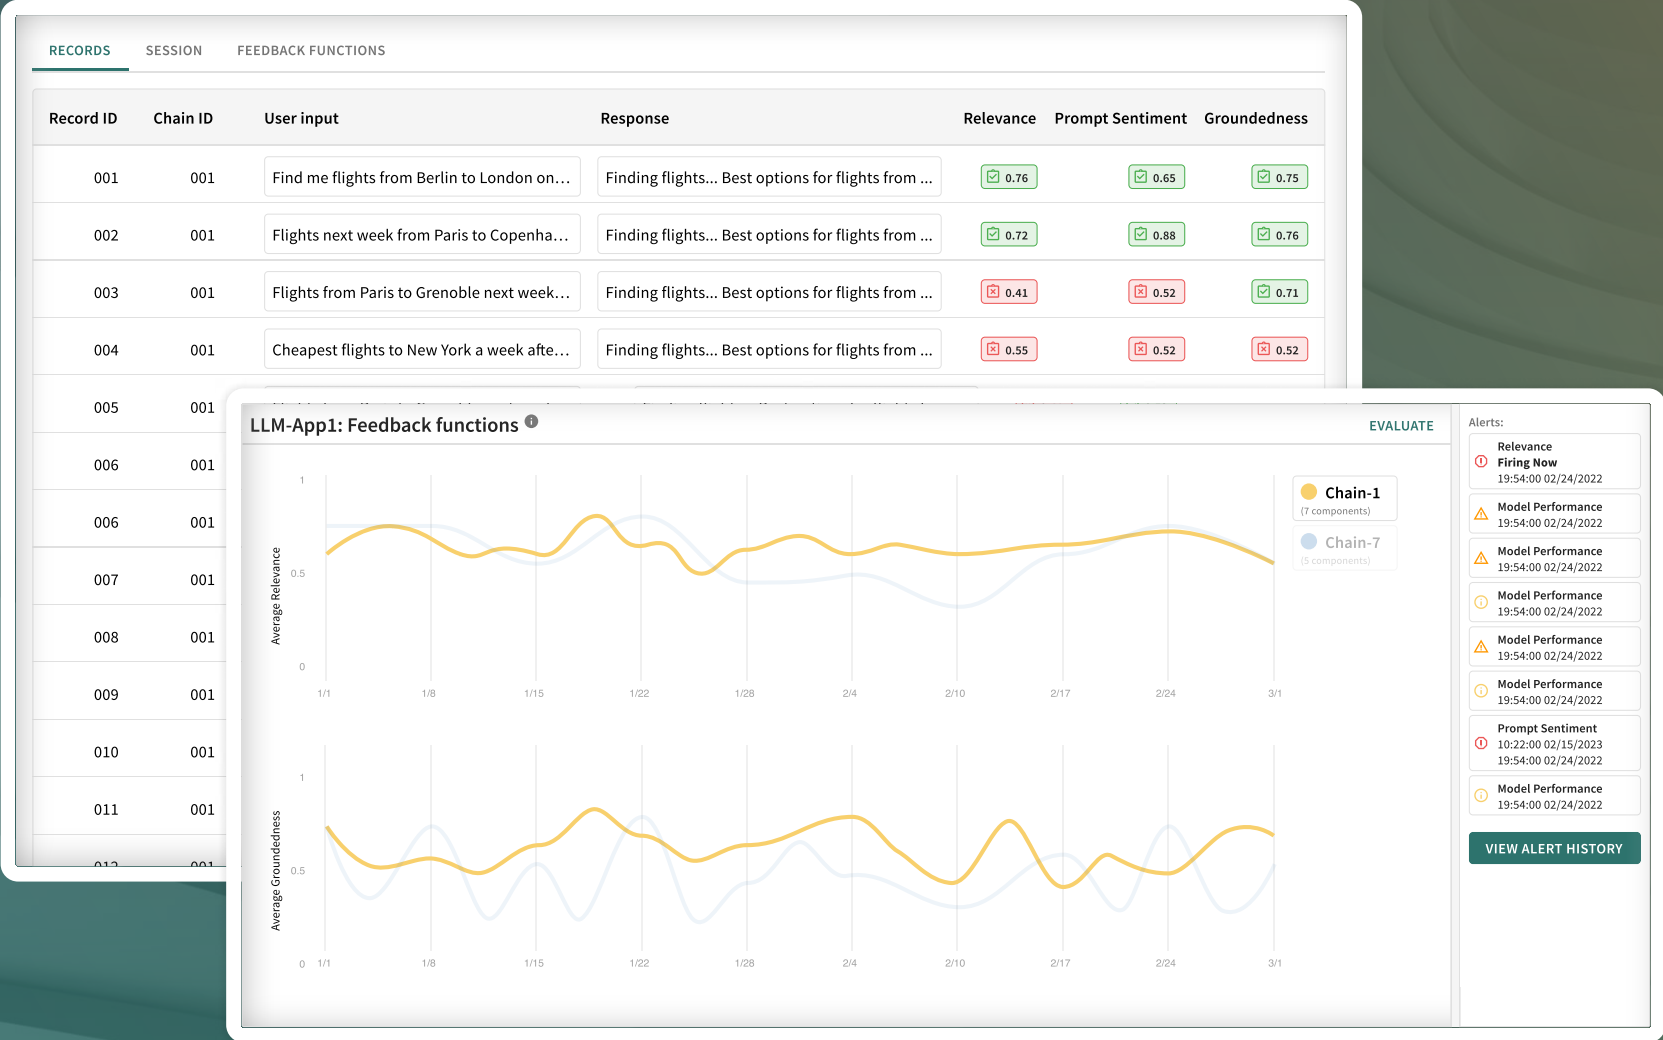


TruEra AI Observability

Cloud-first, comprehensive AI Observability combining monitoring, debugging, testing, explainability, and responsible AI. For predictive and generative AI.

Achieving high performance, trustworthy AI is the next great challenge in artificial intelligence and machine learning. Today, only a minority of AI projects ever make it to production, representing a huge loss in potential value and resources spent. Regulatory pressure on AI quality is increasing in the US, Europe, and Asia. High profile missteps with AI have damaged the reputations of some of the most well known companies.

TruEra provides a comprehensive AI Observability solution that helps enterprises to drive AI quality and get into production, satisfy internal and external stakeholders, and protect corporate reputations. TruEra solutions are based on technology developed over more than nine years of research. They help ensure ongoing model performance and business results throughout the AI lifecycle and are used by leading organizations around the world.



Recognized Leader in AI Observability



As AES expands the coverage and use of our AI models, it's critical that we put in place MLOps tools to efficiently monitor, test, and debug our AI at scale. TruEra's powerful monitoring and testing capabilities enable us to increase the number and types of AI models that we use.

- Sean Otto, Director of Analytics, AES

TruEra: driving AI performance and quality across the full AI lifecycle.

Testing and Evaluation - Generative AI

Thoroughly test and debug Gen AI apps from development to production, to minimize risks like accuracy and hallucinations.

TruEra's Gen AI testing and evaluation capabilities include:

- **Ability to evaluate a broad range of LLM apps.** Evaluate a wide range of Gen AI apps, including retrieval augmented generation (RAGs) and agents. These can be built on popular frameworks such as LangChain or LlamaIndex, or other custom frameworks.
- **Wide range of evaluations.** With TruEra feedback functions, you can measure a wide range of automated, human and traditional metrics on your LLMs. You can also track your own metrics at scale. Evaluate:
 - Relevance
 - Groundedness
 - Custom evaluations
 - Prompt sentiment
 - Language mismatch
 - Transcript length
 - Response verbosity
 - Fairness substitution
 - Toxicity
 - ... + Manual feedback
- **Comparisons and A/B testing.** Compare multiple approaches at the same time to understand tradeoffs between cost, latency and performance.

Testing and Evaluation - Predictive AI

Comprehensive model evaluation and testing technology that promotes quality and transparency, while accelerating model development.

TruEra's testing and evaluation capabilities include:

- **Automated Test Harness** to perform systematic testing, with predefined thresholds.
- **Robust set of analytics** highlighting performance, stability, conceptual soundness, bias, and data quality.
- **Champion/challenger analysis** compares multiple models across a single project to find the top performer.
- **Model explainability** provides a view into the fundamental underpinnings of the model.
- **Feature influence analysis** that shows which inputs are the key model drivers, identifies outliers, and tracks feature trends.
- **Segment analytics** to understand segment level accuracy, see predicted behavior by segment, or see feature importance by segment, and identify areas where the model is error-prone.
- **Fairness and bias analyses** to understand how your model is behaving towards different populations, including protected populations.

Key Benefits



Achieve model transparency

Explain fully how your model works and its key drivers.



Drive model quality

Analyze accuracy, reliability, stability, conceptual soundness, and bias



Build trust, reduce risk

Easily demonstrate to stakeholders, customers, and regulators the effectiveness, quality, and fairness of your models



Get into production faster

Accelerate model development, validation, and approvals; easily report and collaborate



Get the broadest and deepest view of real-world performance

Get a full spectrum view into your models, with analytics that you can't find anywhere else.



Enable fast action - know where to look and what to fix

With more reliable analyses of root causes, your data science and MLOps teams stay focused.

Get started quickly. Monitor your apps and models at scale.

TruEra Monitoring

The machine learning monitoring solution that gets down to the root cause, so that your models stay up and running. Get robust model supervision, views into model performance unavailable anywhere else, and configurable alerting that ensures you're on top of the most critical risks.

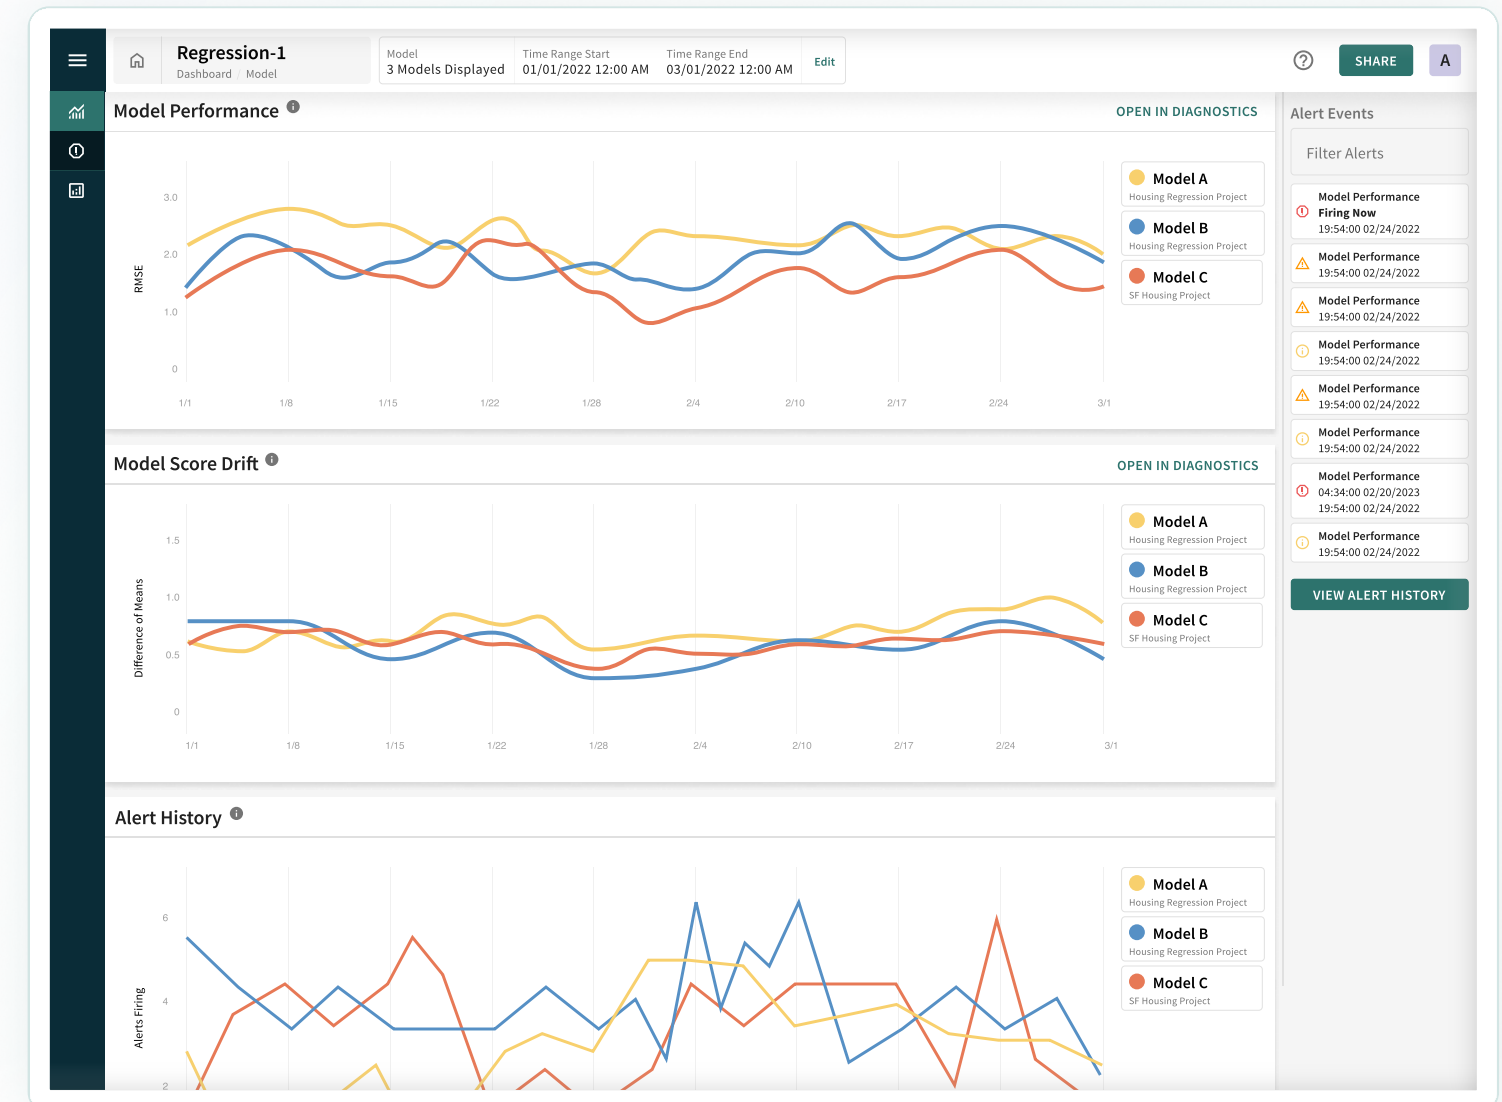
TruEra's predictive monitoring capabilities feature:

- **Performance visibility** with deep model quality dashboards and an overwatch dashboard for supervising 100s of models.
- **Bias tracking** to understand if some groups are being treated differently than others, and why.
- **Root cause analysis** so that you not only know what issues are emerging, but why, so that you can fix them quickly and effectively.
- **Alerts and notifications** for any dashboard metric you choose.
Configure alerts to ensure you're notified for the most important issues, across any of 19 different communication channels.

TruEra is an easy addition to your existing AI stack.

TruEra is quick and easy to deploy, even in complex enterprise environments.

- **Easy, fast SaaS deployments.** TruEra is SOC 2 Type II certified.
- **Deploy on your cloud**, including public cloud providers like AWS.
- **Fast setup and easy user administration** with SSO integration and local TruEra authentication.
- **Easy integration** to data sources, common model development solutions, such as AWS Sagemaker, Microsoft Azure, Vertex.ai, Dataiku, Data Robot, or H2O.ai, and model serving solutions
- **Scale quickly** to analyze high volumes of machine learning models and their data
- **Easily export data via APIs** for integrations to applications or business intelligence solutions, such as Tableau and Looker



Effective monitoring of machine learning is one of the major missing pieces of IT management.

We are excited about a solution that goes beyond current ad hoc tools and allows us to not only quickly and effectively identify emerging issues, but also debug and resolve them.

- Matthew Bennett, CIO, North America,
HarperCollins