Model Monitoring
The Path to Reliable AI



ModelOp

Introduction

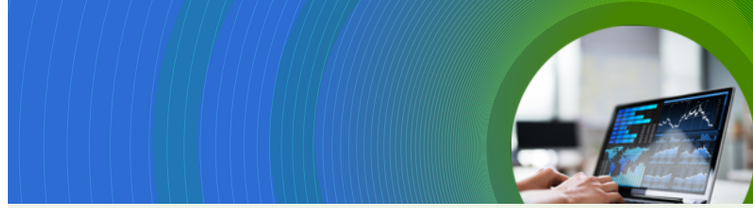
AI initiatives are only as successful as their models. How you treat AI models during their life cycle determines the amount of value you realize.

The model is the central element to AI success. That is inescapable, and is why companies invest so heavily in talent and tooling to create effective, differentiated AI models. With models, there's a sense among some business leaders that If you build it, great new insights and business benefits will come. Not quite. If you build it, you have a built model. You haven't put it into production. You haven't integrated all the data platforms and infrastructure components it needs for input, nor the visualization and other tools it will use for output, nor the quality control checks it needs to run reliably and ensure compliance. The model will not monitor its own performance and optimize output as conditions change.

For models to be effective and valuable they need to be monitored throughout their entire life cycle, from the time they are put into production until they are retired. Monitoring needs to cover the model's **operations**, **quality**, **risk** and **processes** throughout its life cycle.

Using **actionable monitoring** brings automation, including orchestrating issue remediation, to the AI monitoring environment. Actionable monitoring creates value by ensuring models are producing the most optimal results for the business, accelerating problem detection and resolution, managing compliance and risk factors and freeing data scientist teams to do more development and less operations observation and problem detection.

Actionable monitoring makes AI scalable.



Enhance Your AI Investments

With actionable monitoring that observes, detects and remediates, organizations can:

- Make AI operations more scalable
- Improve data scientist productivity by up to 50%
- Increase model contribution by 10% or more by reducing degradation
- Reduce enterprise risk and improve compliance

Why Models Need to be Monitored

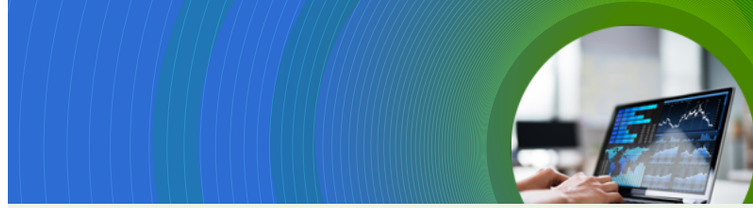
Models need actionable monitoring because AI model performance naturally degrades over time; add to that new risks around interpretability and rising concerns over ethical fairness

Business conditions are not static, so neither is model performance. Without active monitoring, model performance typically decays at a rate of 10% annually, but can be as high as 50% for some models, directly impacting the business value a model returns annually.

Value is harder to measure and more related to use case. For example, consider AI models that are used for fraud detection. Failures or even slight declines in their performance over time can have enormous consequences. A change in accuracy for a model used to assist shoppers with suggested sales or provide other decision support value can swing its value by millions in potential incremental revenue gained or lost. Therefore, these and other models must be rigorously monitored to immediately detect and remediate problems to preserve value and ensure compliance adherence.

As the volume and complexity of models increases, risk grows – risk to outcomes, accuracy, value and especially compliance. Two of the eight largest regulatory fines in 2020 – $400 million assessed against a finance company by the U.S. Office of the Comptroller of the Currency (OCC) and a separate $85 million penalty the OCC levied on a bank – could be attributed to the failure to implement effective risk controls and/or data governance. The growth in AI/ML model development, including more development groups, deployment and data volume make it harder to maintain risk controls and puts organizations at risk of business and regulatory compliance violations.

In addition, models need to be monitored to ensure all their workflows and infrastructure are running as expected. That is why data scientist teams spend much of their time on operations and troubleshooting tasks instead of on development, which limits the number of models that get developed and put into production.



"Deploying models doesn't end with provisioning infrastructure and copying code. Machine learning models are unique in that they must be constantly monitored while in production and regularly retrained, requiring the collaboration of a host of stakeholders from data scientists to ops pros."

Forrester, "Introducing ModelOps To Operationalize AI" August 13, 2020

How Monitoring Produces Value

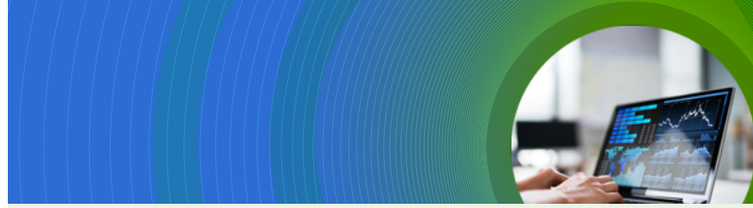
Actionable monitoring creates value in multiple ways, including enhancing model performance, reducing risk, raising data science team productivity and enabling AI expansion.

Actionable monitoring helps organizations get more out of their most valuable AI assets – which includes their data, models, and AI staff. It's well known that data scientists and other data specialist are difficult to find, hire, train and retain, and are paid accordingly. Actionable monitoring helps solve bottlenecks and reduces work efforts so organizations can put and keep more models in production. In our experience, organizations can improve data scientist productivity by up to 50% by supporting them with actionable model monitoring that not only automates monitoring, but problem resolution tasks as well.

Here are some other ways actionable monitoring adds value to AI initiatives:

- **Eliminating model degradation**
  - The 10% annual model degradation is an average and can be mitigated or avoided by continually optimizing inputs and outputs
  - If degradation is eliminated, the model's reliability and business contribution increases
- **Enhancing model performance**
  - Detects problems with infrastructure, thus optimizing model availability
  - Provides a common view of model state and status
- **Reducing business risk**
  - Actionable monitoring prevents unfair bias and unacceptable changes in input data
  - It enforces compliance gates and controls, ensuring business and regulatory requirements are continuously satisfied

In addition, all these automated activities enhance the work efforts of model operators and mean less work for data scientists, capitalizing on the skill sets and time of these roles.



"The sooner institutions get started in building value-based MRM [model risk management] on an enterprise-wide basis, the sooner they will be able to get ahead of the rising costs and get the most value from their models."

McKinsey & Company,
"The evolution of model risk management" 2017

What Needs to be Monitored

It is clear that models should be monitored in a comprehensive, automated way. The question becomes: What specifically needs to be monitored?

Each model has dozens of specific inputs, quality checks and performance thresholds that need to be monitored. To make it manageable, consider that effective monitoring must address four aspects of the model:

- **Operations:** Model is meeting SLA performance and expected decisioning rate
- **Quality:** Decisioning and outcomes are producing optimal results
- **Risk:** Outcomes are ethically fair and unbiased and operating within compliance thresholds
- **Processes:** Operational and governance processes and gates are properly followed

Effective model monitoring includes familiar concepts like uptime and SLA performance, but also requires monitoring many more issues that are specific to AI and ML models, such as input data drift, model concept drift, bias, and much more. Continually monitoring dozens of metrics helps ensure model performance consistency and output quality, and thus reduces risk.

Models are different from applications and other IT assets. They require a monitoring solution that is specific to maintaining a model's operational health and outcomes. Yet, it is important for monitoring activities to be integrated with ITSM solutions, model development tools and software maintenance systems and services to eliminate redundancy and leverage existing investments for remediation activities and enabling auditability.

The following sections provide additional details about what is needed for ModelOps operations, quality, risk and process monitoring.



According to the 2021 State of ModelOps Report:

- Only 32% say their methods are highly effective for maintaining visibility into model state and status
- 42% say lack of efficient processes is a very challenging barrier to operationalizing models

Monitoring Model Operations

Operations monitoring is primarily concerned with the nuts and bolts of AI and ML model execution. It covers the quality and consistency of model inputs, outputs and relevant IT infrastructure.

At a minimum, operations monitoring needs to perform input data validation and provide a dashboard view of each model's status, availability and performance to execution SLAs. Input data validation is essential because data quality is such an important variable to model output results. The relevant performance indicators and reporting metrics are different for AI models than for other software and IT assets. For example, model-specific monitoring and remediation go beyond monitoring thresholds to cover things like:

- Schema (Are the inputs and outputs each consistent with what was planned for?)
- Statistical process control
- Model execution volumetrics (Are model inputs and requests in the range of the expectations?)
- Latency

If these or any other operational metrics fail to fall within pre-set thresholds the monitoring system should issue alerts, and ideally would attempt automated remediation actions. The monitoring solution must be able to easily integrate with enterprise systems and data sources, because automated remediation often requires orchestration across multiple input and output systems.



The Current State of Model Monitoring Needs to Improve

- 50% of organizations say their ability to monitor models to detect and fix issues in a timely way is ineffective or not very effective
- Only 32% say their methods are highly effective for maintaining visibility into model state and status

Source: 2021 State of ModelOps report

Monitoring Model Quality

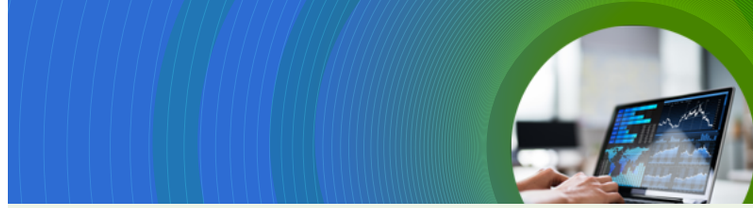
With data and business variables changing all the time, how do you know if your carefully built models are still relevant and able to produce advantageous insight?

Quality monitoring should be aligned to ensuring model decisioning and outcomes are within reasonable, predetermined parameters. Besides the data validation input that occurs as part of a thorough operations monitoring process, quality monitoring also needs to address:

- Data drift – this tracks how data sources are changing over time & how that is affecting model performance
- Model concept drift/inference drift, where outputs are tested to detect whether performance is within thresholds
- Model decay is measured by applying statistical performance monitoring to model outputs and comparing them to known outcomes; this helps determine if inferences and predictions are becoming less accurate, and why

When input data and model execution quality suffer, so does model performance. Depending on the model, inaccurate predictions and insights can lead to:

- Missed revenue opportunities
- Higher operations costs
- Excessive time spent by data team on troubleshooting
- Reduced customer and employee experience
- Brand reputation damage
- Exposure to compliance violations or lawsuits

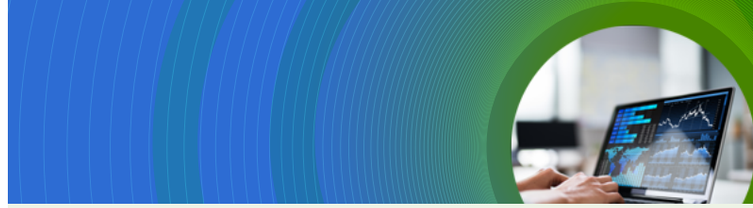


Quality Operations Lead to Quality Results

- 90% of businesses are focused on improving data management resilience to at least some degree over the next year
- 40% of organizations say poor data quality results in wasted resources and additional cost
- 32% say it negatively impacts customer trust and impacts customer experience
- 27% say it hinders compliance with regulatory obligations

Source: “Experian 2021 Global data management research”

Monitoring Model Risk



There are two important forms of risk to address: risk to the models themselves, and the enterprise risk that model performance poses to compliance and liability.

Data professionals report that managing risk and compliance is the most challenging barrier to operationalizing AI models. Risk is a persistent challenge because AI and machine learning models tend to parse and combine data and use it in different ways, which can make documenting provenance and consent difficult.

Managing risk requires continual monitoring of whether the model is operating within established business, risk and compliance ranges. The output must also be continually monitored to ensure results are ethically fair, which is a constant challenge. Actionable monitoring with automated remediation steps enables risk checks and quality controls to be built in, for example to address:

- Ethical fairness monitoring: Is the model demonstrating partiality or bias?
- Interpretability: can you identify the feature weightings that are most important to specific inferences
- Population score indexing (PSI): has the distribution of model inferences deviated from expected ranges
- Characteristic stability: is an input feature being distributed as intended? How much has distribution changed over time?
- Rank order break: consistency of rank ordering with respect too the target variable

Conducting thorough documentation at every stage of the model life cycle, from registration in the model inventory to recording all approvals, validations and other activity, is also important for mitigating risk.

"Sleeping with one eye open is unsustainable. A top complaint of data science, application development & delivery (AD&D) teams, and, increasingly, line-of-business leaders is the challenge in deploying, monitoring, and governing machine learning models in production. Manual handoffs, frantic monitoring, and loose governance prevent organizations from deploying more AI use cases."

Forrester, "Introducing ModelOps To Operationalize AI" August 13, 2020

Monitoring Model Processes

Process monitoring addresses workflow execution. It looks beyond the model itself to continuously monitor the related end-to-end processes associated with the model to ensure all steps are properly executed – plus send alerts and orchestrate remediation steps as required.

Models do not operate in a vacuum. They depend on upstream and downstream processes, and must comply with practices, policies and processes from a variety of business, IT, risk and compliance teams. Without monitoring the ongoing model operations processes, vulnerabilities and exposure to business and governance issues exist.

Effective process monitoring spans activities within operations, risk and quality monitoring. It needs to provide alerts any time a required process step is not completed or not completed in a specified time frame. Processing metrics should be collected and used to identify bottlenecks and detect problematic trends. Some of the specific areas it should address include:

- Model registration
- Risk Management
- Operationalization
- Governance

Process monitoring should have specific capabilities for model governance. Governance tracks and documents the steps in model registration, promotion to production and maintenance. Proper governance ensures audibility, repeatability, and business and regulatory compliance, which is essential for consistent quality performance and scalability.

ModelOps is not DevOps or ITOps by another name and needs its own tooling that addresses the complete life cycle to maximize model uptime, effectiveness and compliance.



"Approximately half of all AI models never make it into production due to lack of ModelOps. Even though data complexity and volume remain a concern, the top hurdles to AI implementations are security, privacy and challenges related to integrating AI systems with existing infrastructure."

Gartner, "Innovation Insights for ModelOps" August 6, 2020

When ModelOps is Done Right…

Maximizing model uptime and optimizing performance through automated monitoring enables organizations to get the full value from their investments in AI.

When model monitoring is done right, barriers to using and scaling AI are addressed – risk threats and compliance burdens are mitigated, and output quality enables organizations to get the full value from their models. That can lead to increased revenue, reduced cost and more success in product and service introductions or whatever the model was developed to support.

In our experience, model monitoring with automated remediation steps produces:

- Increased data scientist productivity by up to 50%, by removing the burden of manual monitoring and remediation
- Improved model accuracy by 10% or more that results in increased revenue, improved customer service, waste prevention or reduced risk
- Faster time to problem detection and model redeployment, plus, automated quality checking and remediation also increases uptime by preventing models from having to be taken offline
- Reduced business risk through the ability to track and retain problems and remediation steps, ensuring models are running in a compliant state throughout their life cycle
- Organizations gain the ability to scale AI through automated and orchestrated monitoring for all models throughout the model operationalization process

As you can see, effective model monitoring produces benefits for the organization that extend well beyond AI operations.



"Model operations (ModelOps) is a must-have capability to operationalize AI at scale."

Forrester, "Introducing ModelOps To Operationalize AI" August 13, 2020 `

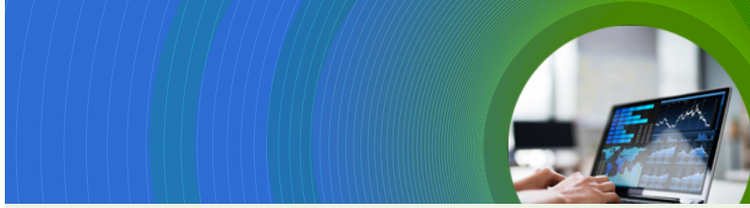
When ModelOps is Insufficient...

When issues aren't promptly detected and remediated, quality problems and delays develop and AI programs can't achieve their full potential value.

Data scientists have many responsibilities and new development often takes priority. As a result, performance problems may go undetected for models in production, which can impact the availability and reliability of the model decisioning. Without automating the monitoring of both model performance and processes, there is exposure for business, risk and compliance failures that impact customer service and business revenue.

One of the biggest dangers to insufficient model monitoring and management is elevated liability and compliance risk. Notably, data science professionals surveyed cited difficulty in managing risk and compliance as their greatest barrier to operationalizing AI, and rated improved control, governance and regulatory compliance their second-most important benefit to automating ModelOps, after cost reduction .

While lack of insight and elevated compliance risk are serious problems for business, they are also preventable through actionable monitoring.



According to the 2021 State of ModelOps Report:

- 64% of organizations say their lack of efficient processes is a barrier to operationalizing models
- 80% have difficulty managing risk and ensuring regulatory compliance for models

How Well Are You Monitoring Now?

Does your model monitoring process need an upgrade?

For most organizations, the answer is yes, because AI efforts, new demands and complexity have outgrown the organizational ability to satisfy them. As recently as 2020 two-thirds of enterprises were not monitoring their models at all, and among those that were, 36 percent were doing so manually. Many organizations have no way of knowing if their model management is effective until an obvious problem develops.

Consider the following questions to assess model monitoring effectiveness at your organization. If your processes and systems can't answer these questions, you may need to update them.

- How many models are in production?
- Where are models running?
- Are decisions being made in a timely manner?
- Are model results reliable and accurate?
- Are compliance and regulatory requirements being satisfied?
- Are models performing within established controls and thresholds?
- How is model performance changing over time?



Calculate Your Potential Payoff from ModelOp Center

See for yourself how much value ModelOp Center can deliver to your business.



Use our **ROI calculator** and get immediate numbers.

Conclusion

AI program effectiveness – and business value – depend on having actionable monitoring with automated remediation steps throughout the model's life.

Traditionally, data scientists have had to spend much more time wrangling data, getting models into production and monitoring and troubleshooting them than have been able to spend on new development. Advancements in automated ModelOps monitoring are changing that. By adopting actionable monitoring with automated remediation steps, organization can gain a competitive advantage by reducing their AI operations costs, taking risk out of operations, improving compliance and freeing data scientists to do more development.

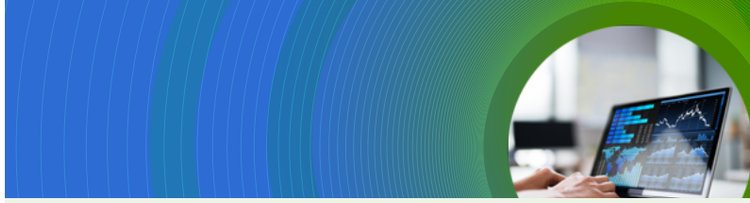
ModelOp can help. ModelOp helps customers govern and scale AI initiatives with the industry's leading enterprise-class ModelOps platform. Large organizations use ModelOp Center to govern, monitor and orchestrate models across the enterprise, to unlock the value from their machine learning and AI investments.

ModelOp Center automates the monitoring of AI and ML model performance, governance and operational processes and orchestrates remediation actions, regardless of where the model is running or type of model, resulting in reliable, compliant and scalable AI.

Automated alerts and notifications are sent for potential and immediate problems. All models are automatically monitored for operational, quantitative and risk management performance. Remediation steps are immediately executed using pre-defined processes that are automated and orchestrated, eliminating model degradation.

In addition, the end-to-end model operational processes (model life cycle) are monitored to ensure business, risk and regulatory controls are adhered to and fully auditable.

Much more information about ModelOp Center, the philosophy behind ModelOps monitoring and insight on how to do it is on our website. For more information visit **www.modelop.com**



For more information, visit
www.modelop.com



To calculate your improved value with ModelOp Center, access the **ROI calculator** and get immediate numbers.

About ModelOp

ModelOp, the pioneer of ModelOps software, enables large enterprises to address the critical governance and scale challenges necessary to fully unlock the transformational value of enterprise AI and machine learning investments. Core to any AI orchestration platform, G2000 companies use ModelOp Center to govern, monitor and orchestrate models across the enterprise and deliver reliable, compliant and scalable AI initiatives.

For more information visit **modelop.com**.



ModelOp

© Copyright 2021 ModelOp