towards a particular class, education, income, etc.
- Ground truth check: Ensure fairness in the model predictions by having a ground truth mechanism.
- Quality assurance: Pre-defined test cases to ensure quality assurance of the model’s performance.

Accountability and Transparency

AI systems provide accurate predictions using the power of the algorithm on which it is built, which are then used as judgments. To ensure that these judgments are fair and easy to decipher, a trail is needed to cover all the key artifacts and knowledge documents. As AI systems get more complex with multiple touchpoints and teams across the organization using the solution, the risks pertaining to their judgment will increase. The need is to have an efficient mechanism covering the following:

- Clear roles and responsibilities: Crystal clear definition of roles and responsibilities within the organizational structure for the people working on the models, along with having a designated model/ product owner.
- Documentation: Complete model documentation containing all the artifacts and documents of the model in a centralized knowledge hub that is accessible across the organization.
- Knowledge sharing: Clear communication by the model/product team highlighting any shortcomings or limitations of the models and clear documentation of how data is being used to ensure transparency.

Robustness and Security

Develop AI systems that provide robust performance and deliver the right predictions to help the business team achieve the desired objectives, and provide safe systems in terms of utility by diminishing the negative impact.

- Safe deployment: Having safe and sound infrastructure systems to ensure frictionless production of the model, along with ensuring contentious monitoring.
- Malware safeguards: The right security measure to ensure that the deployed infrastructure, whether in a local or cloud environment, is safeguarded from malware or cyber-attacks.
- Security framework: Robust security framework on the infrastructure network and a defined policy with a role-based access control mechanism.

Sustainability

Ensure that models consume natural resources efficiently and smartly by optimizing the overall processing power from multiple server infrastructures which support the models. The need is to have climate-efficient models sensitive to environmental changes by efficiently consuming power. This can be achieved by:

- Green credit models: The models should be given green credits based on energy consumed during computation.

- Sustainable ways of working: Support green ways of working by reducing the compute-heavy models and enabling efficient models that minimize the data-center load.
- Efficient usage: Ensure that while building the models, compute resources are optimized and not used unrealistically while training the models.

2. Holistic AI Operating Model

To ensure the success of a particular model, many teams work cohesively to build, manage, and operationalize a model. It requires dedicated efforts from key roles like data science, ML engineering, and data engineering along with key functional teams like IT, DevOps, and business to streamline an AI/ML model's journey from a pilot project to a full-scale production model.

The challenge comes when all the work gets assigned to a single team with no defined roles and responsibilities to ensure the end-to-end inception of the model, from piloting to operationalizing the models. This leaves us with critical gaps like:

- How do we engage with multiple stakeholders if we need to work towards building new AI/ML use cases?
- Which team is responsible for scaling the models into production?
- In case a model misbehaves, which team is responsible for managing it?
- How do we keep a check and ensure adherence to Responsible AI parameters?
- Who will act as change agents to drive the change management process across the organization once models are deployed in production?

This results in a lack of control over ownership and clarity of roles required across the model lifecycle process. This hampers the time to market from pilot to production as the team grapples with daily issues and reduces the model performance, as managing the operations around the models becomes difficult.

As organizations grow big and embed multiple models across their product systems, the chasm across the stages from piloting the models to operationalizing them at scale increases. A holistic operating model for AI/ML functionality will provide a well-defined way of working, emphasizing how multiple teams will engage and having the right mix of skills and roles.

Having a centralized operating model for AI/ML teams will further ensure a new way of working with an AI/ML process with pre-defined roles and responsibilities to ensure holistic ModelOps as the centerpiece coupled with integrated, responsible AI adherence. However, there are multiple types of operating models.

Implementing a hub-and-spoke operating model that is supported by key business influencers and enablers will ensure the effectiveness of the AI/ML initiatives.

The hub-and-spoke model works as a Product-Oriented Delivery (POD) based structure with a centralized hub that acts as a gatekeeper for all the AI/ML initiatives. The operating models' primary purpose is to help establish a functional relationship between the key roles that work towards piloting the key AI/ML use cases. This functional relationship simplifies scaling the AI/ML models as the pre-defined charter enables the key roles to operationalize the models at scale.

The "key influencer" in the operating model plays the role of domain expert or functional team which helps to define and structure the business problems, along with framing the critical parameters for the AI/ML use cases. The team provides functional expertise to enable the models and ensure their success in real-world scenarios. It consists of multiple business enabler teams covering both the functional and operational aspects, focusing on redefining the business problems and working towards improving the revenue or optimizing the cost.

The "key enablers" in the operating model provide the support functions that act as invigorators in the model lifecycle process of piloting the use cases and operationalizing the use cases. It helps facilitate the necessary infrastructure in terms of people, processes, and technology enablement and ensures adherence to the AI regulations across the model lifecycle process. It covers key processes, principles, and templates for AI/ML, which helps standardize the journey of AI/ML models from pilots to end-to-end operationalization. Another critical part of the operating model is change management; this helps streamline the model adoption across business teams, focusing on improving the adoption. A well-defined change management process is essential in democratizing the AI/ML models across the organization.

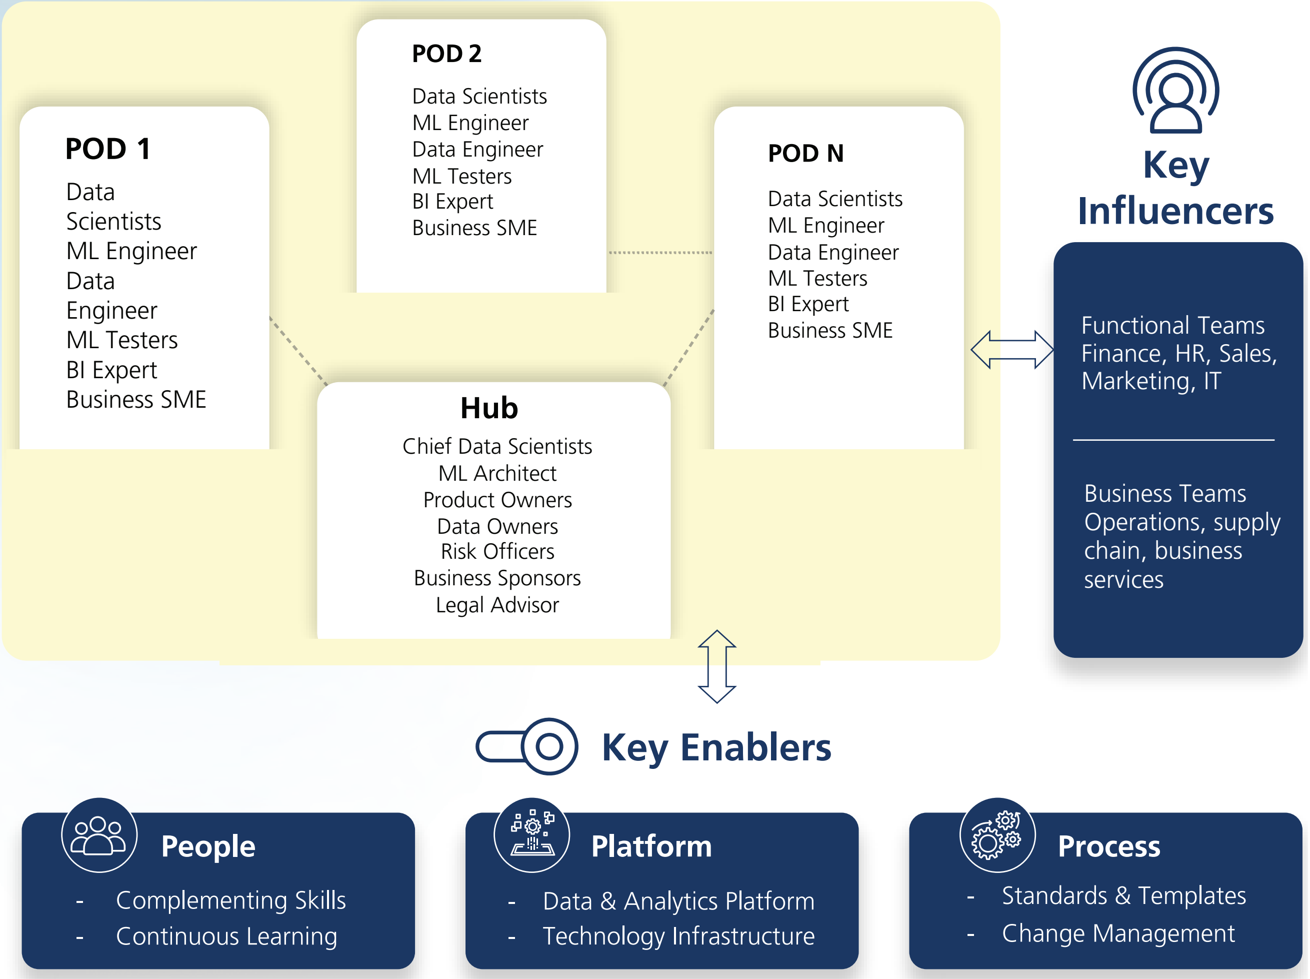

POD 1

Data Scientists
ML Engineer
Data Engineer
ML Testers
BI Expert
Business SME

POD 2

Data Scientists
ML Engineer
Data Engineer
ML Testers
BI Expert
Business SME

POD N

Data Scientists
ML Engineer
Data Engineer
ML Testers
BI Expert
Business SME

Hub

Chief Data Scientists
ML Architect
Product Owners
Data Owners
Risk Officers
Business Sponsors
Legal Advisor

Key Influencers

Functional Teams
Finance, HR, Sales, Marketing, IT

Business Teams
Operations, supply chain, business services

Key Enablers

People

- Complementing Skills
- Continuous Learning

Platform

- Data & Analytics Platform
- Technology Infrastructure

Process

- Standards & Templates
- Change Management

Figure 4 Target Operating Model

Some of the key roles and responsibilities covered as part of the target operating model include:

- **Chief Data Scientists:** Key owner of all AI/ML models, along with leading the Responsible AI council to ensure adherence to AI principles and model governance across models.
- **AI/ML Architects:** Outline and build the overall technology and process roadmap to power the new-age use cases and manage and operationalize the existing ML models.
- **Model Risk Officer (MRO):** Responsible for validating, approving, and managing model risk. The risk officer will ensure the models adhere to policies and regulations. In small teams, Chief Data Scientist acts as MRO.

- **Data Scientist:** Collaborate and lead the data science POD of building AI/ML models leveraging statistical and mathematical knowledge. They work towards analyzing, processing, and modeling data to get actionable insights for business teams.
- **AI/ML Engineers:** Ensure seamless operationalization of AI/ML models and manage the model drift and monitoring.
- **Business SME:** Provide knowledge and expertise with functional experience of working on key business domains. This will help shape the ML models with a business-specific purpose.
- **ML Testers:** Ensure quality assurance to the models which are being built and validate the results.
- **Legal Advisor/Ethics Expert:** Responsible for working closely with the product owners and data scientist to crystalize the model definition and the process, and ensuring adherence to AI and other regulatory compliance.
- **Product Owners:** Work closely towards managing the various PODs and help bridge the gap between business and technical teams. Help streamline and prioritize the execution of the products and ensure timely deliverables.

The outcome being the target operating model helps deliver a holistic ModelOps with integrated responsible AI by enabling:

- Integrated operations between development and operations teams to help scale end-to-end model lifecycle management.
- Effective communication and change management for stakeholders to ensure the seamless adoption of AI solutions.
- Frictionless collaboration between multiple teams with pre-defined roles and responsibilities.
- Ensure a coordinated approach for Responsible AI by making ethics an integral part of AI implementation.
- Enable an agile way of working with a pragmatic governance and expectation management approach.
- Ensure the right skill mix to support the AI/ML projects.
- Improve the speed with which solutions are developed and deployed, reducing time to market.

3. Integrated ML + Responsible AI Process

Building AI/ML models is creative process and requires constant iterations and refinement. In a typical process, the data scientists prepare the data, create the feature, train the model, tune the parameters, and perform validation and QA. Once the model is prepared, the data scientist hands it to the IT team to operationalize it. The challenge comes when there are multiple models and scenarios to scale and manage the models across the organization. If a model misbehaves without having a pre-defined process or threshold parameters to catch the same, it would hamper the business revenue and damage the reputation.

The complexity of piloting the AI/ML use case to operationalize them at scale requires standardization and pre-defined processes to streamline model development, implementation, refinement, and governance and ensure responsible AI parameters.

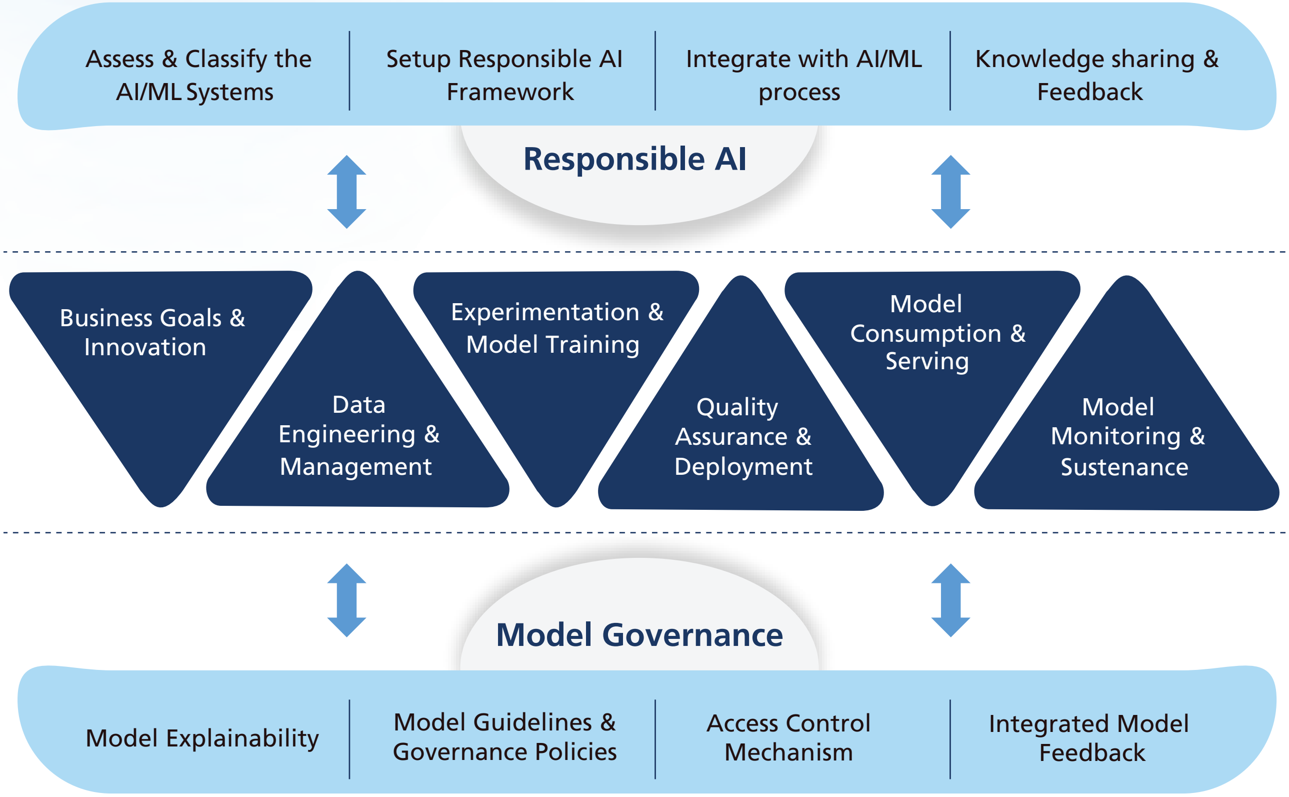
Assess & Classify the AI/ML Systems
Setup Responsible AI Framework
Integrate with AI/ML process
Knowledge sharing & Feedback
Responsible AI
Business Goals & Innovation
Data Engineering & Management
Experimentation & Model Training
Quality Assurance & Deployment
Model Consumption & Serving
Model Monitoring & Sustenance
Model Governance
Model Explainability
Model Guidelines & Governance Policies
Access Control Mechanism
Integrated Model Feedback

Figure 5 End-to-end ML Process

The focus of the end-to-end process is to have a defined ML process with integrated ModelOps, this will help to ensure complete visibility and standardization, right from defining the business goals to experimenting and training to deploying and monitoring the models in a frictionless manner. The detailed aspects of each of the elements in the step cover:

- Business goals: It covers the initial business problem along with success criteria and defined KPIs for the same. This helps to measure the model output with the initial input parameters and showcase the true value generated with the help of the models.
- Data engineering: The key elements of data engineering covers the acquisition of the data, preparing the data for analysis, and performing initial data quality-related checks. A feature store is one key element that helps standardize the data engineering process. This helps to keep all the data sets and feature elements in a centralized store
- Experimentation and model training: The focus is to experiment with the model idea using the data sets, doing the initial explorative data analysis, and building a base experiment. This step also covers infrastructure provisioning, code and model testing, managing the code repository, data and model registry, and training and validating the model.
- Quality assurance and model deployment: This process step focuses on testing the model and its performance in a stage gate-type process. Once the final model adheres to the defined parameters and tests, the model is serialized and pushed into deployment using pre-defined templates.
- Model consumption and serving: Once the model is deployed, the focus comes on how the model is available for consumption across the organization, along with ensuring online or batch scoring and API management.
- Model monitoring and sustenance: Post the model deployment, the key is to continuously monitor the output and performance of the model continuously to ensure robust performance. The drift management mechanism helps manage the model and data and provides the ground truth on model performance. The process covers an integrated feedback loop mechanism and re-training rules in case the model drifts.
- Responsible AI: It helps to ensure that the AI/ML models being built are sound from an ethics and compliance standpoint and meet the global regulatory norms. Also, it focuses on providing an approach towards ensuring responsible AI across AI/ML systems, defining the Responsible AI framework and integrating it with the ML process.
- Model governance: The primary purpose is to oversee the entire ML process and ensure that the AI/ML models adhere to the model guidelines and are easy to explain and decipher. It also ensures an integrated access control mechanism for the stakeholders.

A key ingredient for a successful AI/ML operating model is having a well-defined end-to-end ML process. It helps to standardize the process for AI/ML development and operationalization, along with providing tools/ templates to support the adoption of the process. This helps to determine the model's output by allowing different teams to work independently by focusing on the core skills and competencies. It also creates well-defined roles and responsibilities to help manage unexpected and disruptive changes.

Conclusion

The AI/ML systems have helped us re-imagine the ways of working for organizations across the globe, with businesses leveraging it to solve new-age business problems. But as these systems get democratized in our day-to-day life, government and regulators are focusing on regulating the AI/ML systems. Organizations need to look from a vantage point to effectively and efficiently manage the AI/ML models and their operations across their existing data and analytics platform, focusing on standardization to enable scale. As organizations move towards making AI/ML models sound ethically and compliant with AI regulations, the convergence of Responsible AI with ModelOps remains key to ensuring AI for all across the globe.

Introducing kenAI, our mindful automation accelerator, to standardize and streamline your AI/ML journey across key hyperscaler cloud platforms. KenAI addresses the core challenges of scaling, managing, and governing the models on cloud data platforms. It leverages its pre-built templates and utilities to deliver simplified machine-learning operations leveraging automation, predictability, and Responsible AI.

Reference Links

https://www.microsoft.com/en-us/ai/reswponsible-ai

https://aws.amazon.com/machine-learning/responsible-machine-learning/

https://ai.google/responsibilities/responsible-ai-practices/

https://hbr.org/2020/10/a-practical-guide-to-building-ethical-ai

https://hbr.org/2022/03/how-to-scale-ai-in-your-organization

https://hbr.org/2021/05/getting-ai-to-scale

LTIMindtree

About the author

Deep works as a data and AI consultant and has more than six years of experience working across new-age data and AI initiatives. He has worked with multiple customers across domains to help them scale AI across the organization and adhere to Responsible AI.

Deep Sharma

Sr. Specialist, AI & Data Engineering,
Data & Analytics practice,
LTIMindtree