

Clearcover Accelerates Model Velocity, Navigates Feature Drift and Drives Confidence In Deployed Models

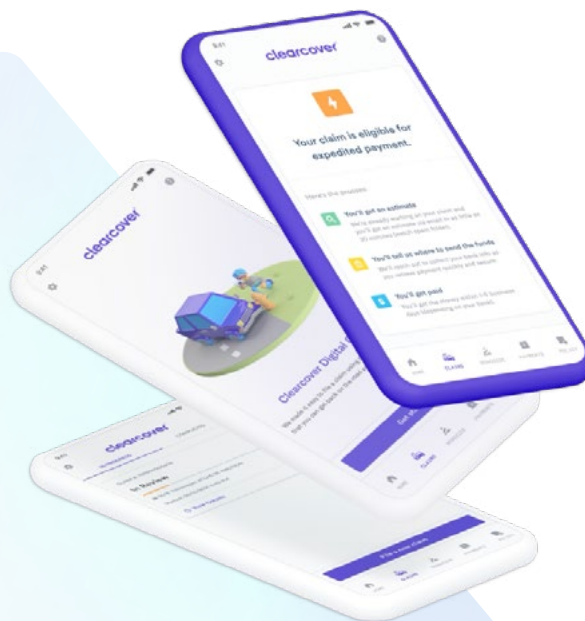
The machine learning (ML) team behind Clearcover Insurance Company's award-winning app and the fastest claims in auto insurance moves to real-time models after implementing Arize's ML observability platform

Introduction

Clearcover Insurance Company is a tech-driven car insurance company that promises exceptional coverage at a great price, empowering customers to make the smartest decisions at every step. Through its simple-to-use mobile app, Clearcover boasts the industry's fastest claims, easy payments, and convenient policy management – a differentiated experience for modern policyholders.

More than 65% of consumers who adopted digital claims handling practices during the pandemic intend to continue those habits post-COVID, according to a McKinsey & Co. survey. Clearcover is capitalizing on this trend through updates like Clear Claims™, an advancement in the digital claims process that provides payment on eligible claims within 30 minutes or less.

Behind this track record of innovation is an [expanding](#) AI team of 10 data scientists and machine learning engineers building and deploying ML models to help power the full customer lifecycle at Clearcover. To help realize its vision for real-time ML at scale, Clearcover selected Arize after a competitive proof of concept in early 2021.



Key Facts

clearcover

Industry: Insurance

About: Tech-driven car insurance company

Models In Production: 10–15

Primary Use Cases: customer churn modeling, customer lifetime value modeling, marketing optimization, potential claims fraud identification

Challenges

The complexity and time required to set up in-house model monitoring emerged as a likely bottleneck in achieving AI goals, necessitating an early (and prescient) shift in strategy to overcome the following constraints:

- >2 weeks to implement new model monitoring dashboards in a business intelligence (BI) tool for every new model
- ~24 hour delays in detecting performance degradation due to BI dashboard not being real-time
- 9–10 business weeks per year for data scientists to proactively monitor models due to lack of real-time alerts
- Lack of monitoring for data drift on models scoring all eligible claims

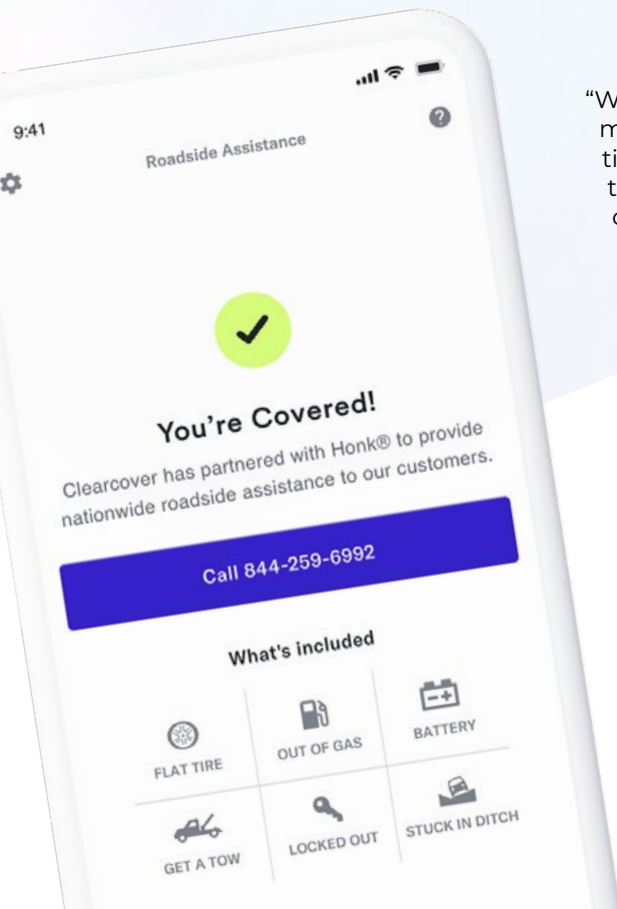
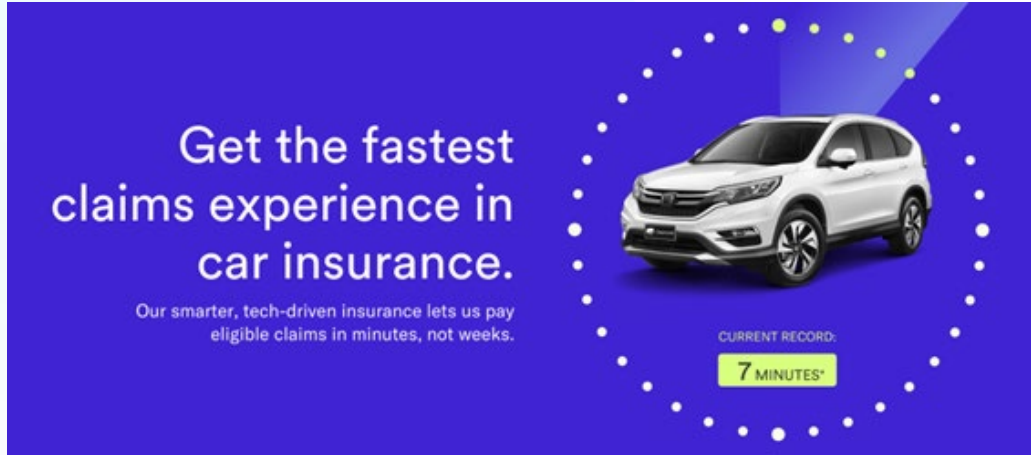
“We recently deployed a model that went from inception to production in 46 days – hardly a small endeavor given the model is relied on to score over 50,000 insurance applications daily. Arize is a big part of that success because we can spend our time building and deploying models instead of worrying – at the end of the day, we know that we are going to have confidence when the model goes live and that we can quickly address any issues that may arise.”

Alex Post, Lead Machine Learning Engineer, Clearcover

Machine Learning Use Cases & Stack

At Clearcover, machine learning models touch nearly every part of the customer journey. This includes:

- Marketing: models score potential customers getting auto insurance quotes on price comparison websites and recommend the next best marketing action
- Customer Retention: customer value estimates and churn predictions are made prior to policy bind and help Clearcover's team better understand both current and prospective customers
- Customer Experience: when a customer submits a claim via Clearcover's mobile app, models predict the risk for potential fraud



"We've known from the beginning that our most meaningful models would be real-time models, influencing the customer as they interact with us," notes Alex Post. "In order to do that, we needed a strong infrastructure to allow us to deploy real time models, access the data we need in real time, and then – more challengingly – align our offline training data with the data that is available in real time."

To that end, Clearcover's stack includes FastAPI endpoints that allow for models to be plugged in through a model registration process; Kubernetes and Amazon S3 for model storage and model registration; Kafka for real-time streaming of results to keep track of predictions and enable offline analysis in a warehouse; and Arize AI for end-to-end ML observability.

Solution

In order to have an efficient training-to-production flow and a foundational ML stack in place to enable real time models, Clearcover implemented Arize for ML observability. Arize enables:

- Automated monitors based on predefined thresholds
- 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

Results

Since deploying Arize AI, Clearcover is achieving improved model performance and significant time savings as AI increasingly contributes to the company's bottom line. That translates to:

- >400 extra hours freed up per year across the ML team
- 10% more models deployed into production per year
- A payback period of under nine months; >150% ROI in first year
- Improved model performance from proactively surfacing feature drift and performance impact score at a cohort-level
- Automated and instantaneous monitoring via a simple integration into Arize's Python SDK, compared to 2–4 weeks to set up monitoring before implementation
- Alerts sent to Slack and email immediately, maximizing internal visibility

Challenges

Of course, this technology stack did not exist on day one – and the company’s success in AI was hardly a given at the outset. “There were a lot of eyes on us as we prepared to deploy our first model,” recalls Alex Post. “We needed to not only show stakeholders it was working but also stay a step ahead if something were to go wrong in production.”

Like many teams, Clearcover initially planned to leverage a business intelligence (BI) tool for ongoing ML monitoring. The rationale: the business intelligence and data engineering teams already leveraged this tool, and adding dashboards for machine learning to monitor results would be straightforward.

Ultimately, Clearcover opted to implement Arize. Several key reasons drove the decision.

BI Tools Are Often Bulky In Practice

At most companies, BI tools are managed by business intelligence or data teams that have expertise in managing data models, preparing the data, and setting up dashboards. Even under the best of circumstances, setting up dashboards is often a time-consuming task. In Clearcover’s case, the ML team faced delays of at least one sprint (around two weeks) to set up monitoring for each new model in coordination with other teams and stakeholders.

BI Tools Are Not Real-Time and Lack Timely Alerts

One of the first models Clearcover deployed into production was designed to predict the retention of a customer during the quoting process. Since such a prediction would be of little value to the business a day or more later, deploying a real-time model was essential.

BI Tools Are Not Built for ML Use Cases

For Clearcover, building out monitoring for feature drift in particular stood out as a dreaded, time-intensive task. For data scientists, the prospect of coming up with logic and then hard coding it for a seemingly endless array of feature-value combinations was an unwelcome distraction from training new models and exploring potential ways AI can solve business problems.

“From the start, Clearcover focused on making AI effectively scale. A lot of places might start with a data scientist, but unless you have someone with engineering skills they may struggle to get things in production. Clearcover took the opposite approach of starting by building the ML capability first to ensure that when dedicated data scientists joined they would be able to provide value right off the block. That required a lot of foresight.”

Carolyn Olsen, Senior Manager of Data Science, Clearcover

Solution

As an early pioneer and leader in machine learning observability, Arize AI tracks 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.

Clearcover selected Arize in early 2021 after a proof of concept as the insurer prepared to deploy its first models into production. Several factors influenced the decision, including Arize’s technical architecture and ease of setup, and the platform’s ability to proactively surface feature drift and performance issues at a cohort-level, and real-time alerts.



Results

With end-to-end ML observability from Arize as a key part of its ML stack, Clearcover now deploys models with confidence. The team sees several key benefits from the Arize platform.

Monitoring Setup Is Now Automated and Instantaneous Instead of Taking 2–4 Weeks Per Model

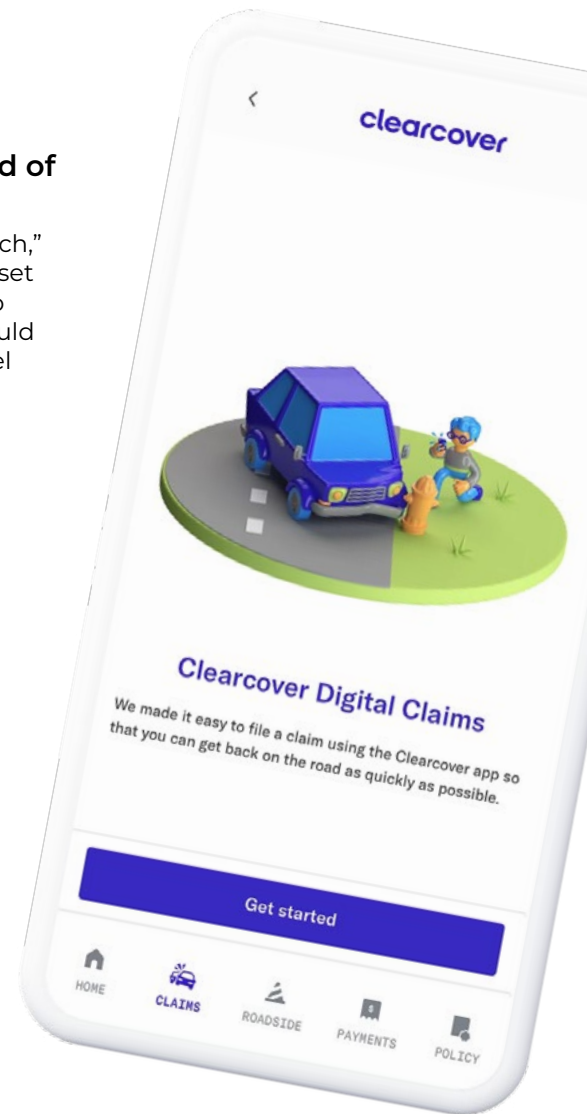
“As an ML engineer, I appreciate that I don’t have to interact with Arize very much,” says Alex Post. “The nature of using Arize’s Python SDK and the way that we’ve set up our infrastructure is that when we need to deploy a model, we don’t need to set up monitoring – it just happens – which is huge for us. It takes time that would otherwise be spent doing repetitive tasks every time we want to deploy a model and allows us to focus on other things that deliver value to the business.”

The ML Team Is Now Deploying 10% More Models Per Year

By not having to explicitly set aside time to proactively check in on model performance or drift, Clearcover’s ML team estimates that it gains 400 hours annually – time that is ultimately put to good use accelerating model development to the tune of 10% more models deployed into production per year. “Knowing that real-time alerts will go out – whether it’s an email or Slack – based on predefined thresholds offers a lot of peace of mind and frees up a lot of time to help the business,” Carolyn Olsen says.

Proactive Drift Detection and Troubleshooting

Clearcover also leverages Arize to monitor and troubleshoot feature and prediction drift, which is critical for certain use cases. For example, the data from insurance price comparison tools often changes as questions or categories get updated or the composition of site visitors shifts. In such cases, timely alerts from Arize on feature and prediction distribution changes combined with visualizations to aid in figuring out underlying causes make a big difference. Clearcover’s data scientists can then retrain a model or revert to a prior model version, minimizing periods of degraded performance.



Conclusion

With an inspired underdog spirit and an impressive track record of innovation, Clearcover is taking the insurance industry by storm. The company’s machine learning team, aided by a stack that includes Arize AI for the critical task of end-to-end ML observability, is a key part of how the company is achieving its mission.

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