STRIVEWORKS

Model Drift and The Day 3 Problem

Why Remediation Is the Key to Unlocking AI Value

AI and machine learning have fast become priorities for organizations across sectors.

In the span of a few years, AI has gone from bleeding-edge research labs to the most mainstream software applications with millions of daily active users. As consumers, we see AI working seamlessly in the background of our email clients as suggested text and our mobile phones as photo enhancements.

Business and public sector applications are just as common and far more impactful. E-commerce recommendation engines now drive double-digit improvements in conversion rates.[1] Banks use fraud-detection algorithms to prevent billions in losses.[2] The US Military trusts computer vision models to spot enemy aircraft in conflict regions.[3]

Research and consulting firm Gartner projects AI market growth will accelerate to a 31.1% compound annual growth rate (CAGR) by next year—racing past the overall software market.[4]

Since the launch of ChatGPT in 2022 demonstrated the capabilities of current-era AI, organizations across sectors have doubled down on the technology as part of their growth strategy. Yet, any organization with more than a few machine learning (ML) models in production knows that the technology comes with business challenges. It takes a lot of effort to create models efficiently and deploy them effectively. **It takes even more effort—which is often overlooked—to maintain models indefinitely**.

What's Different About Machine Learning?

ML is different from other technologies. Its speed and capabilities are revolutionary, but the output of ML models is often difficult to predict or explain. Models are trained—not programmed like traditional software. They often behave in ways that even their creators do not fully understand.

When it works, ML is transformational—but, too often, it flounders. *Scientific Reports* notes that **91% of ML models degrade over time**.[5]

Various issues can arise when models enter production. Models may work well at first, but performance may drop over time as incoming data shifts until it no longer resembles the model's training data. In other cases, changes in human behavior may lead to a mismatch between a model and its data pipeline.[6] Or updates to one model may produce inferences that a formerly effective downstream model no longer understands.

This is old news for ML engineers but a growing surprise to leaders who have invested millions in the AI capabilities of their organizations only to watch them struggle.

At Striveworks, we call this **The Day 3 Problem**—and it's the number one reason AI programs fail to produce any value.

31.1% *Projected AI market CAGR through 2025*

1 https://martech.org/roi-recommendation-engines-marketing/
2 https://stripe.com/resources/more/how-machine-learning-works-for-payment-fraud-detection-and-prevention?
3 https://www.striveworks.com/blog/how-geospatial-machine-learning-transforms-decision-making-with-7-use-cases
4 https://www.gartner.com/guest/purchase/registration?resId=4007140
5 https://www.nature.com/articles/s41598-022-15245-z
6 https://towardsdatascience.com/the-covid-19-concept-drift-using-sydney-ferry-activity-data-32bbff63cb9f

“The Day 3 Problem” Defined

Consider the standard process for organizations using AI/ML.

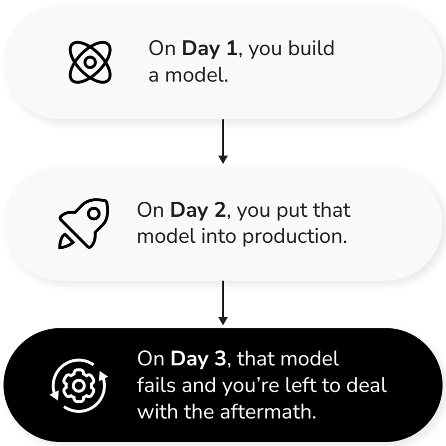

This scenario happens daily at any organization that has more than a handful of models in production. But it is often overlooked.

“Temporal model degradation [is] a virtually unknown, yet critically important, property of machine learning models, essential for our understanding of AI and its applications,” writes Vela et al. in *Nature*. “AI models do not remain static, even if they achieve high accuracy when initially deployed, and even when their data comes from seemingly stable processes.”

Whether models start to struggle on Day 3 or Day 33, each time it happens sparks a round of questions from ML teams and management alike:

- What is going on?
- Is this a short-term issue or a persistent problem?
- Is the problem with our data or our model?
- How inaccurate are the results coming from our model, really?
- What business decisions did we make on incorrect insights?
- How can we stop this from happening?

“Models fail because the world is dynamic,” says Jim Rebesco, cofounder and CEO of Striveworks. “The statistical phrase is ‘non-stationary,’ which means that the data being put into a model in production is different from the data it was trained on.”

ML models are created by training an algorithm on historical data. They ingest thousands or even millions of data points—images, rows of numbers, strings of text—to identify patterns. They excel at matching new data points to similar examples that exist in their training data. But in the real world, models encounter situations that appear different from a set of specially curated data points—and **even slight differences can lead to bad outcomes**.

“Take a predictive maintenance model trained on a particular engine,” says Rebesco. “Maybe we originally deployed this model in the summer and trained it on production data from the same period of time. Now, it’s winter, and you’ve got thermal contraction in the parts, or the lubricating oil is more viscous, and the data coming off the engine and going into the model looks a lot different than what it was trained on.”

What Causes Model Drift?

Many factors contribute to this shift in model reliability, known alternately as model drift or model degradation *(see **Table 1**, p.4)*. In all instances, though, the result of drift is an ML model that generates predictions from real-world data to the best of its ability, but those predictions are inaccurate—making them useless or even harmful. In 2023, a paper from Stanford researchers showed that, for simple math questions, even OpenAI’s flagship GPT-4 dropped in accuracy by **95.2%** over a three-month period.[7]

Inaccuracies of this order have far-reaching consequences for organizations relying on these models. Often, models will continue to generate bad inferences for long periods of time before a human notices the problem. By that point, organizations have made business or operational decisions based on these supposed insights.

7 https://arxiv.org/pdf/2307.09009

While some organizations can tolerate this imprecision, others have much higher stakes. If a streaming service's recommendation engine suggests an unexpected TV show, management might not worry about **The Day 3 Problem**. But a bank that relies on AI to make investments, or a military that trusts AI to distinguish allies from enemies during airstrikes, requires a much higher degree of precision. In these scenarios, a broken model can spur a catastrophe with devastating consequences.

This build-deploy-fail process is a cycle that AI-driven organizations see time and again. With all of the challenges that contribute to model degradation, leaders may wonder if anything can be done to prevent or recover from it.

Often, the answer is to pull the plug. But while taking a model out of production stops the influx of bad insights, it introduces other negative consequences—namely, **delaying the capabilities an organization was relying on or destroying returns on investment.**

Fortunately, there are ways to reduce or eliminate the effects of model drift. For too long, a critical step in the machine learning operations (MLOps) workflow has been ignored—but it's the key to restoring models to strong performance and keeping them reliable in production.

That step is remediation.

Example of Adversarial Adaptations

©Maxar

Table 1: Why Models Drift

Many factors contribute to this shift in model reliability, known alternately as model drift or model degradation.

Cause of Drift	Definition	Example
Natural Adaptations	Data changes in response to outputs from a model.	A financial trading model sells a stock because other models are selling it, driving down its price.
Adversarial Adaptations	An agent's behavior changes in order to evade a model.	An enemy airforce attaches tires to its airplanes to trick models trying to detect them.
Use Case Differences	A model that functions well in one context produces poor results in another.	A sentiment analysis model tracking US-China relations fails to interpret US-Japan relations correctly.
Time Sensitivities	A model trained on data from earlier time periods misses new contextual changes.	A model that understands US-Japan relations in the 1940s fails to produce useful insights for US-Japan relations today.
Chaotic Interference	A change in an upstream model's settings introduces inaccuracies into a downstream model.	A change to an embedding model causes a text classification model using its outputs to label everything wrong.
AI Aging	The process by which random variations in training can contribute to accelerated degradation.	An effective object detection model simply starts to waver more and more in accuracy over time.

Remediation: The Solution to The Day 3 Problem

Organizations with a reliance on ML models need to understand remediation: a critical function for maintaining the long-term performance of ML models and the crux of post-production support.

Simply, remediation is the process of restoring the performance of an ML model and returning it to production.

"In the MLOps space, everybody has been saying, 'You can build and deploy models.' Of course that's valuable, but even the language people are using is wrong," says Eric Korman, cofounder and Chief Science Officer of Striveworks. "They need to think about the **full life cycle of a machine learning model**: build, deploy, and maintain."

Until now, the standard workflow for MLOps has been as follows:

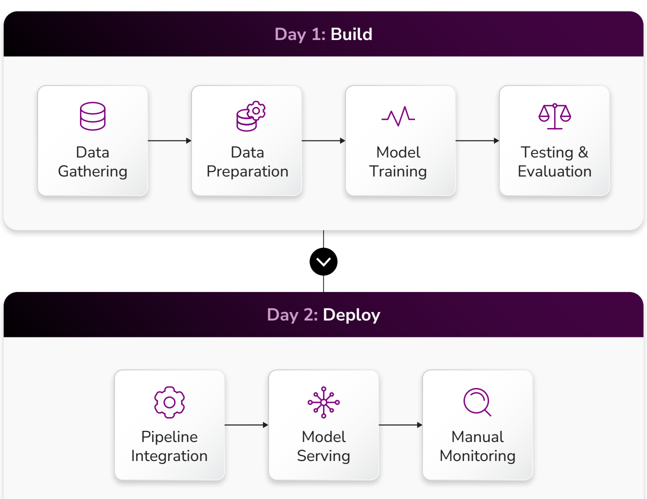

All these steps are critical for getting ML models *into production*. But in a dynamic world where models constantly experience drift, ending the process with monitoring is shortsighted.

Today, at best, teams get alerted when a model's performance is suffering. An additional remediation step is necessary to bring the model back to good standing.

"For too long, the only remediation capability that an AI-powered workflow had was the kill switch—turning off the system," says Korman.

The standard needs to change.

A modern vision of the MLOps life cycle involves making adjustments throughout your workflow so your model's results are continuously accurate and useful. That process may involve model retraining—and likely involves many additional steps to make your workflow more effective.

The Day 3 Solution

1. Identify model candidates for remediation.
2. Determine root causes of performance degradation.
3. Explore and resolve problems with upstream data pipelines.
4. Explore and resolve problems with downstream models and applications.
5. Test and evaluate models for efficacy.
6. Retrain models.
7. Redeploy updated models and data pipelines, including integrations.
8. Monitor models in production on a continuous basis to identify further issues.

Until recently, engineers had no simple, standardized method to determine if a model had drifted. Even today, too many organizations are ignoring the critical need to build remediation into their MLOps workflows.

To unlock the truly transformative capabilities of ML models and advance them into positions of reliability, trust, and value at organizations, ML teams need to place remediation at the center of their MLOps process—and they need the right tools to enable them.

The Ideal Machine Learning Workflow

Organizations aiming for AI success over the long haul need a different MLOps workflow.

The standard approach only takes them to Day 2. **Teams need tools for Day 3 and beyond**.

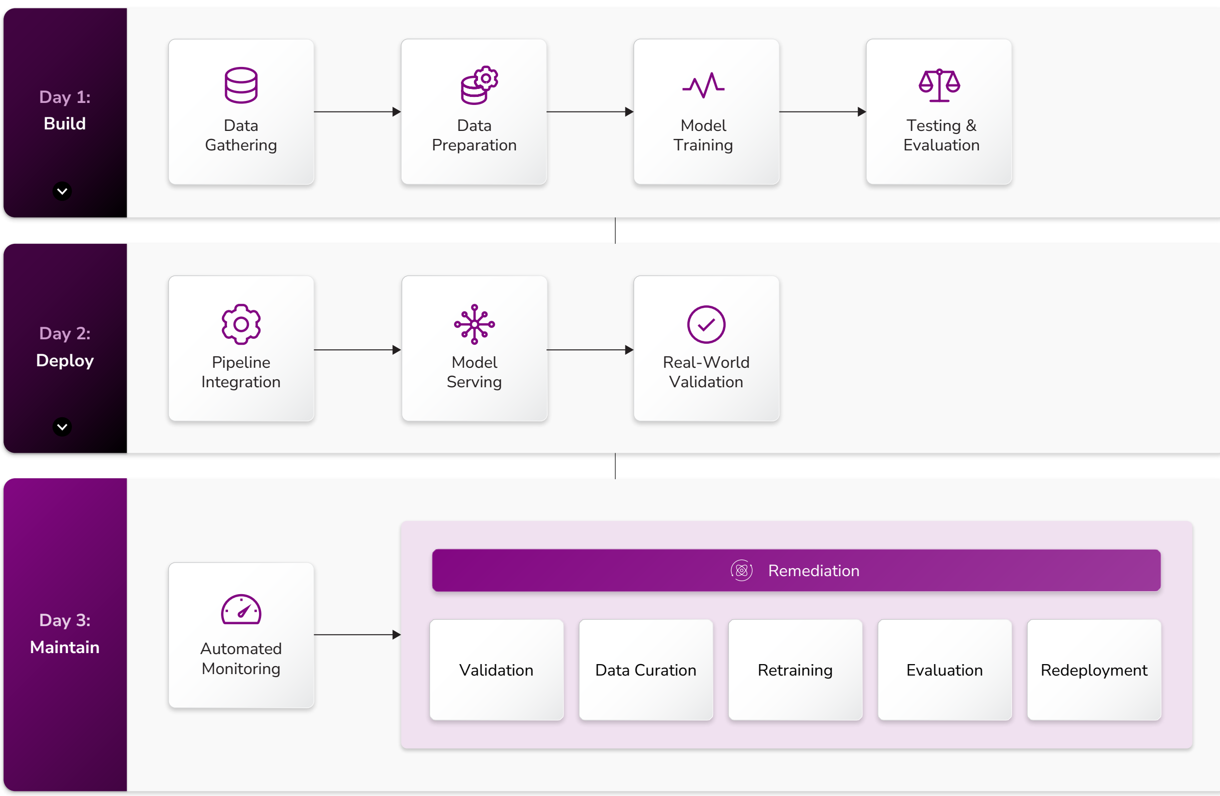

All organizations using AI follow some variation of the first two phases of this process, in which they gather a dataset, annotate it, use it to train a model, and deploy the model into production with any associated pipeline integrations.

At this point, though, most MLOps workflows end—ignoring the critical risk of **The Day 3 Problem**.

The ideal MLOps workflow includes concrete steps for mitigating issues specific to the post-production phase, including automatic monitoring of production models for data drift.

When signs of drift are detected, the workflow initiates a remediation process that validates drift is occuring, curates new training data, and cycles through an accelerated training, evaluation, and deployment process to get an updated, effective model back into production in short order.

How Does Remediation Work?

Remediation is a critical function for maintaining the long-term performance of ML models in production. But how does it work?

It is helpful to think of remediation as the next step in the post-production phase of the MLOps workflow. It includes monitoring production models, evaluating production data and models, and retraining.

Before you can remediate a failing ML project, you need to know that your model has drifted or degraded in the first place. ML engineers and data scientists often have good heuristics to identify problems with models.

However, **automating drift detection** enables ML teams to pay consistent, close attention to all models in production—hundreds or even thousands of active models—preventing the common issue of models generating poor results for weeks or months without anyone noticing.

Once you have correctly identified that a model needs remediation, you have two fundamental questions to ask:

1. Is the problem with the data?
2. Is the problem with the model itself?

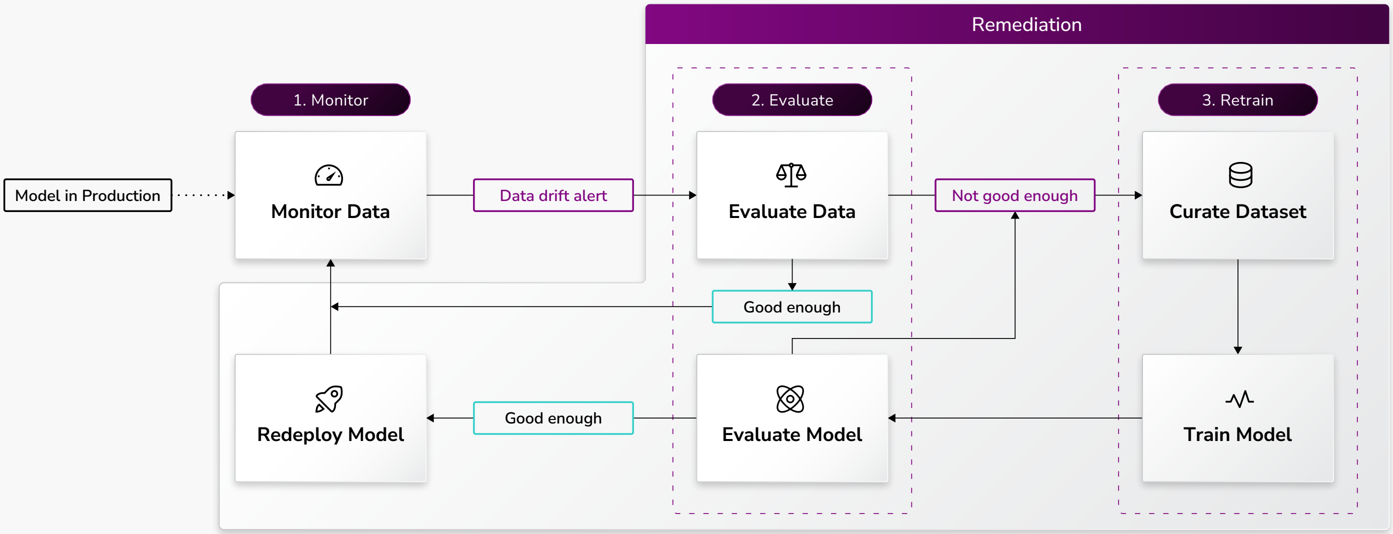

Model post-production starts with an automated process for detecting drift occurring on ML models. Two standard multidimensional tests—Kolmogorov-Smirnov and Cramér-von Mises—are widely used to confirm whether production data is out of distribution with models' training data. If the incoming data falls below critical thresholds, a data team then intervenes to evaluate if, indeed, model drift is underway. If so, teams need a new, more applicable dataset to fine-tune their model—a process that takes notably less time than training a new model. Remediated models are then re-evaluated for efficacy and redeployed into production, where the cycle begins again.

Fixing Problems With Data

If the problem resides in your data, fixing it can prove challenging. Problems with a model may require you to examine its history and configuration, but problems with your data may come from anywhere in your MLOps workflow—from preparation of original training data to the live data your models see in production. To understand your data problem, you need the ability to **audit your data lineage from beginning to end**.

If you have full auditability, then fixing your data problems becomes much simpler: You comb through your data to find anomalies or errors, such as mislabeled datums in a training set or an upstream data source with new parameters. Then, you address the specific issue with a point solution—a new dataset for a new training run or a patch for an application programming interface (API).

Unfortunately, auditability is still immature in most ML workflows. To gain an understanding of data lineage, engineers typically need to alter software significantly or write customized code, especially when their workflows involve calls to external data services.

"Custom coding for data lineage is simply insufficient," says Rebesco. "Today, any organization can ensure that their AI-powered workflows have automated processes that can alert nontechnical users to identify, confirm, and report errors."

Observability, explainability, and auditability have become hot topics in ML for this very reason. Too often, ML models—especially deep learning neural networks—operate as a black box, providing little insight into how they derive their outputs. But this lack of information inhibits users' abilities to identify errors and remediate their ML workflows.

Engineers need to see their data and understand how models are transforming it before they can consider remediation. Fortunately, the standards are starting to change. Striveworks, for example, has developed a patented process that gives engineers access to their full data lineage, including calls to external services.

WHAT IS DATA LINEAGE?

Data lineage refers to the full history of data points in an ML workflow. By inspecting a workflow's data lineage, ML teams can observe and understand the factors contributing to errors and anomalies, including model drift.

Striveworks has been assigned **US Patent 11853304** for a unique data lineage process that enhances transparency and auditability for ML models. This process automatically captures activity involving data throughout workflows—including calls to external services—saving time and resources for data science teams.

"Without this process, developers would need to build a custom one and rigorously enforce compliance across their teams," says **Matthew Griffin**, the Striveworks software engineer who was awarded the patent. "This process ensures that happens as part of the normal workflow, so developers can focus on their real objectives."

Fixing Problems With Models

In comparison to fixing data problems, fixing a model appears straightforward: You can retrain your current model or replace it with a new one. But even these options are more complex than they seem.

> According to Nature, "Retraining a model on a regular basis looks like the most obvious remedy to AI aging, but this is only simple in theory. To make retraining practically feasible, one needs to, at least:
>
> – develop a trigger to signal when the model must be retrained;
>
> – develop an efficient and robust mechanism for automatic model retraining; and
>
> – have constant access to the most recent ground truth."

Continuous monitoring of production models can create that initial trigger for retraining. The Striveworks platform incorporates automatic drift detection that compares incoming data against your model's training data. When it finds that production data departs significantly from training data, it alerts users that data drift has occurred.

Then ML teams have a few options to restore the quality of their workflow's outputs. They can try another model that may be more appropriate to the real-world data. They can source more data, often from a third party, that may prove to be a better match for the model.

But perhaps the most effective way to remediate a model is to retrain it on its actual production data.

For example, the Striveworks platform writes all model inferences into a **persistent inference store.** When it comes time to retrain a model, ML teams can take the model's recent inferences and use their own best data to form a new training dataset. Because this data is coming directly from the real-world deployment of the model, it is uniquely suited for training that model for the scenario it has encountered in production—certainly more so than generic, third-party data.

Choosing the Best Model

Of course, ML teams remediating models are compelled to wonder if the models in production are even the most appropriate options for their use cases.

"When remediating, you want to ask some hard questions," says Rebesco. "Is this model better than the old one? Will it continue to be better? What are its failure modes? Is there an ability to evaluate and compare the models with other ones?"

This uncertainty around model choice makes testing and evaluation critical in any MLOps workflow. Many models may show similar characteristics but deliver varying results, even after fine-tuning with training data.[8]

Testing and evaluating multiple models can bring additional assurance that the model your team is about to deploy is the best choice available for its specific application. In an ideal workflow, ML teams perform these evaluations when selecting a model for initial deployment. But **testing and evaluation are also critical after retraining**—when teams may wonder if their remediation efforts did the job of making their models more suitable for their specific circumstances.

In organizations with multiple models in production, this function can save days of effort for ML teams. Yet, tools for testing and evaluation are only now becoming available. Striveworks released its open-source testing and evaluation service, Valor (**https://github.com/Striveworks/valor**), to get evaluation tools into the hands of the data science community. This service enables engineers to compare various candidate models and determine the best option available from their catalog to deploy into production—a vital ability as more models become available every day.

Striveworks Model Inference Store

8 https://www.nature.com/articles/s41598-022-15245-z

Predicting Lightning-Induced Wildfires—A Remediation Case Study

Lightning is the leading cause of wildfires in non-tropical forests,[9] responsible for destroying 4,206,980 acres in 2022 alone.[10] Standard approaches to wildfire identification are ineffective for large-scale prevention. The US alone has over 751 million acres of forest—much too large an area for firefighters to inspect each lightning strike location.[11] Likewise, lightning strikes often spark small ignitions that smolder for several days before smoke becomes visible.

If firefighters could apply the latest technology to narrow their searches, they could save critical time and even eliminate blazes before they begin.

In 2022, a Fortune 500™ government contractor commenced a project to use data from the National Oceanic and Atmospheric Administration Geostationary Lightning Mapper (GLM) to predict the likelihood of lightning-induced wildfires.

Depending on environmental and atmospheric conditions, each lightning strike has a range of potential to ignite a wildfire. The data science team combined GLM data with raster imagery and tabular weather data to track more than 150 features in an AI model for predicting the risk of wildfires. Because these datasets varied by data type and resolution, data fusion was necessary. Models for this project were built and deployed using the Striveworks platform.

The project proceeded according to plan, and the model performed well—until the winter and spring of 2022/2023, when the area under observation experienced unusually wet conditions. The ML team saw a classic representation of **The Day 3 Problem**: A model that once performed well on training data was now drifting in response to real-world conditions.

	Goal	Actual
Peak 14-day moving average performance (F1 score)	0.65	0.80+
Accuracy	0.85	0.87+
Recall	0.65	0.90+

87%

Accuracy above project scope expectations

25

Minutes end-to-end, from data processing to automated report delivery

1.1M

Predictions in 6 weeks

9 https://www.preventionweb.net/news/lightning-leading-cause-wildfires-boreal-forests-threatening-carbon-storage#:~:text=9%20November%20 2023-,Lightning%20is%20the%20leading%20cause%20of,boreal%20 forests%2C%20threatening%20carbon%20storage&text=New%20research%20 shows%20that%20lightning,will%20increase%20with%20climate%20change

10 https://www.nifc.gov/fire-information/statistics/lightning-caused

11 https://forests.org/so-how-much-forest-is-there-in-the-us-and-canada/

The ML team commenced remediation, starting with identifying potential failure modes. The data scientists then trained a new model using an appropriate range of target data. The team then deployed this remediated model back into production using the Striveworks platform. The new production model exceeded all customer goals for accuracy, F1 score (a measure of predictive performance), and recall (also known as the true positive rate).

The improvements generated immediate value. The team used this model consistently to aggregate 96 hours of multisource data in just 25 minutes—a capability that is functionally impossible otherwise. The entire process was automated to deliver this analysis to firefighters just in time for daily mission planning meetings, ensuring that they focus their limited resources to scout areas most likely to develop into a blaze.

Ultimately, this pilot project reaffirmed the critical importance of remediation for sustaining the performance of ML models as they face fluctuations in real-world data. It also underscores the Striveworks approach to ML post-production, including the continuous monitoring of production models that makes rapid remediation possible.

WHAT IS DATA FUSION?

Data fusion is a methodology that combines data from multiple disparate sources to create a more nuanced, accurate insight.

Fusion is challenging, as each data source may have a distinct format, structure, quality, or other characteristic that impedes standardization. It requires large amounts of configuration and data handling as well as heavy computational resources.

The wildfire detection project fused the following data types to predict the likelihood of lightning-induced fires:

- Tabular data from the NOAA Geostationary Lightning Mapper (GLM)
- Tabular representations of lightning strike data
- Three-dimensional and time-series data from the NOAA
- Raster imagery of forest cover

The process has wide-ranging applications in fields such as health care, agriculture, finance, and cybersecurity.

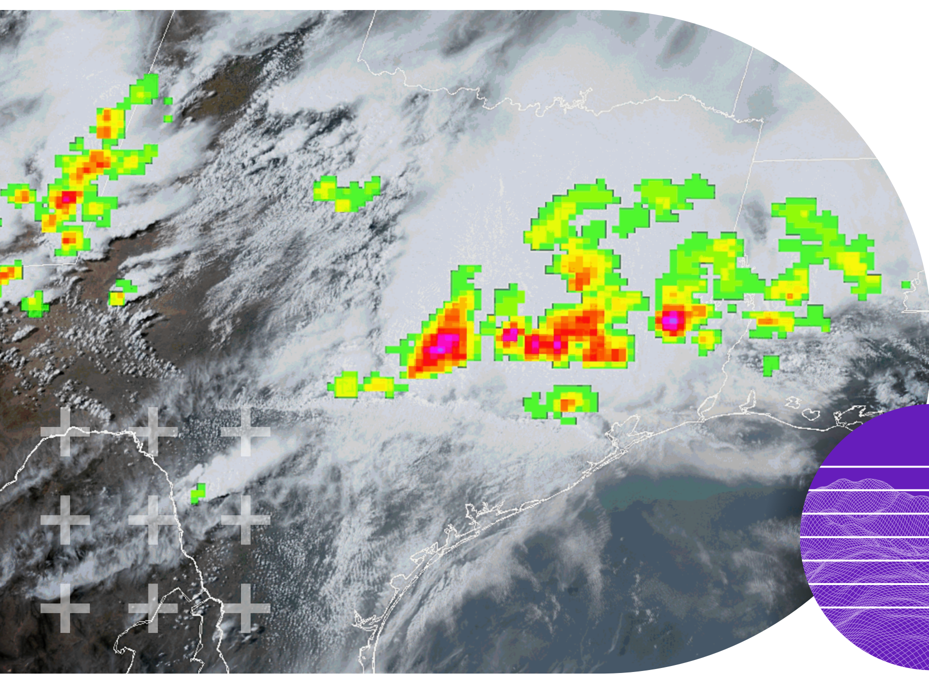

The Tool That Keeps Machine Learning Models Reliable

The Day 3 Problem is a persistent and pernicious concern for any organization moving toward greater dependence on AI. Drift, degradation, and obsolescence are unavoidable with ML. Any time a model encounters data from the volatile real world, this problem will appear and demand that project owners find a solution.

In their paper in *Scientific Reports*, Vela et al. set out several ways to combat model degradation:

- Treat ML models as dynamic systems.
- Choose stable models to put into production in the first place.
- Monitor and test models often.

All of these are important changes for overcoming the challenges of **The Day 3 Problem**, but ML teams still need solutions and tools to put these ideas into practice.

> "Training a model that performs well on Day 1 or 2 is table stakes. Ensuring that [a] model keeps performing as promised on Day 3 and beyond—without toil—is the real test of an MLOps platform, and where Striveworks truly shines."
>
> ***Kutta Srinivasan***
> *Tech Executive, Advisor, and Investor*

Remediation: Beyond Retraining

Retraining is critical, but it only forms one part of the full remediation process. Organizations cannot rely on retraining alone. To make their models consistent and reliable over the long term, ML teams need a solution that delivers three key functions:

- It monitors models in production for early signals of drift.
- It tests and evaluates models for suitability.
- It audits data lineage and retrains models on production data.

As this approach matures, it will only become more open and streamlined. Both current and future ML teams need to remediate a range of models from various vendors in production in various environments, and they need it to happen automatically.

"This is what we mean by 'disappearing MLOps,'" says Rebesco. "It's our vision for Striveworks: enabling the MLOps life cycle to become so trusted and seamless that it all happens under the hood. Remediation is the missing key that makes it possible."

Until then, the most foundational point to remember when dealing with **The Day 3 Problem** is a natural one: Models make mistakes, and those mistakes have consequences. No organization would accept ongoing human errors for weeks or months without intervening. Models need similar consideration and similar opportunities to improve. At a higher level, organizations need to become fault tolerant, understanding their range of ML failure modes and their range of options to address them.

With appropriate processes and tooling in place, solving **The Day 3 Problem** is set to become merely another step in the MLOps process. Forward-thinking ML teams no longer need to treat it as a persistent blockage dragging down their AI program, but a blip in their jobs of delivering consistent, reliable results from AI for months or years into the future. ■

About Striveworks

Build, Deploy, and Maintain AI for an Unpredictable World

AI is driving a new Industrial Revolution. But most AI tools only work when the world looks the same tomorrow as it did yesterday. That's rarely the case.

Striveworks supports machine learning operations (MLOps) for an ever-changing world. We empower organizations to rapidly build models, deploy them in one click, and maintain them to sustain results at scale—even when the world changes around them.

As a result, Striveworks was recognized as an exemplar in the National Security Commission on Artificial Intelligence Final Report. In 2023, Striveworks placed on the Deloitte Technology Fast 500™ as one of the most rapidly growing technology companies in North America. In 2024, Striveworks was honored with a Built In Best Places to Work award—for the third year running.

References

1. Fagella, Daniel. "The ROI of Recommendation Engines for Marketing." MarTech. Last modified October 3, 2017. https://martech.org/roi-recommendation-engines-marketing/.
2. "How Machine Learning Works for Payment Fraud Detection and Prevention." Stripe. June 27, 2023. https://stripe.com/resources/more/how-machine-learning-works-for-payment-fraud-detection-and-prevention?.
3. "Geospatial Machine Learning Transforms Decision-Making (7 Use Cases)." Striveworks (blog). https://www.striveworks.com/blog/how-geospatial-machine-learning-transforms-decision-making-with-7-use-cases.
4. "Forecast Analysis: Artificial Intelligence Software, Worldwide," Gartner. Accessed April 23, 2024. https://www.gartner.com/guest/purchase/registration?resId=4007140.
5. Vela, Daniel, Andrew Sharp, Richard Zhang, Trang Nguyen, An Hoang, and Oleg Pianykh. "Temporal Quality Degradation in AI Models." Scientific Reports, 12, no. 11654. 2022. https://doi.org/10.1038/s41598-022-15245-z.
6. Ma, Victor and Drew Jarrett. "The COVID-19 Concept Drift, Using Sydney Ferry Activity Data." Towards Data Science. April 23, 2020. https://towardsdatascience.com/the-covid-19-concept-drift-using-sydney-ferry-activity-data-32bbff63cb9f.
7. Chen, Lingjiao, Matei Zaharia, and James Zou. "How Is ChatGPT's Behavior Changing Over Time?" 2023. arxiv.org/pdf/2307.09009.
8. Vela, Daniel, Andrew Sharp, Richard Zhang, Trang Nguyen, An Hoang, and Oleg Pianykh. "Temporal Quality Degradation in AI Models." Scientific Reports, 12, no. 11654 (2022). doi: https://doi.org/10.1038/s41598-022-15245-z.
9. "Lightning Is the Leading Cause of Wildfires in Boreal Forests, Threatening Carbon Storage." United Nations Office for Disaster Risk Reduction. November 9, 2023. https://www.preventionweb.net/news/lightning-leading-cause-wildfires-boreal-forests-threatening-carbon-storage#:~:text=9%20November%202023-,Lightning%20is%20the%20leading%20cause%20of,boreal%20forests%2C%20threatening%20carbon%20storage&text=New%20research%20shows%20that%20lightning,will%20increase%20with%20climate%20change.
10. "Lightning-Caused Wildfires." National Interagency Fire Center. 2023. https://www.nifc.gov/fire-information/statistics/lightning-caused.
11. "Forester's Blog: So, How Much Forest Is There in the U.S. and Canada?" Sustainable Forestry Initiative (blog). August 9, 2022. https://forests.org/so-how-much-forest-is-there-in-the-us-and-canada/.

STRIVEWORKS

STRIVEWORKS

STRIVEWORKS.com

3520 Executive Center Drive, Suite 300
Austin, TX 78731

sales@striveworks.com