

CASE STUDY

America First Credit Union Pioneers A New Era of AI-Led Growth

With machine learning (ML) observability in place to quickly detect and diagnose the root cause of model performance degradation, America First Credit Union's ML team ships AI with confidence

Introduction

Founded in 1939, America First Credit Union is one of the largest, most stable and most progressive credit unions in the country—a proudly member-owned, not-for-profit cooperative financial institution.

With a mission to provide personal financial services of a superior quality to members-owners, America First has long embraced automation and technology to accomplish that goal. The credit union's early investments in AI are no exception, helping the organization thrive amid evolving consumer expectations and a changing market for financial services.

To help in the critical task of monitoring and troubleshooting machine learning, America First Credit Union selected Arize as its ML observability partner in early 2021.



Key Facts



Industry: Financial Services

About: Credit union with over \$17 billion in assets and 1.2 million members

Models In Production: >30; deploying ~2 new models per month

Primary Use Cases: Loan origination, ongoing credit risk, overdraft risk, potential fraud prevention, product recommendations, and more

Challenges

As America First Credit Union scaled its AI efforts, the limitations of using a business intelligence (BI) tool for in-house model monitoring necessitated an early and prescient shift in strategy. Among the constraints:

- 1-2 day delays to detect model drift and performance degradation
- IT costs from compute-heavy tasks taking up over half of the business intelligence tool server
- Opportunity costs from the team spending time monitoring models instead of building new ones or automating other processes

"The strategic importance of ML observability is a lot like unit tests or application performance metrics or logging. If it isn't part of your workflow and automated, then it's just one more thing you have to do to get a model into production—another box to check rather than a strategic asset. We use Arize for observability in part because it allows for this automated setup, has a simple API, and a lightweight package that we are able to easily track into our model-serving API to monitor model performance over time. When investing a lot into a model and the productionization of that model, it's important to get that maximum insight into how it is performing for future versions. No one wants to be on a team that gets an email from a product owner that says 'are you guys sure this is working' because something broke or changed, often in a data stream that you don't directly control, clobbering model performance."


**Richard Woolston, Data Science Manager,
America First Credit Union**

Machine Learning Use Cases & Stack


Originally part of the business intelligence organization, the data science and machine learning team at America First Credit Union first took shape over five years ago after an early initiative—a member churn model built for the marketing team—proved successful.

Fast forward to today and America First Credit Union has over 30 models in production spanning the full member lifecycle across its \$17 billion asset portfolio, influencing everything from loan origination (i.e. auto, home) to ongoing credit risk, overdraft risk, credit card fraud prevention, member loyalty, product recommendations, and much more.

America First Credit Union: ML Stack for Data Scientists



In building its core stack and strategy, “we look at the capability maturity index—how smooth and how repeatable is something—and then we iterate on the next thing to make it better,” says Richard Woolston. “One of the core focuses initially was just on getting the software development lifecycle down—so getting the continuous integration and continuous deployment (CI/CD) humming—then standing up the feature store, data catalogs, and ultimately getting our in-house monitoring to run much faster.”



Challenges

For its first few models, America First Credit Union relied on ML monitoring built in-house using business intelligence dashboards. While the solution worked as designed, something more robust was needed to scale the organization’s ML practice.

Ultimately, America First Credit Union opted to partner with Arize as its ML observability partner in early 2021 to help remove a few key barriers to achieving its vision.

BI Tools Are Not Real-Time and Are Not Built for ML Use Cases

Business intelligence tools help measure known, defined metrics. Because ML models are built on data to act on data, relying on business intelligence tools for monitoring can lead blindspots when it comes to detecting unknown issues. Models

- Business counterparts sometimes catching model performance issues before the ML team
- An estimated 2–3 full time employees (FTEs) needed to build out and maintain a more robust ML observability solution in-house to meet future needs

Solution

After assessing whether to build or buy and undergoing a competitive review process of leading vendors in the space, America First Credit Union selected Arize for ML observability. Arize enables:

- Automated monitors based on customizable thresholds
- Fast implementation via an integration with Arize’s Python SDK
- Model versioning and environment comparisons
- Concept and feature drift monitoring and troubleshooting to compare across training, validation, and production environments
- Data integrity checks to ensure the quality of model data inputs and outputs with automated checks for missing, unexpected, or extreme values
- ML performance tracing to quickly pinpoint the source of model performance problems and map back to underlying data issues
- Bias and fairness tracing to ensure models are not generating potentially biased or unfair outcomes for protected segments of interest

routinely encounter unforeseen issues in the real world (i.e. training-serving skew) that can cause performance to fall below what is expected during training. Since BI tools are not tied directly into production systems, 24-48 hours of degraded model performance has the potential to impact business results. As a result, moving to automated and real-time monitoring and alerts is increasingly becoming table stakes.

Building More Robust ML Monitoring In-House Is Potentially Time-Consuming and Costly

Most technical teams have a natural tendency to build over buy when it comes to needed infrastructure. However, machine learning is multiple orders of magnitude more complex than data engineering or software engineering. Given the amount of logic in SQL procedures that feed into its models, building an ML observability solution in-house at America First Credit Union would likely mean dedicating multiple machine learning engineers for around a year of building on top of open source frameworks—a non-trivial and intense task with real opportunity costs in terms of forgoing other high-value projects.



AI Best Practices and Standards Are Evolving

In heavily regulated industries like financial services, compliance is paramount. A whole body of law and industry standards exist to ensure fair and transparent use of AI in underwriting, for example. For a mission-driven institution with nearly a century of built-in member trust like America First Credit Union, going far above and beyond what is required by regulations is also vital. With this in mind, the company is always on the lookout for better testing and validation to anticipate the needs of auditors and backstop robust first, second, and third-line defenses.

Results

Since deploying Arize, America First Credit Union is seeing improved model performance and time savings. That translates to:

- A payback period of under two months; estimated >500% ROI in first year
- ~2-3 FTEs saved in year one by not needing to build and maintain an in-house ML observability solution to meet future needs
- Hundreds of team hours gained annually from Arize enabling greater automation and fitting into America First Credit Union's CI/CD pipeline
- Improved model performance from proactively surfacing feature drift impact and performance impact at a cohort-level
- Reduced onboarding and communication costs with product owners and auditors from having standardized monitoring in place
- Freed up to focus on the important aspects of ML observability: poor-performing segments and fairness, both of which are critical in financial services
- Significant savings to the business monthly from AI systems performing as designed

"Machine learning is a discipline where few notice when everything is performing perfectly—and everyone notices when things go wrong. In that sense, it's not a question of whether you need ML observability—you do—it's more a matter of whether to build or buy. For us, Arize was the clear choice in terms of cost efficiency and freeing us up to achieve our broader vision."

**Richard Woolston, Data Science Manager,
America First Credit Union**



Solution

An early innovator and leader in ML observability, Arize AI processes hundreds of billions of predictions a month on behalf of clients. By connecting offline training and validation datasets to online production data in a central inference store, Arize's ML observability platform helps ML teams streamline model performance management, drift detection, data quality checks, and model validation.

America First Credit Union first selected Arize in early 2021 after a competitive review of monitoring platforms. Key factors influencing the decision include Arize's ease of implementation, ability to handle enterprise scale, robust capabilities to surface drift and performance issues at a cohort level, and Arize's shared technical philosophy.

Results

With robust ML observability from Arize now part of its ML stack, America First Credit Union's ML team is able to keep its focus on building new models and helping product teams achieve business goals. The organization sees a few key benefits from the Arize platform.

Model Drift and Performance Degradation Are Now Caught Immediately Instead of Taking 1-2 Days

Lagged reporting in America First Credit Union's business intelligence tool is now supplanted by real-time alerts and monitors from Arize. When performance falls below a set baseline, a monitor fires—kicking off an automated process where a Jira ticket is created and an on-call team member is then paged to investigate, creating an artifact trail. "It's nice because we are not working from an email, but rather from the tools we are already using," says Richard Woolston.

Improved Model Performance From Faster Root Cause Analysis

Once alerted to an issue, America First Credit Union's ML team can then leverage Arize to help figure out why it is happening—filtering through the cohorts where drift impact and performance impact is highest and adjusting accordingly. Instead of time-consuming SQL querying and writing ad-hoc scripts to slice the data and then visualize it, insights are a few clicks away. "We're catching issues earlier with Arize, and that makes a big difference," notes Woolston.

Time Savings Leads To More Models In Production

By not needing to dedicate several machine learning engineers to the task of building a more robust ML monitoring solution to support real-time models or other use cases, the team can now focus more exclusively on solving other infrastructure challenges or getting new models from research to production. In all, the team is now deploying two new models per month and is applying machine learning in new and exciting ways across the business.



Proactive Drift Detection and Troubleshooting

When it comes to monitoring and troubleshooting lending and lifetime value models, navigating delayed actuals is always a challenge. America First Credit Union leverages Arize to monitor drift as a useful proxy metric. “Drift is also one of the first things that we look at when first deploying a model as predictions start to come across over the hours or days to know that things are stable,” notes Richard Woolston. “If not, we can go ahead and roll back and identify what pieces were missing and fix the offline or online feature store and iterate.”

Collaboration Across Teams Is Easier

The data science and machine learning team at America First Credit Union supports multiple internal stakeholders, including product teams who share responsibility for the results of models. Many resemble citizen data scientists in the sense that they are partnering with the ML team, leveraging deep subject matter expertise in areas like credit risk or lending to find the right questions and train and troubleshoot bespoke models. With Arize, the team can more easily share performance data and alerts across teams.



Additional Safeguards Help Ensure AI Fairness and Compliance

America First Credit Union is among the earliest clients of Arize to use the platform to monitor model fairness, automatically adding an Equality Credit Opportunity Act (ECOA) feature group to every prediction without actually using these features in the model. Arize Bias Tracing, a new tool that is purpose-built for measuring whether a model is generating potentially biased outcomes for segments of interest using prevailing fairness metrics, is also now available to supplement these efforts—as are robust explainability tools to better understand why a model made a particular decision.

Conclusion

America First Credit Union has come a long way since its founding in 1939, but one thing has not changed: a steadfast dedication to its member-owners. The organization’s early adoption of artificial intelligence, aided by an ML stack that includes Arize, promises to help further this mission well into the future.

[Sign up for a free account](#) today to start your ML observability journey.

