



Introduction to Vertex Forecasting and Time Series in Practice

Course description

In this comprehensive course, you will gain a deep understanding of time series data and its unique characteristics, and learn how to leverage Google Cloud's Vertex AI for robust and scalable forecasting solutions. This course covers the end-to-end lifecycle of time series forecasting, from data ingestion and preparation to model training, evaluation, deployment, and monitoring. You will explore both Vertex AI's powerful AutoML for time series and the flexibility of custom model development (including statistical, machine learning, and deep learning approaches) within the Vertex AI ecosystem. Through practical insights and hands-on labs, you will address real-world forecasting challenges, understand key evaluation metrics, and implement best practices for MLOps in time series.

- Course level: Intermediate
- Duration: 2 days

Activities

This course includes presentations, hands-on labs, demonstrations, and group exercises.

Course objectives

In this course, you will learn to:

- Define time series data, identify its components, and recognize common challenges in forecasting.
- Prepare and preprocess diverse time series datasets for optimal performance on Vertex AI.
- Utilize Vertex AI AutoML for automated time series forecasting, understanding its configuration and output.
- Implement and train custom time series models (e.g., ARIMA, Prophet, XGBoost, LSTMs) using Vertex AI Custom Training.
- Evaluate forecasting model performance rigorously using appropriate metrics and visualization techniques.
- Deploy trained forecasting models to Vertex AI Endpoints for real-time and batch predictions.
- Monitor deployed forecasting models for data and concept drift, ensuring ongoing accuracy and reliability.
- Apply MLOps principles and advanced techniques like hierarchical and probabilistic forecasting in practice.

Intended audience

This course is intended for professionals working with time-dependent data and seeking to implement scalable forecasting solutions on Google Cloud:

- Data Scientists
- Machine Learning Engineers
- AI/ML Practitioners
- Business Intelligence Analysts working with time-dependent data
- Solution Architects interested in designing cloud-based ML forecasting systems
- Data Analysts aspiring to specialize in time series forecasting

Prerequisites

We recommend that attendees of this course have:

- Basic understanding of machine learning concepts (e.g., supervised learning, model evaluation).
- Familiarity with Python programming and common data science libraries (e.g., Pandas, NumPy).
- Basic knowledge of Google Cloud Platform (GCP) fundamentals (e.g., Cloud Storage, BigQuery).
- Prior experience with SQL is beneficial but not strictly required.

Course outline

Module 1: Introduction to Time Series and Vertex AI

- Time Series Data Fundamentals: Definition, components (trend, seasonality, cycle, noise), characteristics (autocorrelation, stationarity), types of time series problems.
- Challenges in Time Series Forecasting: Irregular intervals, missing data, exogenous variables, cold start problem, dynamic environments.
- Overview of Vertex AI: Unified ML platform on Google Cloud, key components (Workbench, Training, Endpoints, Feature Store, Pipelines, Experiments).
- Vertex AI for Time Series Forecasting: Advantages of using Vertex AI for scalable, managed, and robust forecasting solutions, integration with other GCP services.
- Demo: Navigating the Vertex AI Console, setting up a new project, and exploring key functionalities.

Module 2: Data Preparation for Time Series Forecasting on Vertex AI

- Time Series Data Ingestion: Loading data from various sources such as BigQuery, Cloud Storage, and external databases into Vertex AI Workbench.
- Data Cleaning and Preprocessing: Strategies for handling missing values (interpolation, imputation), outlier detection, and data type conversions.
- Feature Engineering for Time Series:
 - Lag Features: Creating lagged values of the target variable and relevant features.
 - Rolling Statistics: Calculating moving averages, standard deviations, and other statistics over defined windows.
 - Date and Time Features: Extracting granular information (day of week, month, year, quarter, hour), creating holiday indicators, and workday flags.
 - Exogenous Variables: Incorporating external factors like weather data, promotional events, economic indicators, or competitor activity.
- Data Splitting Strategies: Time-based splits (train, validation, test sets), walk- forward validation for robust model evaluation.
- Hands-on Lab: Preparing a sample time series dataset, performing feature engineering, and splitting data for forecasting on Vertex AI.

Module 3: Exploring Vertex AI AutoML for Time Series Forecasting

- Introduction to Vertex AI AutoML Forecasting: Leveraging automated machine learning for efficient and performant time series forecasting.
- Configuring AutoML Forecasting Jobs: Defining the target column, time column, forecast horizon, context window, and specifying categorical/numerical features.
- Understanding AutoML Model Training: Overview of the underlying ensemble methods, hyperparameter optimization, and built-in algorithms used by AutoML.
- Interpreting AutoML Results: Analyzing model evaluation metrics (RMSE, MAE, MAPE), examining residual analysis plots, and understanding feature importance.
- Limitations and Best Practices for AutoML: When to use AutoML, considerations for highly complex or specialized time series scenarios, and tips for optimizing performance.
- Hands-on Lab: Building, training, and evaluating an AutoML forecasting model using a real-world dataset on Vertex AI.

Module 4: Custom Models and Training on Vertex AI

- Introduction to Custom Training on Vertex AI: When and why to use custom models, leveraging managed infrastructure for custom code execution.
- Common Time Series Models:
 - Statistical Models: ARIMA, SARIMA, Prophet (Facebook Prophet), Exponential Smoothing (ETS).
 - Machine Learning Models: XGBoost, Random Forest, LightGBM (adapted for time series features).
 - Deep Learning Models: Recurrent Neural Networks (RNNs) like LSTMs and GRUs, Encoder-Decoder architectures, and Transformer networks (brief overview).
- Developing Custom Training Scripts: Writing Python training code, using pre-built TensorFlow/PyTorch containers, or creating custom Docker containers.
- Configuring Custom Training Jobs: Specifying machine types, GPU accelerators, distributed training setups, and container images.
- Hyperparameter Tuning with Vertex AI Vizier: Utilizing Vertex AI Vizier for efficient hyperparameter optimization through Bayesian optimization, grid search, or random search for custom time series models.

- Hands-on Lab: Developing a Python training script for a Prophet or custom scikit-learn based time series model, and executing a custom training job on Vertex AI.

Module 5: Model Evaluation and Interpretation

- Key Forecasting Metrics:
 - Error-based: Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE).
 - Percentage-based: Mean Absolute Percentage Error (MAPE), Symmetric MAPE (SMAPE).
 - Scale-free: Mean Absolute Scaled Error (MASE), R-squared (coefficient of determination).
- Visualizing Forecasts: Creating actual vs. predicted plots, visualizing confidence intervals, and plotting residuals over time.
- Residual Analysis: Checking for patterns, autocorrelation, heteroscedasticity, and normality in forecasting errors to diagnose model shortcomings.
- Backtesting and Walk-Forward Validation: Implementing robust evaluation strategies for time series data to simulate real-world performance.
- Hands-on Lab: Analyzing model performance metrics and visualizing forecasts for both AutoML and custom models, comparing their strengths and weaknesses.

Module 6: Deployment and Monitoring

- Deploying Models to Vertex AI Endpoints: Creating and managing real-time prediction endpoints for trained models.
- Online vs. Batch Predictions: Choosing the appropriate prediction strategy based on latency, throughput, and operational requirements.
- Making Predictions: Interacting with Vertex AI Endpoints using the Vertex AI SDK for Python, gcloud CLI, or REST API for inference.
- Model Monitoring on Vertex AI: Setting up and configuring drift detection for time series models, including data drift (changes in input features) and concept drift (changes in the relationship between features and target).
- Setting up Alerts and Logging: Integrating with Cloud Monitoring and Cloud Logging for operational insights, anomaly detection, and automated alerts on model performance degradation.
- Hands-on Lab: Deploying a chosen time series model (AutoML or custom) to a Vertex AI Endpoint and making online predictions.

Module 7: Advanced Concepts and Best Practices

- Hierarchical Forecasting: Techniques for forecasting at multiple aggregation levels (e.g., store, region, product category) and ensuring consistency across hierarchies.
- Probabilistic Forecasting: Generating prediction intervals and quantiles to provide uncertainty estimates alongside point forecasts.
- Interpretable AI for Time Series: Explaining model predictions using techniques like SHAP (SHapley Additive exPlanations) or LIME (Local Interpretable Model-agnostic Explanations) where applicable.
- MLOps for Time Series Forecasting: Designing and implementing CI/CD pipelines for time series models, versioning models and data, and automating model retraining and redeployment.
- Ethical Considerations in Forecasting: Addressing bias, fairness, and the societal impact of forecasting predictions.
- Case Studies: Exploring real-world examples of Vertex AI Forecasting applications across various industries (retail, finance, energy, healthcare).

Module 8: Course Wrap-up

- Review of Key Concepts: Comprehensive summary of Vertex Forecasting functionalities, time series analysis principles, and practical best practices.
- Next Steps and Additional Resources: Pointers to official Google Cloud documentation, community forums, advanced courses, and relevant research papers.
- Q&A and Discussion: Open session for participant questions, sharing experiences, and discussing future learning paths.
- Course Survey: Opportunity for participants to provide feedback on the course content and delivery.

Your course overview ends here.
But your cloud journey doesn't!

Unlock advanced training and certification pathways with **CloudThat**.



About CloudThat

Established in 2012, CloudThat is an award-winning company and the first in India to offer cloud training and consulting services for individuals and enterprises worldwide. Recently, it won Google Cloud's New Training Partner of the Year Award for 2025, becoming the first company in the world in 2025 to hold awards from all three major cloud giants: AWS, Microsoft, and Google. CloudThat notably won consecutive AWS Training Partner of the Year (APJ) awards in 2023 and 2024 and the Microsoft Training Services Partner of the Year Award in 2024, bringing its total award count to an impressive 12 awards in the last 8 years.

CloudThat specializes in Cloud Migration, Data Platforms, DevOps, IoT, and cutting-edge technologies like Gen AI & AI/ML. As a Microsoft Solutions Partner, AWS Premier Tier Training Partner, Google Cloud Platform Partner, and collaborator with leading organizations like HPE, Oracle, Nvidia and Databricks, CloudThat has trained over 850,000 professionals across 600+ cloud certifications, empowering students and professionals worldwide to advance their skills and careers.

To know more about our cloud certification training, email at sales@cloudthat.com or call us at [+91-8880002200](tel:+91-8880002200) or [+1-8555588830](tel:+1-8555588830) for those from outside India.