CME Group AI in Action: US Treasury TCA Analytics

October 2023

CME Group

James Boudreault, CFA
Jason Chauval
George Deamont

Google

Mark Scannell

Agenda

1   What is Transaction Cost Analytics (TCA)

2   AI enabled Treasury TCA Toolset

3   Interactive Survey

4   Q & A

Key Speakers

James Boudreault, CFA
Global Head of Data Science at CME Group

George Deamont
Manager of Data Science at CME Group

Jason Chauval
Executive Director of Analytic Solutions at CME Group

Mark Scannell
Financial Services Solution Architect at Google

CME Group Transaction Cost Analytics Development Program

CME Group is shortly launching a new analytics service via Google GCP starting with a comprehensive suite of client TCA **covering all US Treasury trades, Brokertec CLOB** and **Brokertec Stream** and across our **major UST futures markets**, followed soon by FX Futures, Equity, Major commodity markets and EBS cash FX.

The service will include
- **Full Order and Trade level mark to market calculations** and **market impact by time and order book update** count
- **Parent level aggregation**: Parent order summations by same trader ID, same Firm, Session ID in same direction within continuous time frame of 5 mins
- **Peer group benchmarking** including MTM of Futures/Cash equivalents on each transaction (including Basis Spreads cash-futures)
  - Peer grouping by Order Fill Ratios, and Market Maker classifications,
  - Peer group rankings using explainable machine learning attribution models powered by Google Vertex Ai.
- **Best Execution Total Cost analysis** using aggregate methods with VWAP, TWAP, Optimal Execution calculations.

External client access will come in 2024.
- Clients will eventually be able to access their data via CME Analytics Hub, Looker Dashboards and leverage Google Vertex Ai/Big Query tech stacks to perform more advanced analysis as they need.
- Clients using other cloud platforms, or their own databases can simply import the data table views from GCP to their own structures
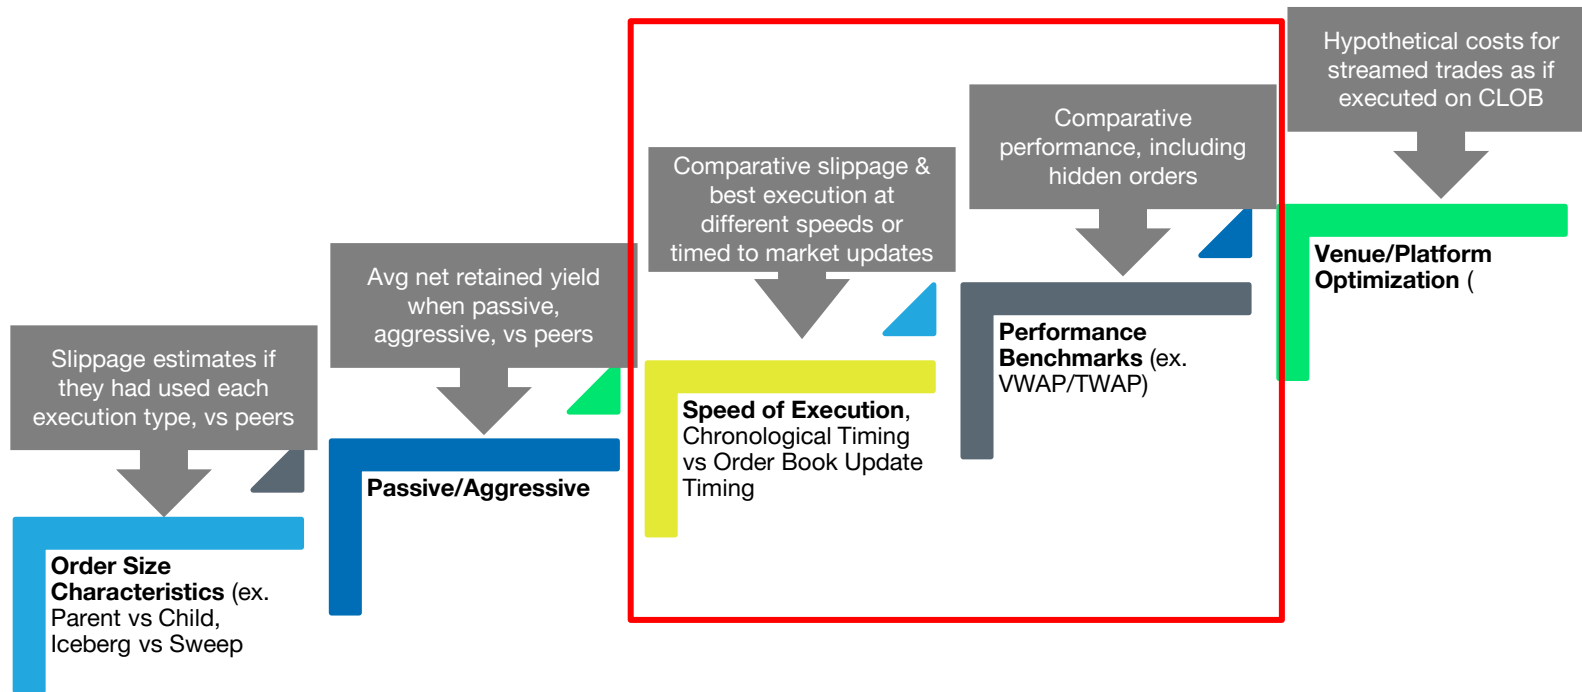
The google analytics platform will **expand in time to other products**, and more complex analysis.

What is TCA? (Transaction cost analysis)

TCA stands for Transaction Cost Analysis, which is a method used to measure and evaluate costs associated with executing financial transactions
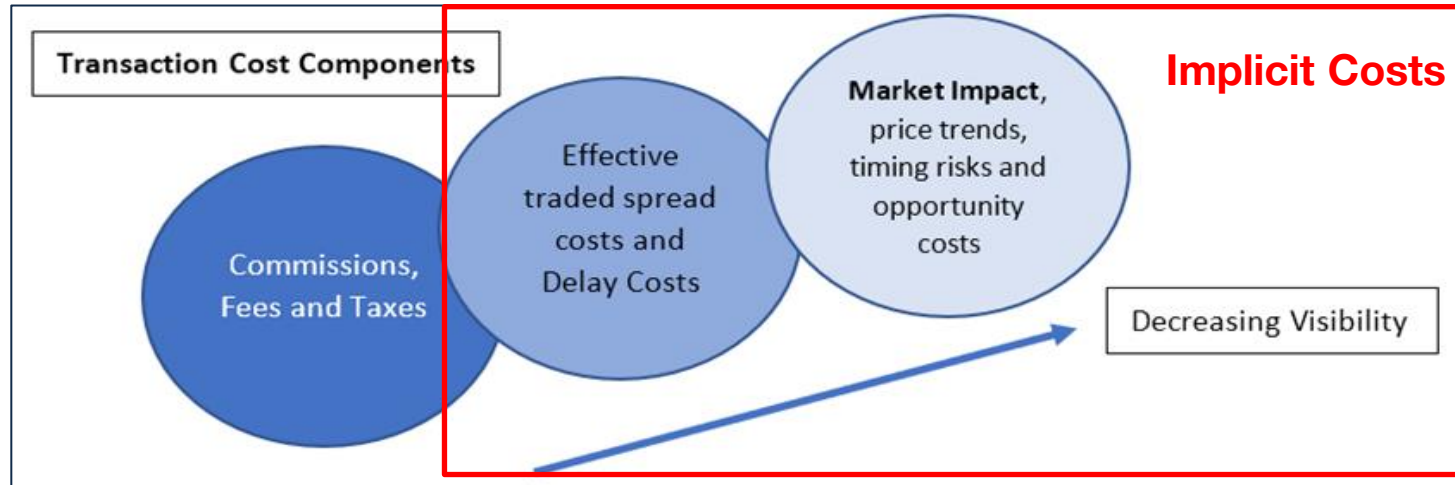
Intelligence at every step of the trade allows clients to find the optimal trade modality to meet their needs

Slippage estimates if they had used each execution type, vs peers

Avg net retained yield when passive, aggressive, vs peers

Comparative slippage & best execution at different speeds or timed to market updates

Comparative performance, including hidden orders

Hypothetical costs for streamed trades as if executed on CLOB

Order Size Characteristics (ex. Parent vs Child, Iceberg vs Sweep

Passive/Aggressive

Speed of Execution, Chronological Timing vs Order Book Update Timing

Performance Benchmarks (ex. VWAP/TWAP)

Venue/Platform Optimization (

**Evaluation of prior trades provides intel for clients for the *how to trade* and *where to trade* decision**

AI - Treasury TCA on GCP

Leverage the latest technologies in **automation and AI** to establish a suite of analytics and models around helping customers better understand and **enhance trade execution activities**. Build a set of foundation models for **peer group benchmarking** and **slippage rate reduction**, providing explainable results and helping customers **capture unrealized value.**

*slippage rate = the difference between the expected price of an order and the price when the order actually executes

Sample of AI and ML Achievements

Data Science Created the First Neural Network at CME Group
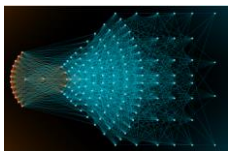
1. **Neural Network**: An ML algorithm modeled after the structure of the human brain. Using interconnected nodes or "neurons" to process and analyze data, making predictions or classifications based on identifiable patterns. NN "learns" by adjusting the strength of connections between neurons through training on a dataset, improving its accuracy over time. Neural networks are commonly used in image recognition, natural language processing, and other complex tasks.

   · To aid Peer Group Benchmarking and trade execution Slippage Rate. Optimizes the strength of connections through training on a data set; allows for improved accuracy over time.

2. **XGBoost:** A ML framework used for gradient boosting in model development. Known for its speed, accuracy, and ability to handle large data datasets. Uses decision trees to build predictive models; whilst training weak models and iterating to improve via correcting the errors of previous models.

   · Used XG Boost to train the Slippage Rate Reduction Model for Treasury TCA Tool. Produces the best results alongside our Auto ML model, resulting in enhanced tool performance and customer outcomes.

3. **Diverse Counterfactual Explanations (DiCE):** A Machine Learning algorithm that provides explanations for a given prediction, helping users understand why a model made a prediction and how the prediction could differ under alternate circumstances.
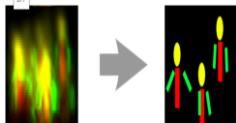
   · Created DiCE ML model for explaining and measuring Slippage Rate. Helps users understand how to lower slippage rates.
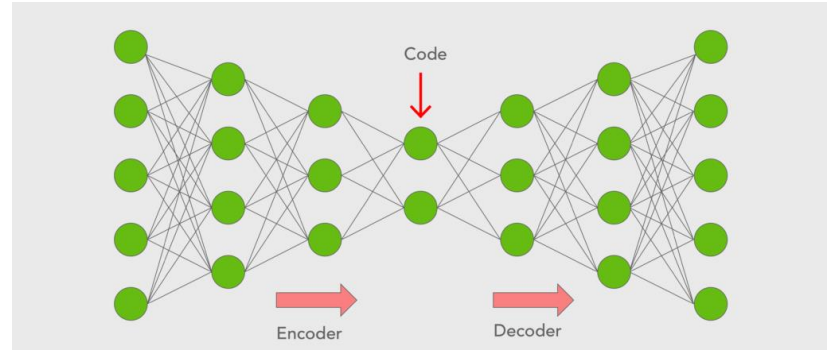
Autoencoder Neural Network

Neural network model trained to learn efficient encoding of data

Capable of capturing non-linear behavior of data

Useful technique for dimensionality reduction, and more robust compared to other methods like PCA

Peer Group Benchmarking

K-Means Clustering
- Convert participant summary data to 2-D embeddings using Autoencoder
- Cluster participants using K-Means algorithm
- Evaluate model using silhouette scores
- Match peer groups over time using jaccard similarity index

Additional Methods
- Gaussian Mixture Models
- DBSCAN
- PCA/UMAP/t-SNE

Clusters in 2-Dimensional Space
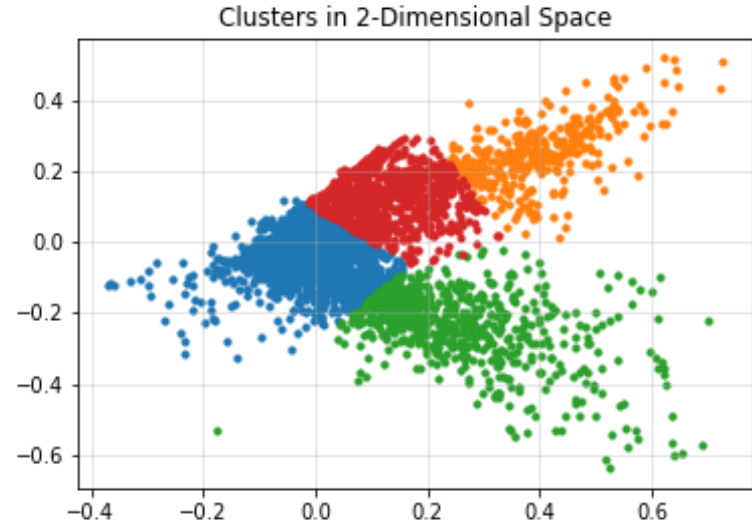
Explainable AI

XGBoost Classifier
- Model trained to classify participants into peer groups using original trading data.
- Needed because the clustering model cannot be explained using the 2-D embeddings.

SHAP (SHapley Additive exPlanations)
- Game theoretic approach to explain the output of any machine learning model.
- Used to determine the most impactful features at the peer group level, and compare peer group mean/standard deviation with population statistics.

LIME (Local Interpretable Model-agnostic Explanations)
- Learns a local linear model around an individual observation that requires explanation.
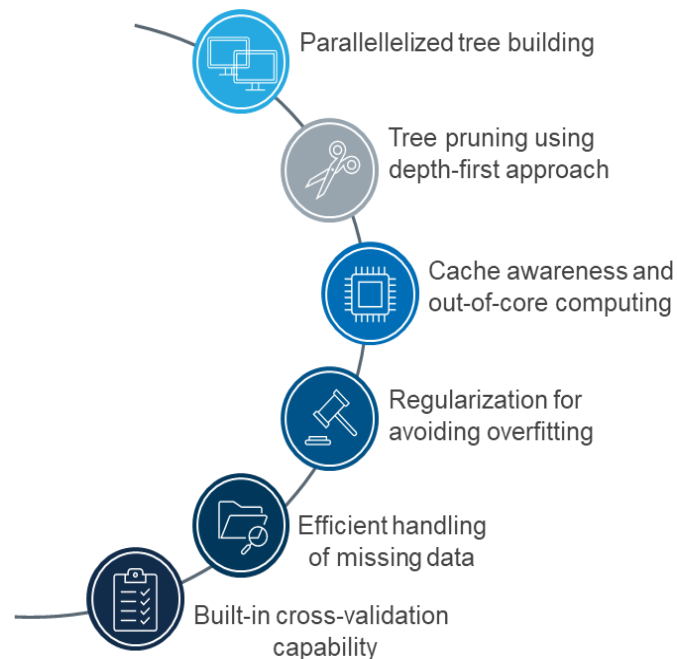- Used to explain individual predictions.
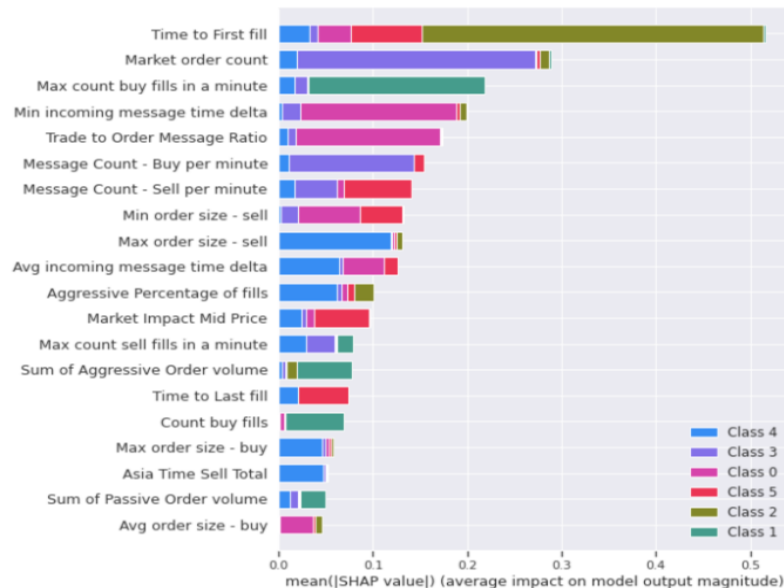
Slippage Rate Reduction

XGBoost Regression Model

- Used XGBoost to train a regression model to predict slippage rate using trade parent data.
- Main Features: number of orders, quantity, duration, passive/aggressive, volatility.

DiCE (Diverse Counterfactual Explanations)

- Explainable AI using counterfactual explanations instead of feature importance.
- Helps the client understand how they can change their trading behavior in order to lower their slippage cost.

Parallellelized tree building

Tree pruning using depth-first approach

Cache awareness and out-of-core computing

Regularization for avoiding overfitting

Efficient handling of missing data

Built-in cross-validation capability

Explainable AI Using Diverse Counterfactual Explanations (DiCE)

Counterfactual Examples

ML model's decision boundary

**Original class:
Loan rejected**

**Desired class:
Loan approved**

Original input
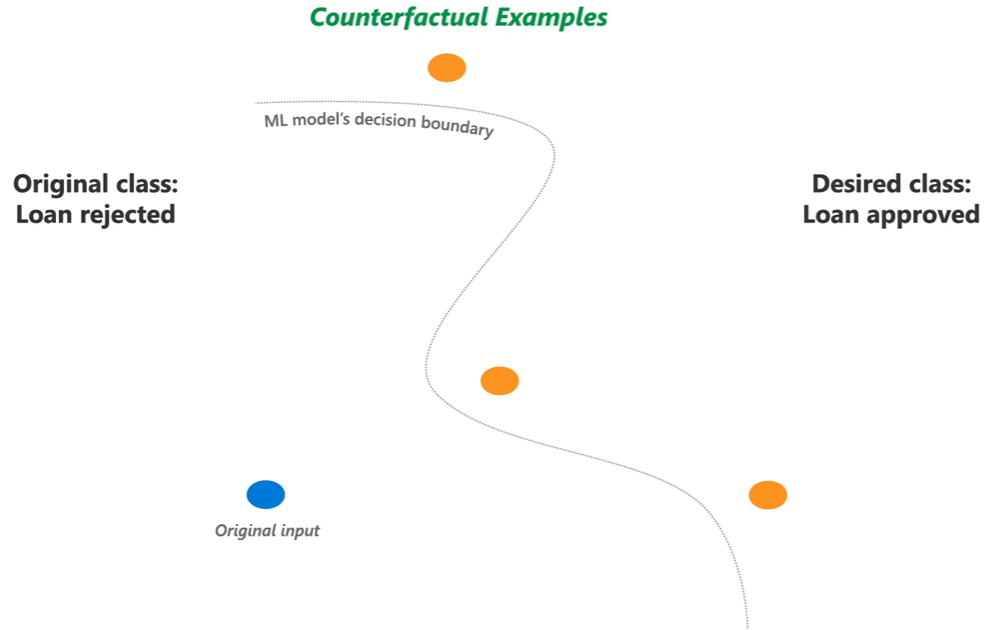
CI/CD High Level Flow
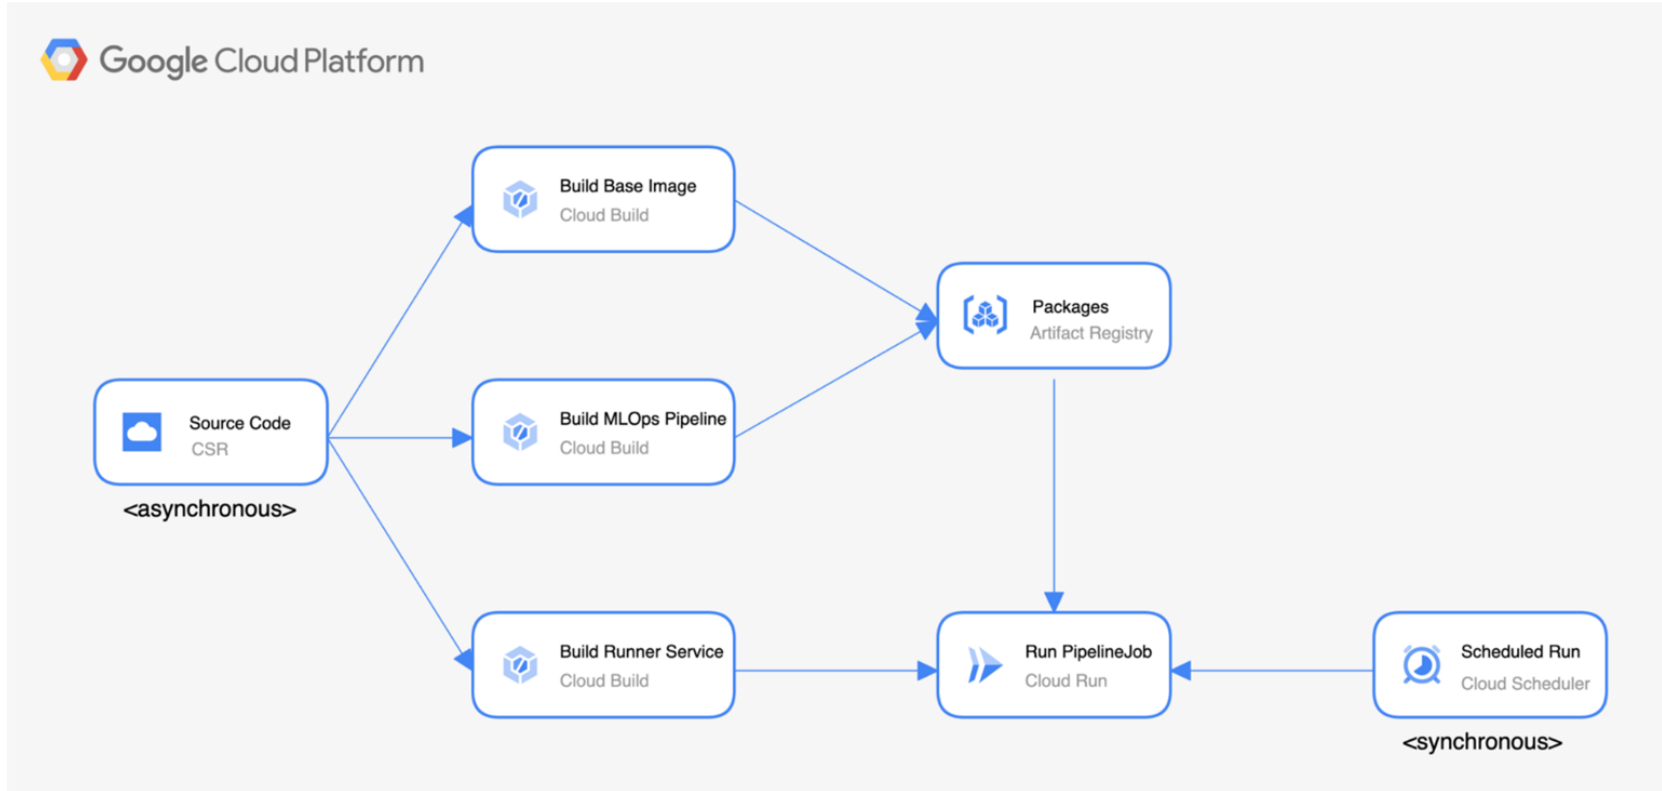
Full MLOps Pipeline

CI/CD and Orchestration

Source Code — CSR → Build Pipeline — Cloud Build Triggers → Run Tests — Cloud Build → Packages — Artifact Registry → Run PipelineJob — Vertex AI Pipelines ← Scheduled Run — Cloud Scheduler

<asynchronous>   <synchronous>

Data Preparation

Raw Data — BigQuery → Data Preprocessing → Preprocessed Data → Ingest Data — Cloud Storage → Vertex AI Featurestore → Serve Data — BigQuery

Embeddings Model

Hyperparameter Tuning —Yes→ HP Tuning — Vertex AI Vizier → Vertex ML Metadata

No → Vertex AI Training → Model Selection — Vertex AI

Vertex AI Training → Vertex AI Model Registry → Batch Prediction — Vertex AI → Embeddings Data

Vertex AI Tensorboard

Ingest Data — Cloud Storage → Vertex AI Featurestore → Serve Data — BigQuery

Clustering Model

Hyperparameter Tuning —Yes→ HP Tuning — Vertex AI Vizier → Vertex ML Metadata

No → Vertex AI Training → Model Selection — Vertex AI

Vertex AI Model Registry → Batch Prediction — Vertex AI → Peer Group Assignments

Vertex AI Tensorboard

Ingest Data — Cloud Storage → Vertex AI Featurestore → Serve Data — BigQuery

Explainability Model

Hyperparameter Tuning —Yes→ HP Tuning — Vertex AI Vizier → Vertex ML Metadata

No → Vertex AI Training → Model Selection — Vertex AI

Vertex AI Model Registry → Batch Prediction — Vertex AI → Shapley Values

Vertex AI Tensorboard

Explainable Result — BigQuery → XAI Endpoint — Cloud Endpoints

Go to
www.menti.com
Enter the code
9212 6709

Q&A

Thank you

Appendix

Speaker Bio

James Boudreault, CFA

James Boudreault, CFA serves as Global Head of Data Science at CME Group. He leads the department that provides analytics relating to customer, product, and industry landscapes in which CME Group operates. He is responsible for data science and developing client-facing tools. He has expertise in data, analytics, cloud technologies, leadership, innovation, and economic thought leadership.

Boudreault previously held roles at CME Group in the Products & Services, Research & Product Development, and Business Development areas. Prior to joining CME Group, he held various positions at Wellington Management and State Street within their portfolio management, trading, and business development teams.

Boudreault holds a bachelor's degree in finance from the University of Massachusetts at Amherst. He earned the Chartered Financial Analyst® designation, is a board member of the CFA Society Chicago Board of Directors and served on multiple advisory groups for the Society including co-chair of the finance committee. He is a board member of the Corporate Advisory Board of Southern Illinois University's Pontikes Center for Advanced Analytics and Artificial Intelligence. He is a retired captain in the United States Army and is an Iraq War veteran.

CME Group

Speaker Bio

Jason Chauval
Jason is Executive Director in charge of Analytic Solutions for CME Group Data services division.

Jason has over 30 years experience in the financial markets covering FX and Metals Options Market making, CTA managed Futures, e-FX algorithmic trading and and risk management as well as advanced transactional analytics.

Jason led the development of kdb+ based analytics at EBS over the last 9 years and is now responsible for developing expanded transactional analytics across all CME Products as well as other client facing areas such as Clearing, Risk and margin analytics and expanding our quant analytic libraries to help clients achieve best execution needs.

Jason is a keen base jumper and free diver and enjoys pressing wildflowers in the evening as well as knitting and looking after stray cats.

CME Group

Speaker Bio

George Deamont

George Deamont is a Manager of Data Science within the AI/ML workstream at CME Group. He is a dedicated professional with nearly 10 years of experience in the field of software engineering and distributed computing. Equipped with a Bachelor's degree in Engineering Physics and a minor in Computer Science from the University of Illinois at Champaign-Urbana, he has established himself as an expert in Python programming and is well-versed in various aspects of software development.

With a strong background in machine learning, George has hands-on experience in utilizing regression models to tackle complex data analysis challenges. He has successfully implemented machine learning solutions and has a deep understanding of underlying principles.

In addition to his technical mastery George is proficient in working with cloud platforms such as AWS and Google Cloud enabling him to develop scalable and streamlined results for clients. He is committed to adhering to best practices in software development, including meticulous documentation of code and processes through testing and version control.

Speaker Bio

Mark Scannell

Mark Scannell works with strategic Financial Services customers and partners to identify, build, and scale repeatable patterns for building on and using Google Cloud Platform.

Mark Scannell has worked in Google Cloud for more than five years, including four years as a Strategic Cloud Engineer in both Data & Analytics and Application Development. BeforeGoogle, Google Mark worked at variety of banks (including Bear Stearns and BTMU) and vendors selling into banks (including Progress Software, StreamBase Systems, and Bloomberg) primarily focused on front office trading. Has a bachelor of Mathematics from the University of Waterloo in Canada.