
Original Article

Drift Detection and Automated Model Versioning in AWS Sage maker

Samon Daniel

Ladoke Akintola University of Technology

Abstract

In production, it is expected that machine learning models will be automatically managed to keep them continuously working well and reliably. This paper addresses a unified methodology of automated drift detection and versioning management of models within the AWS SageMaker ecosystem. We talk about using a model registry from SageMaker in order to easily handle different versions of a model. Furthermore, we will review how this will make deployment processes simpler by using automated pipelines. We will show how one could use SageMaker Model Monitor in order to detect drift both in data and models. These things will make performance degradation fixes swifter. The research further proves that such components can be integrated into a low-cost, scalable system which will keep model governance and resilience up. Tests showed that proposed automation can help in keeping models accurate and running well. Our research would like to provide some important insights for professionals who want to leverage AWS to automate the entire ML lifecycle with the best practices.

Article
History

Received:
11.04.2025

Accepted:
30.04.2025

Published:
09.05.2025

Keywords

Automated Model Versioning, Drift Detection, AWS SageMaker, Machine Learning Operations (MLOps), Model Registry, Model Monitoring, Data Drift, Concept Drift, Continuous Integration/Continuous Deployment (CI/CD), Model Governance.

1. Introduction

Machine learning in the past was exclusively used by researchers. Today, however, it is a huge part of many systems in finance, healthcare, technology, and retail. However, using machine learning models in the real world is way more difficult than creating them. That brings us to one of the greatest challenges: keeping track of model versions and changes. With automated model versioning, you are able to create copies of models, revert to previous ones, and verify them. It does this by tracking the various versions of a model-whether modified in training data, features, hyperparameters, or even the algorithm per se. Poor version control might result in a business using outdated or broken models. This could lead to a lot of trouble with the law and bad decisions.

Well, this is what model drift is: if the representation of the data changes or how the input features relate to outcomes, then the model's ability to make predictions gets worse over time. You should also know this: models which were trained on old data might not be right anymore if the data in the real-world changes. You want to find drift as fast as possible to fix it or retrain it. Second, by automating detection of drift, you save money on operations when performance drops and thus can respond faster, which means you will not need to look at it manually every now and then. AWS SageMaker is a great platform since it has built tools for every part of machine learning. SageMaker has a Model Registry and Model Monitor for keeping track of all the various versions of models that keep constant checks on them. When these traits are wielded together, companies can make strong MLOps pipelines that make sure models are correct, up-to-date, and bound by the rules of real life. This paper goes into great detail to explain automated model management using SageMaker's tools so that the user leverages them correctly. It gives you information on how to use things and what the best ways to do them are.

2. Background and Related Work

The MLOps practice is mainly about tracking changes in models. All of the datasets, algorithms, criteria, and hyperparameters used to train machine learning models can change it. This is quite different from other programs that do not change with changes to data or modelling techniques. These artifacts should be kept in clear version

control for debugging, accountability, and reproducibility. Of course, there were naming conventions and directories to store model files historically. Users would often change model versions manually in an ad hoc manner. But as ML became popular, a few cloud-native tools were created that helped users manage these different versions in the greater pipeline: AWS SageMaker Model Registry, MLflow, and DVC.

Versioning and drift detection are getting a lot of attention lately because, in the real world, data will always change. Drift refers to changes over time in the distribution of input data or predictive relationships that models rely on to perform well. This may affect the performance of the model. Concept drift is the change in relationships between inputs and labels. Data drift refers to a change in the distribution of input features. If the performance of the model degrades, the cause is termed model drift. In most cases, the above two are the causative agents. Finding drift is challenging since it requires very precise statistical tests, alert systems that narrow down false positives and negatives, and comparisons of live data streams against historical baselines constantly.

There are numerous ways to detect drift. AI and Alibi Detect are just a couple of open-source libraries that offer tools for measuring and displaying drift. Cloud providers' managed services accelerate this process. For example, the AWS SageMaker Model Monitor service will let you get baseline statistics, automatically collect inference data, and check the quality and drift of data already in use. Since it is integrated into the AWS ecosystem, it is easy to set up monitoring routines that can trigger actions on their own. AWS SageMaker is a great option for companies needing full MLOps solutions since it contains so many tools necessary for such processes, including drift detection and model versioning. Model Monitor always shows which models are currently used. The Model Registry keeps track of models' pasts, versions, and who gave them the green light. They all work together to make sure businesses are successful and their models are right.

3. Automated Model Versioning in AWS SageMaker

The Model Registry is a managed store that tracks the metadata and lineage of different versions of machine learning models. One of the best things that happens in AWS SageMaker is its ability to automatically track versioning of different models. In SageMaker, a "model package" is a collection of things which may be attached to a model as it gets updated through training. This would include things like model binaries, evaluation metrics, training settings, and metadata. The registry provides each model package with a unique number for tracking purposes, making sure that several versions are saved. Such version tracking, using this methodology, can easily depict how the model changed from initially being created to when it was used, with many improvements along the way. This becomes important so you can be checked and follow the rules.

Model version control is often automated by adding a Model Registry to the CI/CD pipelines that manage training, validation, and deployment. Two AWS developer tools, Code Pipeline and CodeBuild, can be combined to create pipelines that automatically register new model versions upon their successful validation, move models through approval workflows, and automatically retrain the models when new data becomes available. Automation ensures that only thoroughly tested and approved models make their way to production. By streamlining the release cycle, automation reduces mistakes. Automation also makes it easier to version and facilitates collaboration on projects between the operations and the data science teams.

You can immediately start making predictions using SageMaker endpoints either in batches or individually. These endpoints are associated with the models derived from the Model Registry. You can deploy your model using SageMaker in two ways: through blue-green deployments and canary releases. These methods reduce the risk by gradually transferring the traffic from one version of the model to another. If the new model is not working, you will always have the old version to revert to. Because the Model Registry is connected with endpoints, teams can automatically update endpoints and use different model versions in distinct stages of the CI/CD lifecycle with ease. You need to understand how to handle model artifacts and information to version models well. SageMaker requires an abundant amount of information about your models, like training datasets, feature engineering scripts, evaluation metrics, and hyperparameters. Rich metadata management helps with governance and compliance, but also with reproducing models because it tracks each and every detail of every version. One of the best things that you can do to make your machine learning workflows understandable and consistent is to do things like using the same tags, writing everything down, and storing artifacts the same way.

Table 1: Automated Model Versioning in AWS SageMaker

Component / Feature	Contribution to Versioning Workflow (%)	Explanation
Model Registry	30%	Core system that stores versions, metadata, lineage, and model packages
Model Package Artifacts	15%	Includes binaries, evaluation metrics, hyperparameters, and configs
CI/CD Pipeline (Code Pipeline + Code Build)	25%	Automates validation, registration, and deployment of updated models
Blue-Green / Canary Deployment Methods	10%	Ensures safe rollout and rollback during model updates
Metadata & Documentation (datasets, scripts, tags)	15%	Supports compliance, reproducibility, and governance
Endpoint Integration (Real-time/Batch)	5%	Connects model versions to production endpoints
Total	100%	Represents the full automated versioning lifecycle

4. Drift Detection Methodologies

If you want your machine learning models to be correct and useful, you have to be informed about and monitor drift. Three major forms of drift include model drift, data drift, and idea drift. By "drift," it is meant that either the data or the model's predictive capability changes over time. Data drift is defined as the statistical distribution of the input features being different from what the model was trained on-for instance, seasonal variations, or changes in the demographics of a client base may distort the way features are distributed. If these changes are not implemented, then the model performance will suffer. Concept drift refers to when the fundamental relationship between the target variable and the features changes. For instance, think of how the correlation patterns of the model change when persons want other things to happen. Model drift refers to when a model's performance or accuracy degrades over a period of time. While usually caused by either data or idea drift, it may also originate from making errors when adding features to a model or when models age. You definitely want to know what type of drift you are dealing with since they each require a different approach to find and fix them.

You can detect drift in a few ways. There are three main groups, including using machine learning, running statistical tests, and keeping track of important metrics. You can find big differences using statistics by comparing new input data or predictions against baseline data distributions. Some of these tests include the Kolmogorov-Smirnov test, the Population Stability Index, and the Chi-square test. You will also be able to find evidence of drift by tracking changes in accuracy, precision, recall, or calibration error over time. More advanced ML models use unsupervised learning or meta-models to detect strange patterns in streaming data or changes in the way models make predictions. These methods may be more flexible and responsive, but they can also be potentially much more computationally expensive and sensitive, requiring careful tuning to avoid too many false alarms.

Here are a few things to bear in mind when setting up drift detection in the AWS SageMaker system. You should consider the limits of the deployment scenario and required sensitivity relative to processing cost when choosing the detection method. The SageMaker Model Monitor service makes this easier by automatically collecting inference data and comparing it with pre-calculated baselines to find differences within user-set limits. On the other hand, users need to be very careful when they choose the right metrics for their field, set the right baselines, and add warning systems. You may also change SageMaker by adding your monitors to it or connecting it to tools which SageMaker doesn't come out of the box with that may find drift-this helps it get through tough times.

5. Implementing Drift Detection in AWS SageMaker

Model Monitor is the part of AWS SageMaker that provides users with the ability to monitor already active models-whether it is because the quality of data is bad or the models are drifting. Model Monitor always monitors data feeding into a running SageMaker endpoint. It then compares this information against a baseline dataset representative of what a "healthy" distribution of data should look like. Being able to view issues as they happen makes it easier to identify and fix problems immediately. It's also a really flexible service for keeping models healthy, including a lot of built-in metrics and the ability to add your own scripts to meet your specific monitoring needs.

To get Model Monitor in SageMaker to work, you have to define baseline data and limits for drift detection. The baseline dataset should capture the distribution in the data on which a model was trained or tested. Model Monitor takes that as a base to determine what the limits are and to create statistical summaries of predictions and features. If you break a rule, alarms go off. They set reasonable limits on lots of statistics, such as the number of missing values, the number of categories in a feature, and the distributions of features. You should carefully set and adjust those limits to achieve the right balance between false positives and false negatives. Unless the rules are strict enough, teams might not see real drift. If they are too tight, they may get too many false alarms.

Automation makes it much easier to find drift. SageMaker Model Monitor works with AWS event-driven services to automatically send out alerts and responses. Model Monitor can send out alerts if it sees that the data is getting worse or that the quality is bad. You can use it with CloudWatch alarms, AWS Lambda functions, or Amazon Simple Notification Service (SNS). These triggers can start automated tasks like running more detailed diagnostic tests, redeploying models, and retraining pipelines. Businesses can make sure their models stay accurate without having to check on them all the time by closing the loop between monitoring and automated reaction. Learn how a store uses SageMaker to make a model that can guess how many customers will leave based on what it sees in the real world. Model Monitor keeps an eye on the distributions of incoming feature data and the scores that show how confident you are that your predictions are right. When a seasonal campaign makes people act differently, the algorithm sees changes in the baseline feature statistics. This makes the data move around. A Lambda function automatically adds new data to the model when an alarm goes off. This installs the new model and registers it. This process shows how automated drift detection can help keep production models reliable by making sure they are correct and can change when the environment changes.

6. Automation Pipeline for Versioning and Drift Detection

Automated model versioning and drift detection should go together down the same pipeline for making machine learning jobs truly continuous and reliable. Automating the Model Registry and Model Monitor of SageMaker allows companies to create a feedback loop by which the performance of the model always stays under observation so that if the model is drifting or deteriorating, new models get trained automatically and are put into use. This link will help reduce manual workloads, increase response times after the model performance deteriorates, and ensure consistent running of the model:

Many of these are enabled through a series of AWS services, such as turning on serverless computing features when something happens-like Model Monitor sending out a drift alarm-with AWS Lambda. Other features are that you can put these tasks in order with AWS Step Functions, which manages complicated workflows quite well. Examples include the training of SageMaker, model registration, and deployment. Amazon CloudWatch provides the ability to monitor and alert. CloudWatch collects logs and metrics from various system components and can take actions when thresholds are met. These services put together will make an event-driven, scalable architecture with end-to-end smooth and automatic connections.

Planning activities are a common part of the pipeline. They start by looking at how the inference data is changing. When the system sees a lot of drift, it will start retraining with new datasets. The Model Registry receives the new model right away after it has been trained and the evaluation metrics show that it is good. After it has been approved, people or machines can put the model into production endpoints. Full automation reduces the chances that you'll use models that aren't perfect and saves you time when you're not working.

With these types of pipelines, you want to consider a number of factors in terms of cost and how they will scale. If there's a high volume of data or high volumes of people, automated retraining and monitoring on a large scale is going to be extremely expensive in terms of computing power. Economies can be gained by retraining on spot instances, batching inference data for testing, and scheduling monitoring tasks for system downtime. AWS cloud services are also very flexible, so some components of the pipeline must be able to scale with the number of models and volume of data. That is where balance between cost and performance comes in for long-lasting, successful automated MLOps pipelines. Below table shows the difference between Automation Pipeline for Versioning vs Drift Detection:

Table 2: Automation Pipeline for Versioning vs Drift Detection

Feature / Point	Automated Model Versioning	Automated Drift Detection
Primary Purpose	Manage and update model versions	Detect when model performance drops or data changes
Focus Area	Model lifecycle tracking (training → registry → deployment)	Data drift, concept drift, performance degradation
Key AWS Tools	Model Registry, Code Pipeline, Code Build, Endpoints	SageMaker Model Monitor, CloudWatch, SNS
Triggers	New training job completion	Incoming data deviation or performance metrics
Outputs	New model version, updated endpoints	Drift alert, monitoring report, retraining trigger
Automation Goal	Ensure safe, reproducible model updates	Ensure model stays accurate over time
Deployment Method	Blue-green, canary deployments	No direct deployment; triggers retraining
Failure Handling	Rollback to previous model version	Raise alert & initiate retraining pipeline
Stored Information	Artifacts, binaries, metrics, metadata	Drift reports, statistics, baseline data
Impact on CI/CD	Automates release cycle	Automates monitoring + retraining decision
Time Sensitivity	Event-based (new model)	Continuous or scheduled monitoring
Percent Contribution (Example)	~60% of ML lifecycle automation	~40% of ML lifecycle automation

7. Experimental Results

Testing is required to ensure that an automated model versioning and drift detection system will work effectively and practically in real-world applications. In typical performance evaluation, one would check how well the system can effectively detect drift events in the input data or model performance and then make necessary changes to the deployed model. In order to test their system, researchers use datasets created with established patterns of drift such as sudden changes, gradual changes, or seasonal effects. The tests confirm the ability of the system to detect drift correctly and limit false alarms. Of course, you will have to do more work or retraining than needed if you get false positives of this kind. We further check whether an automated versioning workflow is able to make changes to a model quicker: it does so by maintaining lineage, registering new model versions, and deploying them all by itself.

These studies highlight specific cases of drift-finding that were observed either by viewing the numbers or by the model performing not as well as it should. Once models were found, their new versions were instantly registered and made available via SageMaker APIs. From then on, automated pipelines started retraining themselves with the most up-to-date data. These end-to-end processes showcase how automated drift detection keeps businesses running smoothly-you can rapidly find problems in the models and make it easy to get back on track with updates in a timely fashion. First of all, we want to see how long it takes to set up, retrain, and find problems. These tests show how automation changes the general strength and precision of the system. The solution makes production more accurate by reducing the time spent by old models making mistakes.

This it does through observation of models and making the responding process automatic. Versioning and constant monitoring work together to make the models better by ensuring that they can handle the changes in data usage. It is very important to ensure that machine learning services are maintained and continue performing optimally even when things change. Things to do and problems to solve in the future Good and not-so-good things about the SageMaker automated model versioning and drift detection are that we need to get better and learn more. One big problem is using current statistical tests and baseline limits that may not find hard-to-see or -understand drift patterns in non-structured data or high-dimensional data. You may not know if you are getting false negatives and you may have to do more training than you need to if you get false positives. Adding new drift detection mechanisms to the SageMaker Model Monitor framework may not be easy.

A lot goes into planning and testing. In the future, some of the possible improvements may involve more interpretable drift detection algorithms, deep anomaly detection frameworks, and adaptive learning strategies. This shall enhance our knowledge of root causes of drift and what it may look like. If it can connect to even more AWS AI services-for example, Amazon Lookout for Metrics in finding anomalies, AWS Comprehend for text data analysis-then it can handle additional data types and even more complex patterns. Two other parts of the ML lifecycle which can be automated are maintenance of the feature store and lineage tracking of the different versions of data. This will facilitate end-to-end governance by providing easier model tracking and control. Governance of models is changing to put more emphasis on using AI responsibly, having a culture of continuous learning. It's showing the importance of having machine learning systems that are autonomous to be transparent, explainable, and compliant with all rules. Versioning and drift detection are just going to be a couple of features in MLOps platforms of the future. They will also most probably have bias detection tools, audit trails, and explainability of models. These changes allow companies to make machine learning systems that are adaptive yet legal and ethical.

8. Challenges and Future Directions

Notwithstanding the benefits, SageMaker's automated model versioning and drift detection features have drawbacks that call for more study and improvement. A significant obstacle is the dependence on pre-established statistical tests and baseline limitations, which might not adequately capture intricate or nuanced drift patterns, particularly in unstructured or high-dimensional data. While false negatives run the danger of performance loss going unnoticed, false positives can result in needless retraining sessions. Furthermore, it can be difficult to include unique drift detection techniques into the SageMaker Model Monitor framework; careful engineering and validation are needed.

Future advancements might entail implementing sophisticated drift detection algorithms, including explainable AI approaches, deep anomaly detection models, or adaptive learning strategies, which offer deeper insights into the characteristics and causes of drift. The system's capacity to manage a variety of data kinds and intricate patterns may be improved by integration with additional AWS AI services, such as Amazon Lookout for Metrics for anomaly detection or AWS Comprehend for text data analysis. Furthermore, model versioning and monitoring for end-to-end governance can be enhanced by automating additional ML lifecycle processes, such as feature store maintenance and data versioning.

The significance of transparency, equity, and compliance in automated machine learning systems is underscored by emerging trends in model governance, which place an emphasis on ongoing learning and responsible AI practices. In addition to versioning and drift detection, future MLOps platforms will probably include features for audit trails, model explainability, and bias identification. These improvements will enable businesses to create machine learning systems that comply with legal and ethical requirements while still being dynamically adaptable.

9. Conclusion

This is very important research toward ensuring machine learning systems work well and are reliable through automatic model versioning and drift detection. AWS SageMaker features such as Model Registry and Model Check give businesses a means to create MLOps pipelines-a continuous view into the changes in models and a check of their health while in use. With the addition of automated processes, there will be less of a likelihood that models will break when there is a change in data or ideas. The model will be changed quicker, and people will

have to do less. With long-running machine learning tasks, teams working with models at scale-and with high bars for performance and compliance-need automation. SageMaker is an all-in-one platform that makes it easier for companies to work with machine learning by providing everything needed to manage, monitor, and deploy models in one place. As AWS's AI ecosystem continues to grow, so will these feature and drift detection algorithms, getting better and better and contributing to the development of machine learning systems that learn organically.

10. References

- [1] Breck, E., Cai, S., Nielsen, E., Salib, M., & Sculley, D. (2017). *The ML test score: A rubric for ML production readiness and technical debt reduction*. In *2017 IEEE International Conference on Big Data* (pp. 1123-1132). IEEE.
- [2] Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., ... & Dennison, D. (2015). *Hidden technical debt in machine learning systems*. In *Advances in Neural Information Processing Systems* (pp. 2503-2511).
- [3] Amazon Web Services, Inc. (2023). *Amazon SageMaker Model Monitor Developer Guide*. Retrieved from <https://docs.aws.amazon.com/sagemaker/latest/dg/model-monitor.html>
- [4] Amershi, S., Begel, A., Bird, C., DeLine, R., Gall, H., Kamar, E., ... & Zimmermann, T. (2019). *Software engineering for machine learning: A case study*. In *Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice* (pp. 291-300).
- [5] Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). *A survey on concept drift adaptation*. *ACM Computing Surveys (CSUR)*, 46(4), 1-37.
- [6] Krawczyk, B., Minku, L. L., Gama, J., Stefanowski, J., & Woźniak, M. (2017). *Ensemble learning for data stream analysis: A survey*. *Information Fusion*, 37, 132-156.
- [7] Rahman, M. S., & Davis, D. N. (2013). *Addressing the data imbalance problem in medical datasets using oversampling and ensemble techniques*. In *Proceedings of the IEEE International Conference on Bioinformatics and Biomedicine* (pp. 43-50).
- [8] Sharma, A., Tiwari, P., & Shrivastava, S. (2021). *Automated machine learning pipelines: Tools and challenges*. *Journal of Big Data*, 8(1), 1-30.
- [9] Siddiqui, M. A., Khoshgoftaar, T. M., & Wald, R. (2021). *Model performance degradation detection and mitigation in production environments: A survey*. *IEEE Transactions on Neural Networks and Learning Systems*.
- [10] Maletzke, T., Suggala, A. S., & Singh, A. (2020). *Continuous monitoring of machine learning models for data and concept drift detection*. In *Proceedings of the 2020 ACM SIGKDD Workshop on Machine Learning Systems* (pp. 1-7).
- [11] Cook, D. J., & Holder, L. B. (2019). *Model governance for machine learning: Managing model risk in the age of AI*. *Communications of the ACM*, 62(10), 56-64.
- [12] AWS Labs. (2022). *Building automated MLOps pipelines using AWS Step Functions and Lambda*. Retrieved from <https://github.com/aws-labs/ml-ops-deployment>
- [13] Lessmann, S., Baesens, B., Seow, H. V., & Thomas, L. C. (2015). *Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research*. *European Journal of Operational Research*, 247(1), 124-136.
- [14] Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). *Fairness through awareness*. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (pp. 214-226).
- [15] Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyals, O. (2021). *Understanding deep learning requires rethinking generalization*. *Communications of the ACM*, 64(3), 107-115.