AWS re:Invent

AIM335-R

Accelerate time-series forecasting with Amazon Forecast

Eric Greene
Sr. ML Solutions Architect
Amazon Web Services

Gunjan Garg
Sr. Software Development Engineer
Amazon Web Services

AWS re:Invent

aws

Agenda

- Overview
  - Time-series forecasting
  - Amazon Forecast
  - Lab: Setup and data import
- Predictors
  - Amazon Forecast recipes
  - Lab: Train a predictor
- Evaluate and deploy
  - Evaluation techniques
  - Lab: Deploy and evaluate a predictor
- Next steps

Time-series forecasting

Time-series forecasting

The science of predicting future values based on past values in a time-ordered sequence of data

Forecasting

Accurate forecasting is relevant in a broad set of business scenarios

Product demand	Electricity, gas, or water usage	Staffing at call centers	Transportation needs in an area
Cash flow	Cash in an ATM	Flight ticket prices	Ads impressions
Inventory planning	Pay-as-you-go service usage	Deliveries per ZIP code	Patient volume
Real estate prices	Broadband usage for telecom	Selling price of crops	Sensor network monitoring

Forecasting

Impact of under- and over-forecasting

	Inventory planning	Workforce planning	Capacity planning	Financial planning
Over forecasting	Cost of holding excess inventory	Cost of more labor	Cost of over provisioning	Cost of overreporting
Under forecasting	Cost of customer dissatisfaction	Cost of overtime	Cost of unmet demand	Cost of undercutting

Forecasting at Amazon.com

Accurate forecasting is critical for delivering on customer promises

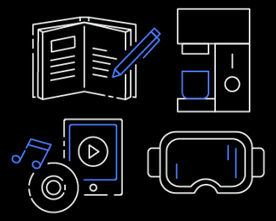

Product availability

Demand forecasting for over 400M products across 10,000 ZIP codes

Lower price

Inventory and fulfillment cost reduction to provide customers low prices

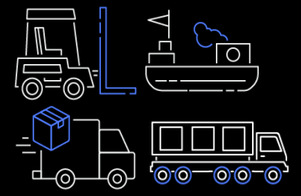

Fast delivery

12 shipping options, with free same-day delivery

Forecasting problems

Traditional statistics can predict demand for some products with reasonable accuracy

Everyday household products

Products specific to seasons

Forecasting problems

Forecasting may be inaccurate in some situations. New or recently promoted products may be hard to accurately model.

High price variability

Regional vs. national demand

Slow-moving products

New products

Highly seasonal products

What is Amazon Forecast?

The AWS ML stack

Broadest and deepest set of capabilities

AI services

Vision	Speech	Language	Chatbots	Forecasting	Recommendations
Amazon Rekognition Image Amazon Rekognition Video Amazon Textract	Amazon Polly Amazon Transcribe	Amazon Translate Amazon Comprehend and Amazon Comprehend Medical	Amazon Lex	Amazon Forecast	Amazon Personalize

ML services

Amazon SageMaker

Ground Truth	Notebooks	Algorithms + AWS Marketplace	Reinforcement learning	Training	Optimization	Deployment	Hosting

ML frameworks + infrastructure

Frameworks	Interfaces
TensorFlow, mxnet, PYTORCH	GLUON, Keras

Infrastructure
Amazon EC2 P3 and P3DN, Amazon EC2 G4, Amazon EC2 C5, FPGAs, DL containers and AMIs, AWS IoT Greengrass, Amazon Elastic Inference, AWS Inferentia

Amazon Forecast

The technology that powers the world's largest e-commerce business

Get started with a few clicks
Point Amazon Forecast to your data

Train your custom ML model
Choose a model or let Amazon Forecast auto-select the best one for your data through AutoML

Download accurate forecasts
Retrieve forecasts through your private API

Amazon Forecast

Behind the scenes

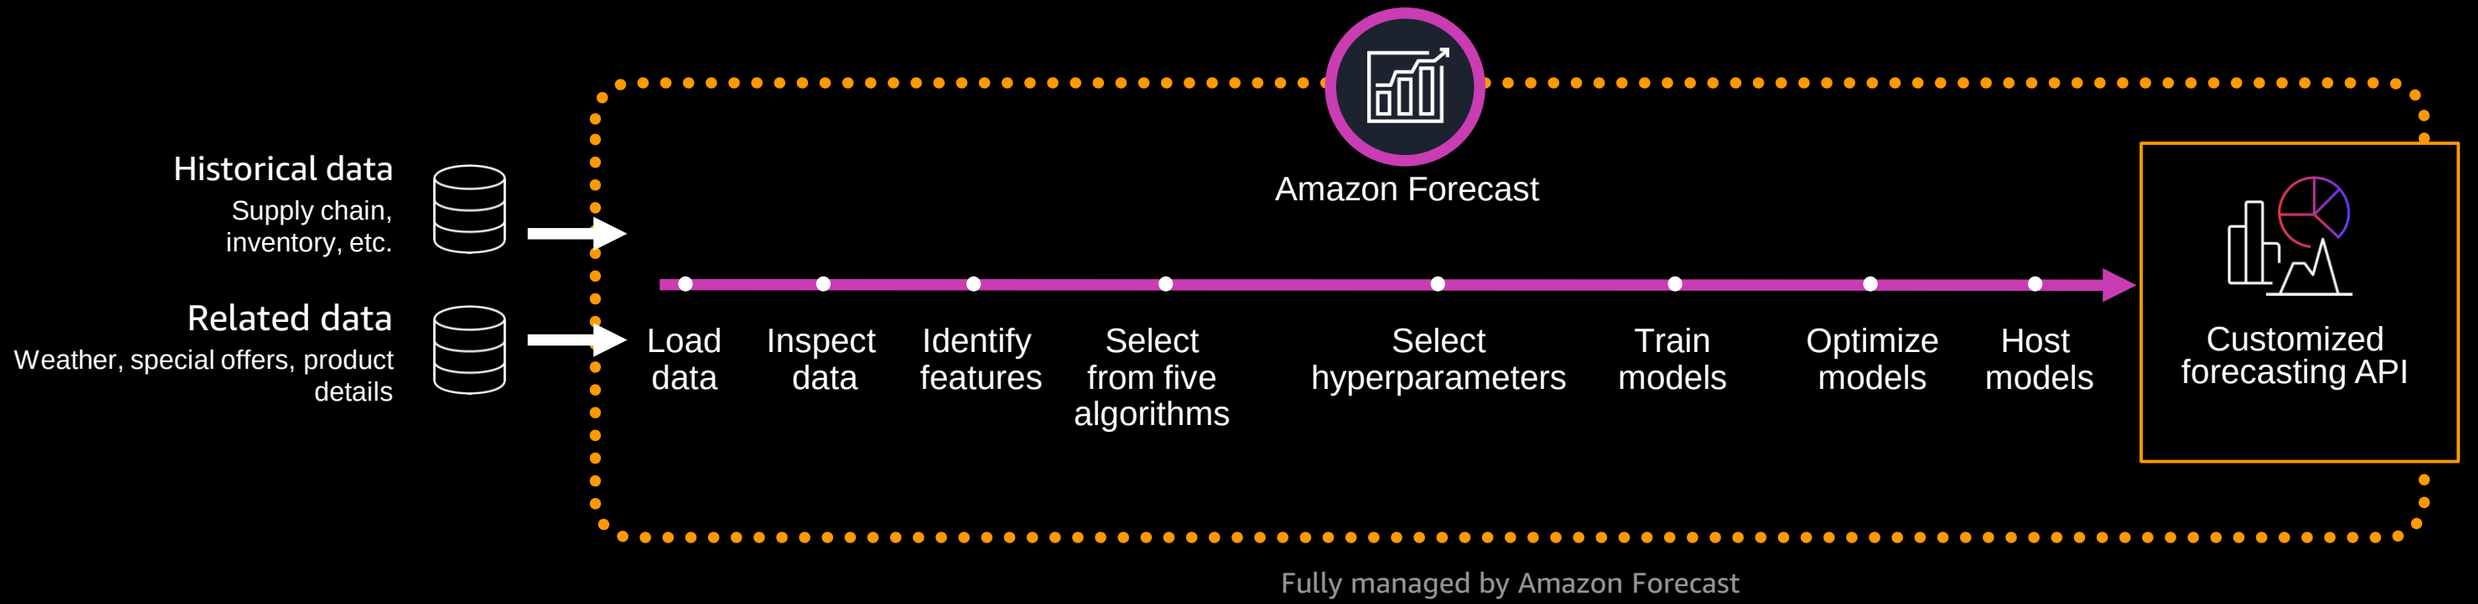

Amazon Forecast terminology – Data

Dataset group – Container for all datasets.

Dataset – The collective data of a given type. There are three types of datasets: target time series, item metadata, and related time series.

Target time series – Historical dataset where each row contains an item identifier, date/time, and the measured value to be predicted.

Item metadata – Dataset of item attributes. This dataset contains attributes that do not change about an item. These attributes establish relationships or similarity between items. (Optional)

Related time series – Historical dataset where each row contains an item identifier, date/time, and values that are *related* to the value being predicted. Used to correlate things like promotions, weather, or any related events in time to your prediction. (Optional)

What is a Notebook?

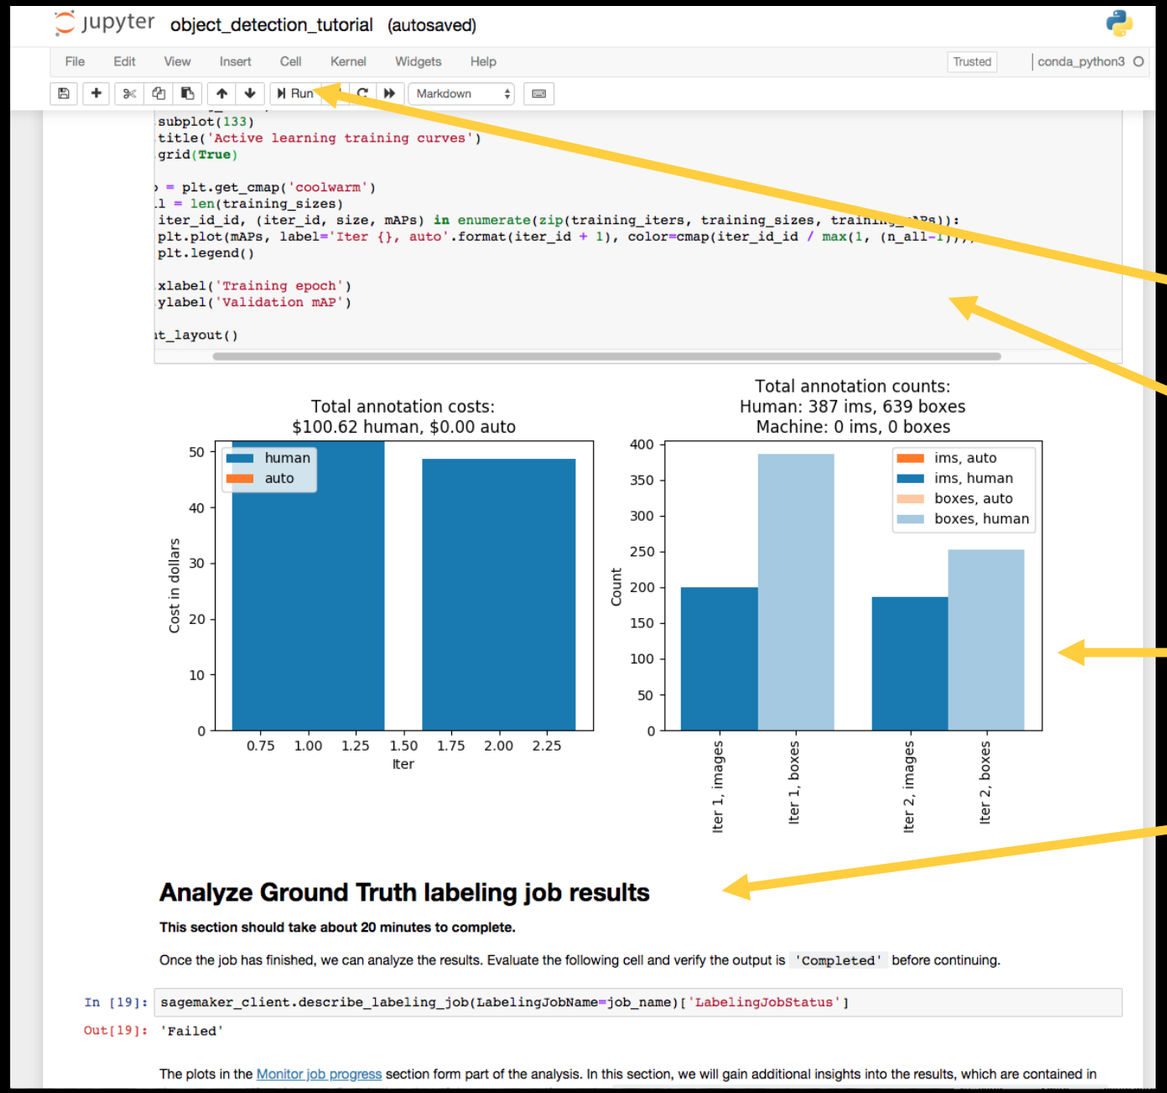

Notebooks are where data science discovery, modeling and analysis work is done. Notebooks contain:

Run button

Executable code

Graphs & charts

Documentation

Lab: Setup and data import

Instructions: Setup and data import

1. Open the workshop URL given to you and log in with your credentials. Open the console and go to the Amazon SageMaker dashboard.
2. Navigate to the notebooks instances and click “Open Jupyter” for “AmazonForecastWorkshop.”
3. Go to the notebooks directory.
4. Open the notebook “9. Amazon Forecast Workshop.”
5. Read the instructions and complete sections 1–2.
6. Custom coding sections are optional.
7. Use hints to help you complete the optional parts of the lab.
8. Still need help? Raise your hand.
9. To speed up completion, run all of step 3 prior to next lecture.

Amazon Forecast algorithms

Amazon Forecast terminology – Training

Predictor – The model that was trained with the datasets in the dataset group using the selected algorithms.

Algorithms – Different statistical and mathematical formulas and techniques used to train predictors.

Forecast – The output predictions from a model. Could be a probabilistic or point value depending on algorithms used.

Forecast frequency – The interval between measured values to be forecast. This could be anywhere from five minutes to one month. If the data is of greater frequency than the forecast frequency, it can be aggregated.

Forecast horizon – How far into the future the model will be predicting values. For example, a forecast horizon of seven days means a model will predict values from today until next week.

Algorithms powering Amazon Forecast

ARIMA

Auto-**r**egressive **i**ntegrated **m**oving **a**verage:

De facto statistical method

Works well with small number of time series. Classical approach to model autocorrelations.

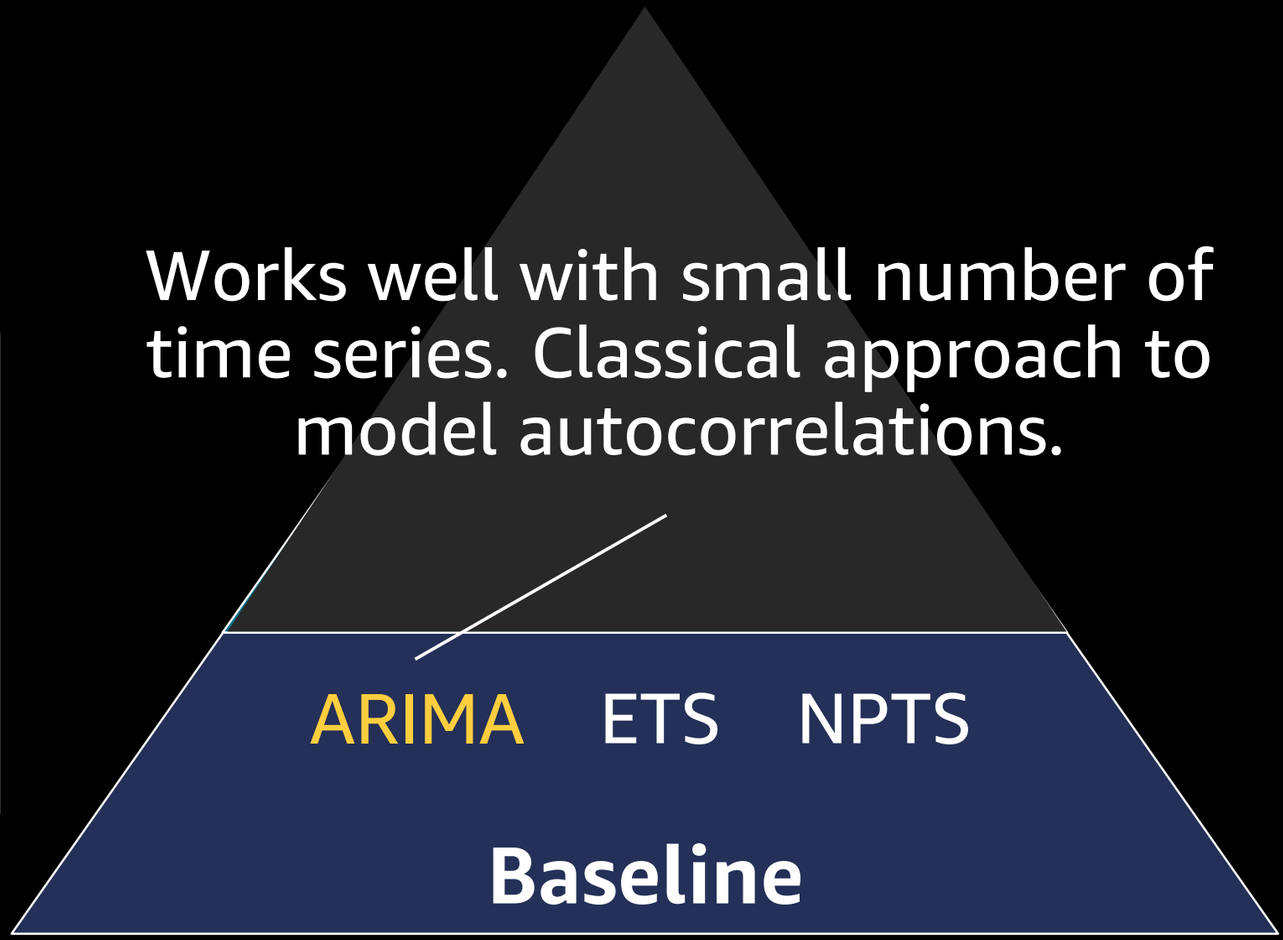

Algorithms powering Amazon Forecast

ETS

Error trend seasonality:

Statistical method that uses exponential smoothing

Works well with small number of time series. Finds trend, seasonality, and residual.

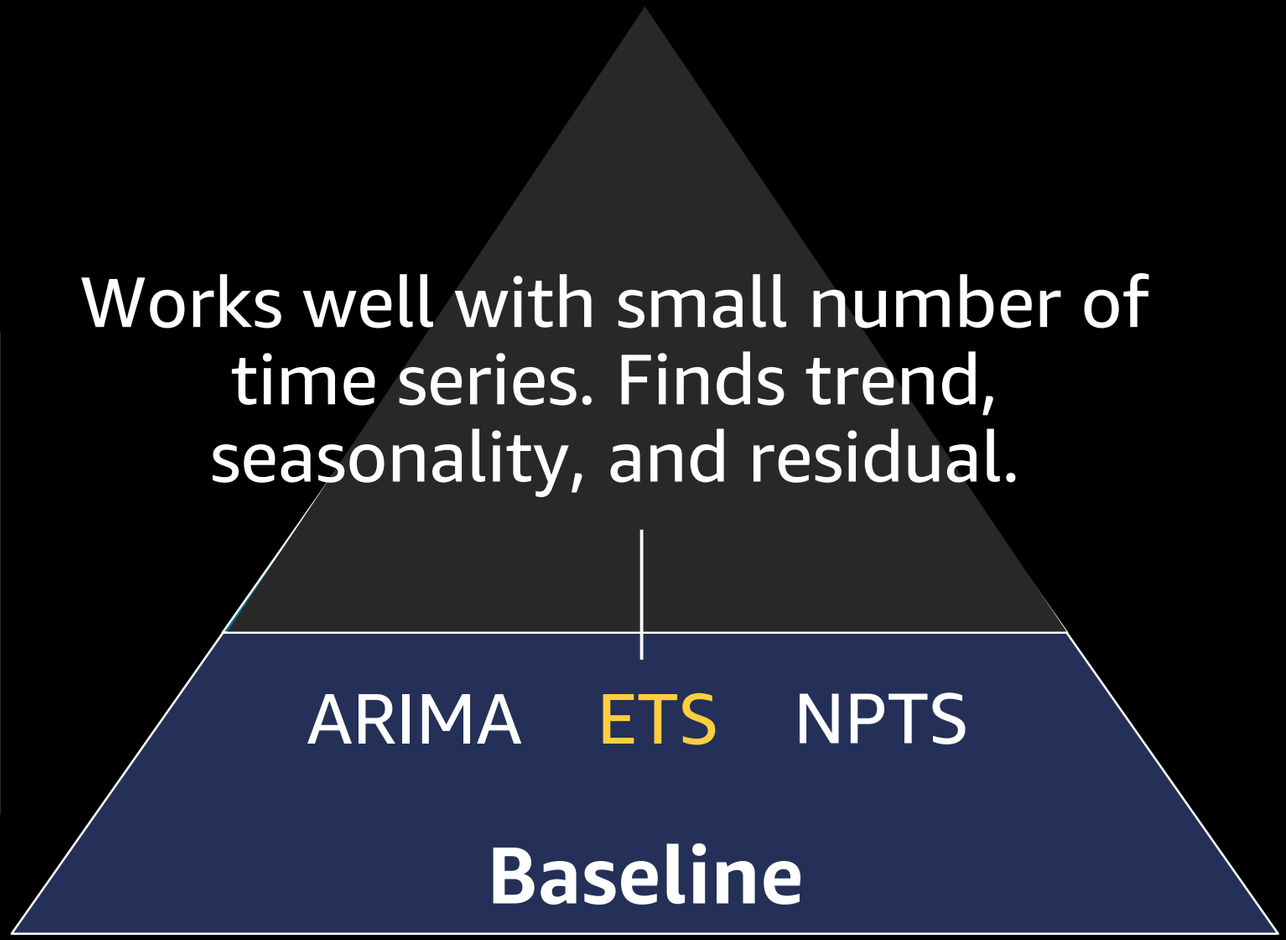

Algorithms powering Amazon Forecast

NPTS

Nonparametric time series:

Another statistical method

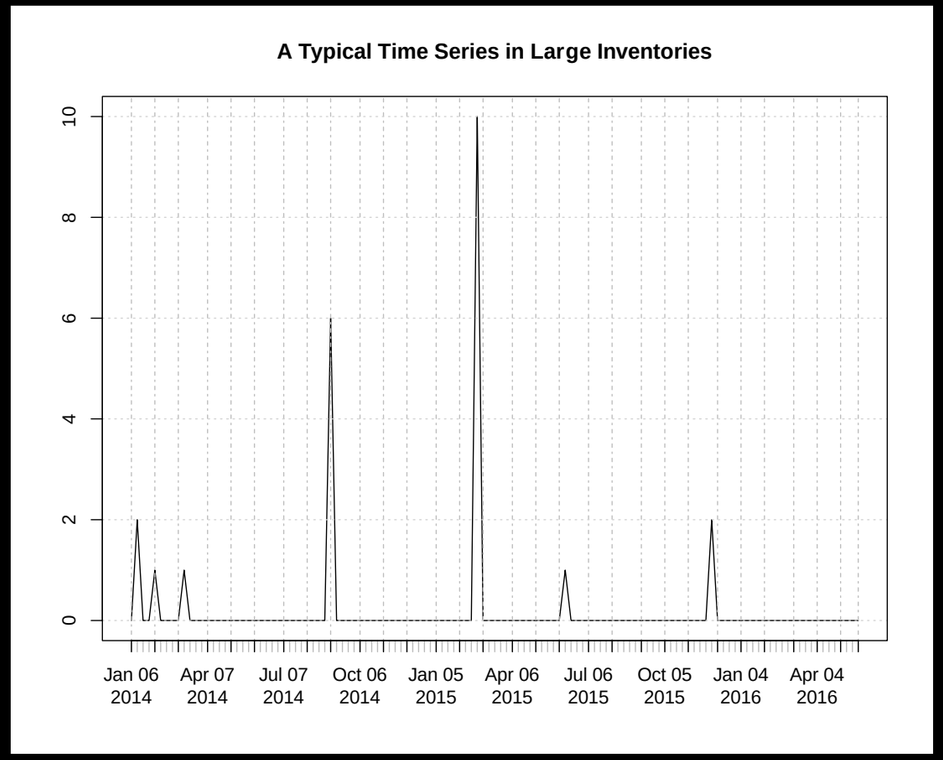

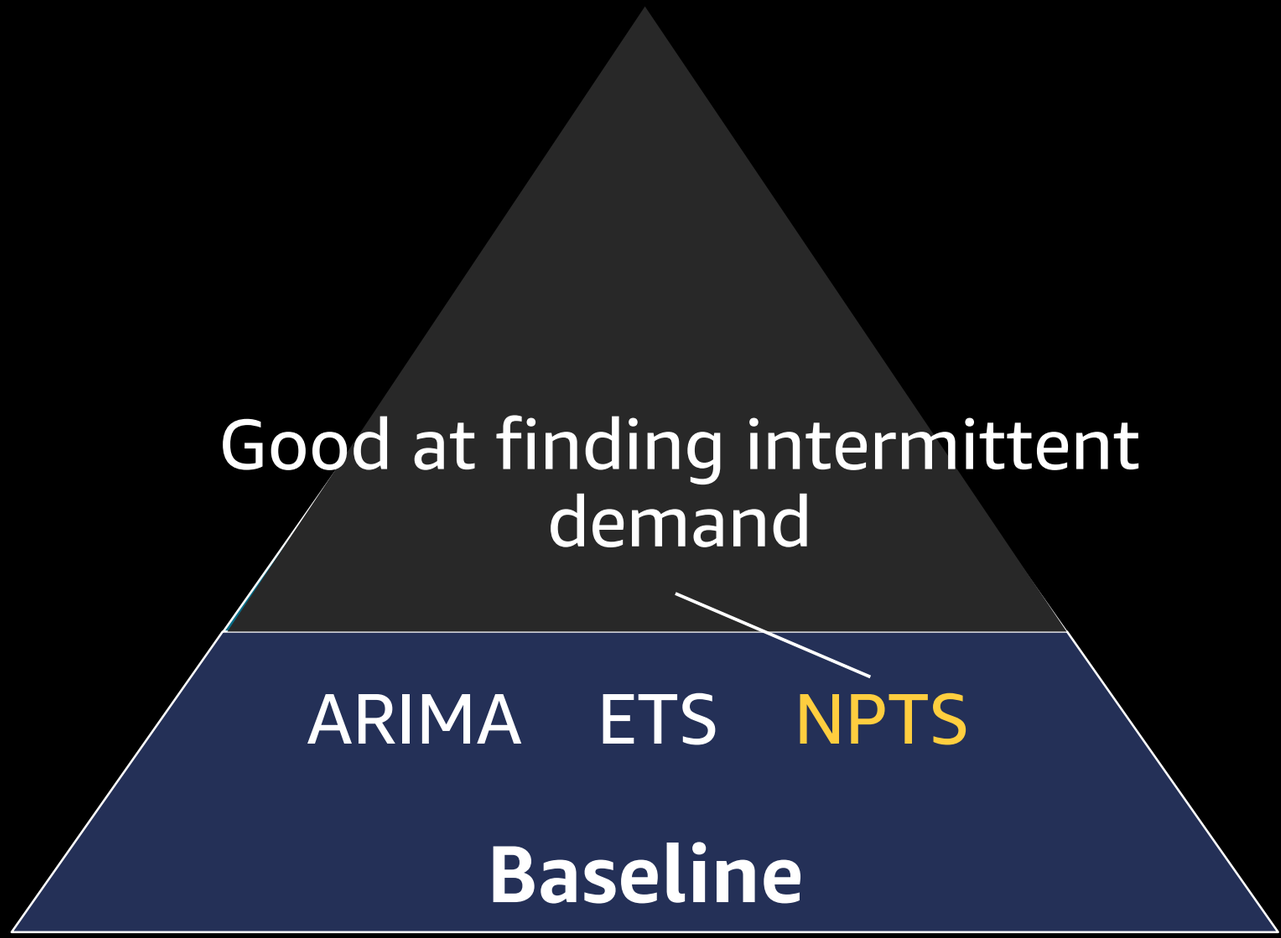

Algorithms powering Amazon Forecast

Prophet

Additive model with
Gaussian likelihood

Can find trend, seasonality, cyclical, and holiday effects

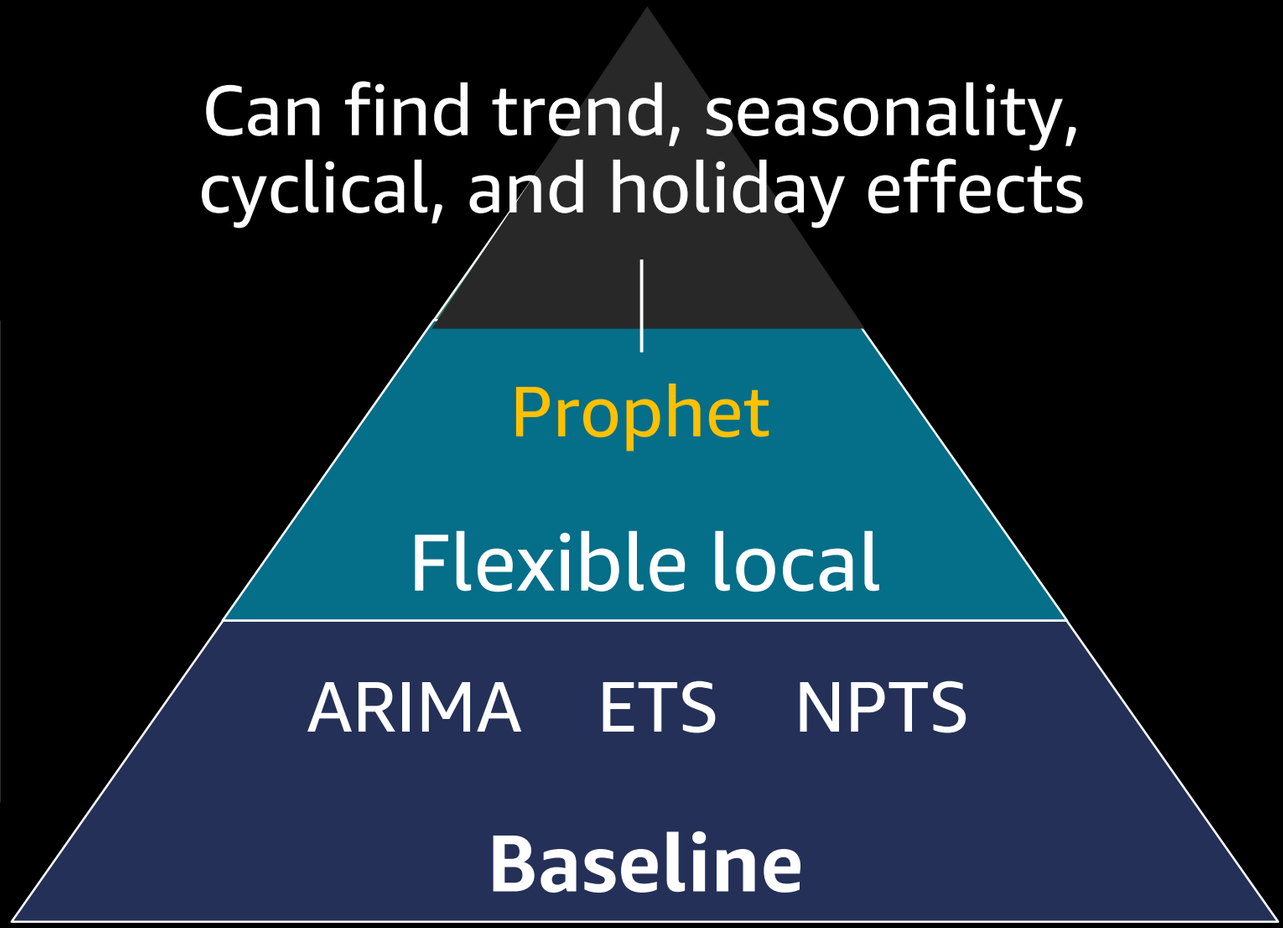

Algorithms powering Amazon Forecast

DeepAR++

General model that can use related time series and attributes

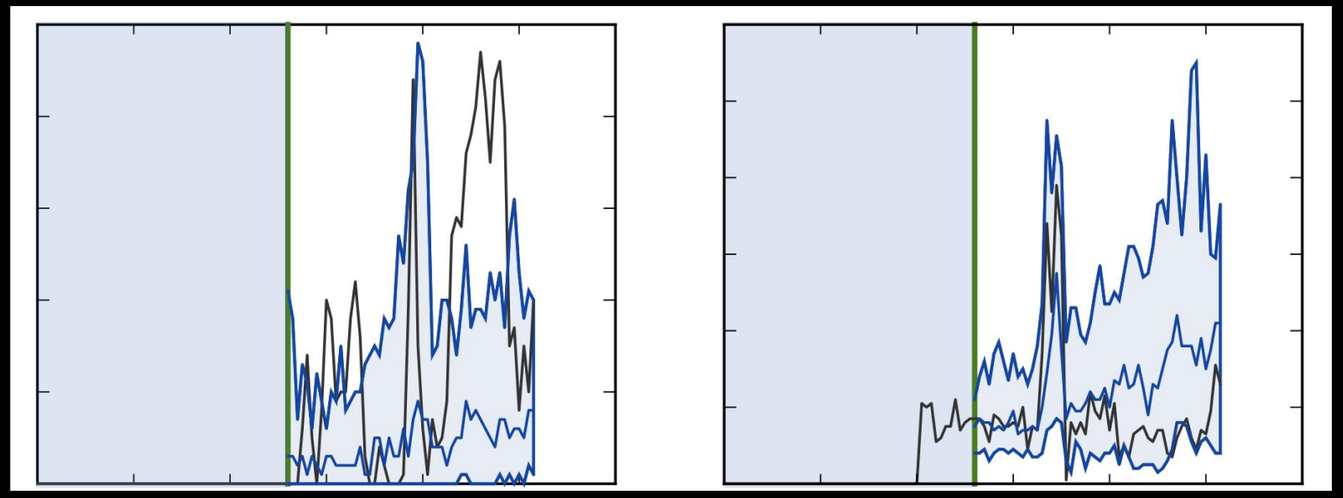

Good for many related time-series and cold-start problems

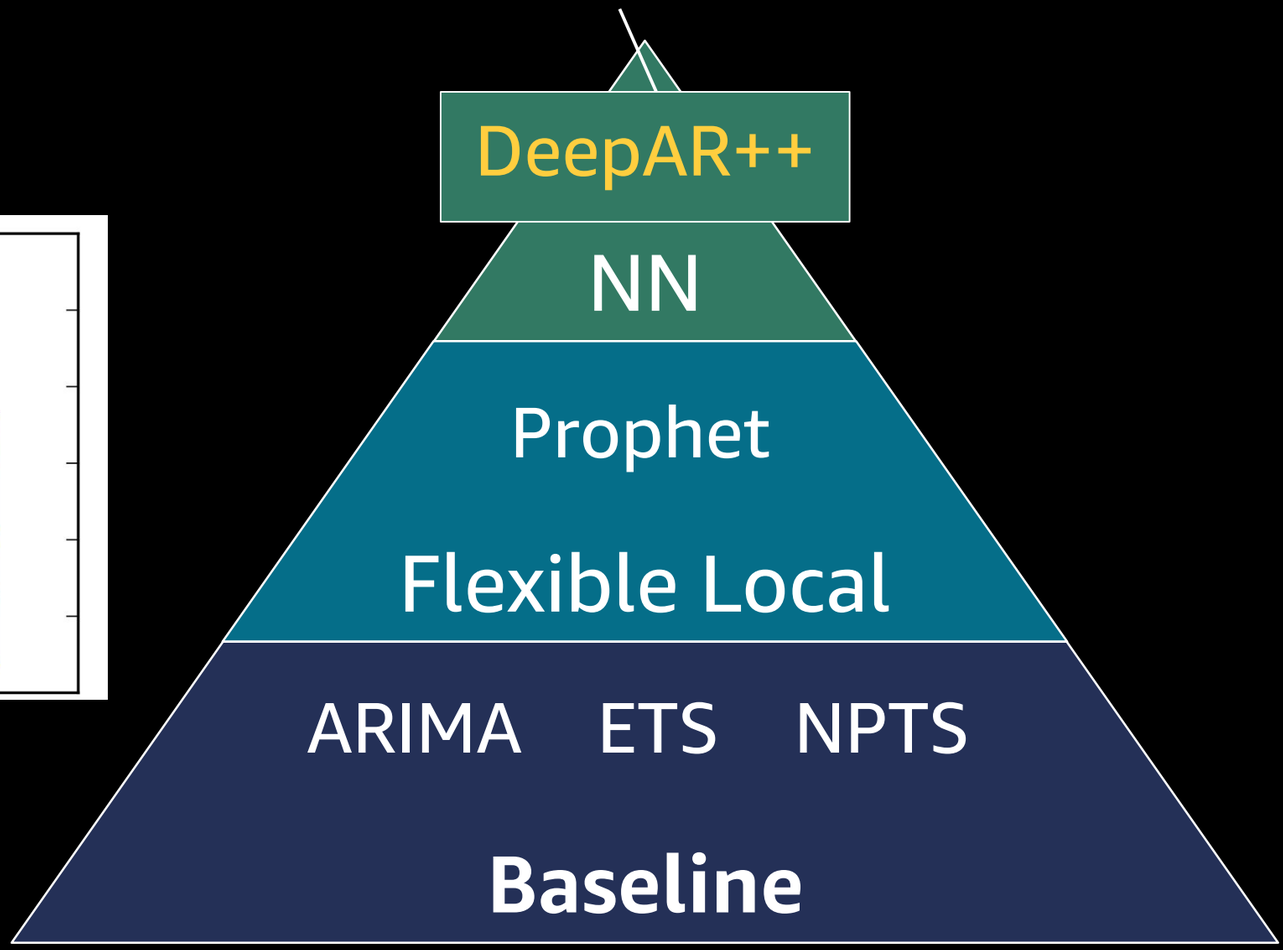

Lab: Train a predictor

AWS re:Invent

aws

Instructions: Train a predictor

1. Go back to the same notebook as before and complete section 3.
2. Section 3 will take some time as a model is being trained. We will talk about evaluation prior to the predictor training finishing.
3. Go to the Forecast Console and explore the data and training jobs you have created while waiting for training.
4. Need help? Use the hint command or raise your hand.

Evaluation techniques

Evaluation – Understanding quantiles

Quantiles are used in probabilistic forecasts to provide ranges of predicted values

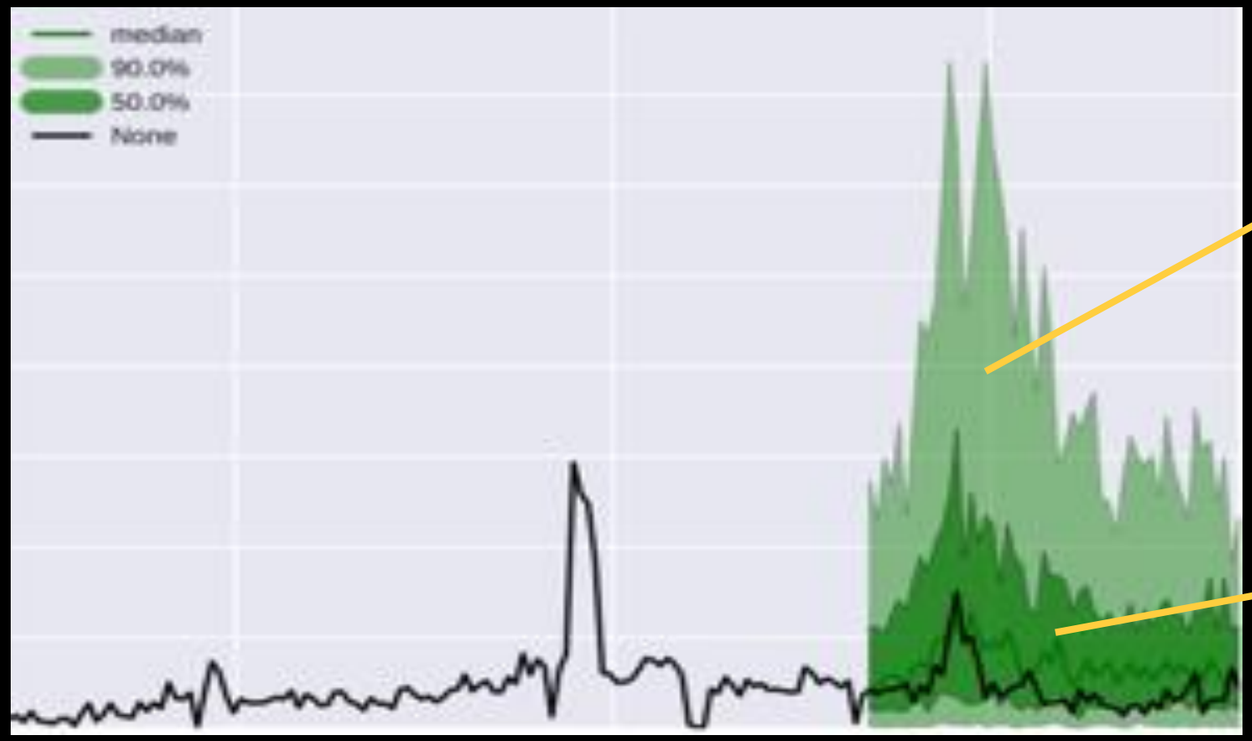

Examples

A P90 quantile predicts that 90% of the time, the true value will be less than the predicted value

A P50 quantile predicts 50% of time the true value will be less than the predicted value

Evaluation – Understanding error and loss

Error/loss functions are used to evaluate machine learning models by calculating the error between true and predicted values. Error functions are chosen based on the algorithms and type of data.

Weighted quantile loss – The weighted quantile loss (wQuantileLoss) calculates how far off the forecast a certain quantile is from actual demand.

RMSE (root mean square error) – Calculates the difference between the actual target value and the predicted (forecasted) mean value.

Many more – MAPE, MASE, log loss

Lab: Deploy and evaluate a predictor

AWS re:Invent

aws

Instructions: Deploy and evaluate a predictor

1. Go back to the same notebook as before and complete sections 4-7.
2. You can skip step 8 (“clean up”). This will happen automatically.
3. Need help? Use the hint command or raise your hand.

Next steps

Amazon Forecast in production

Automate into a repeating process

Define your integration points and create data pipelines for input data sources

Continually infer new predictions based on time-series frequency

Train models often to capture changes in the environment

Automate evaluation and trigger manual evaluation based on error thresholds

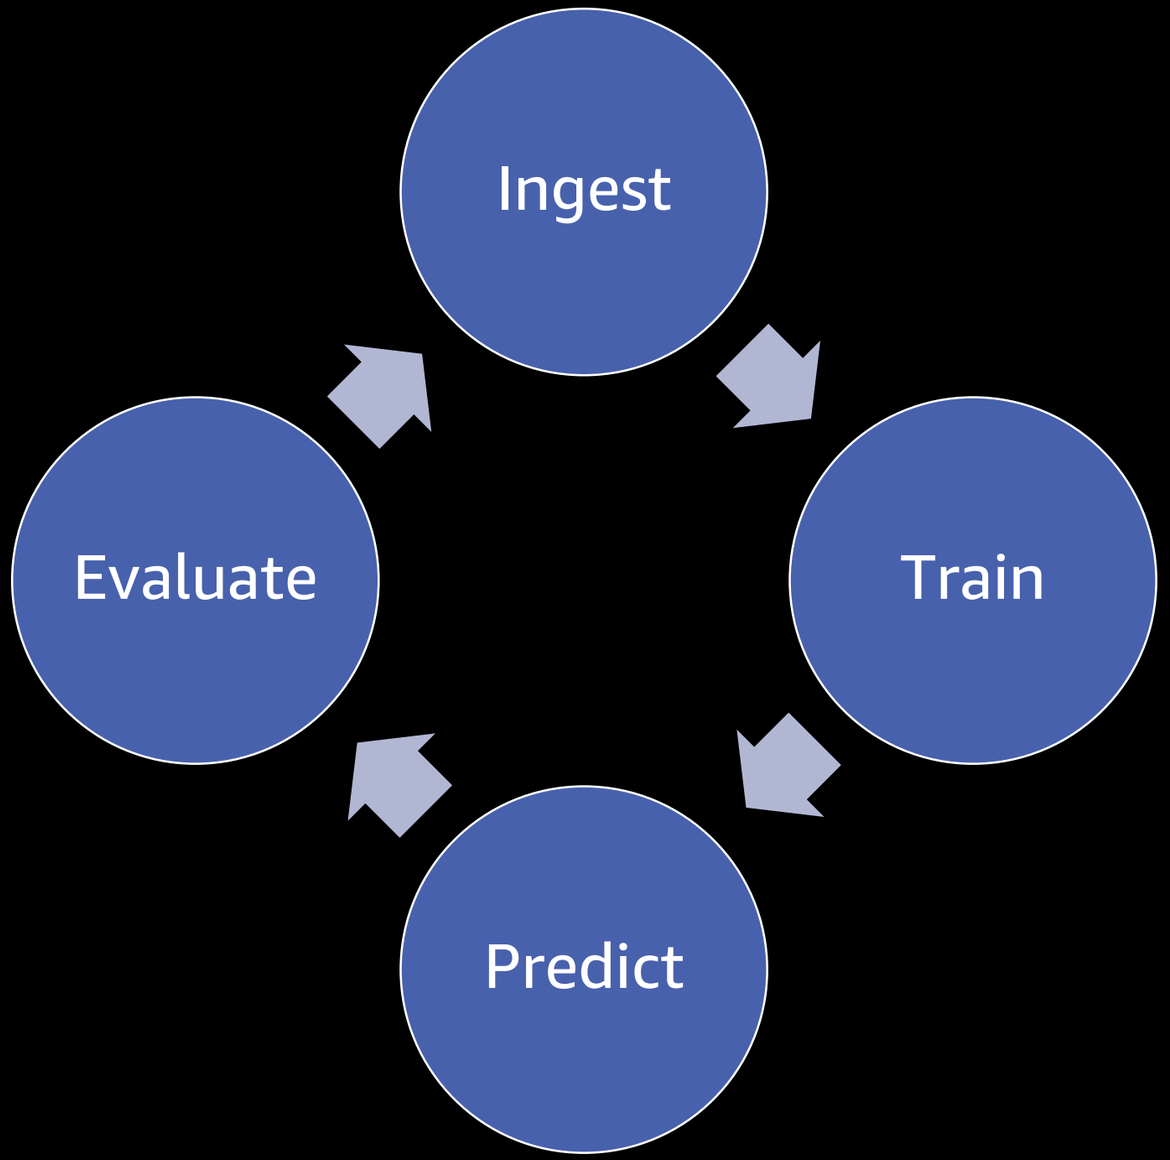

Amazon Forecast and Amazon Redshift cluster management

Manages a set of EC2 instances

Maintains cache pools of EC2 instances

Needs to forecast cache pool size

1. Amazon EC2 cache pool demand changes are published to S3 bucket
2. New data is ingested into Amazon Forecast and new forecast predictions are stored in S3 bucket
3. A Lambda function copies new forecasts to a DynamoDB table
4. The cluster scaling logic reads forecasts and adjusts the cache pool size based on projected demand

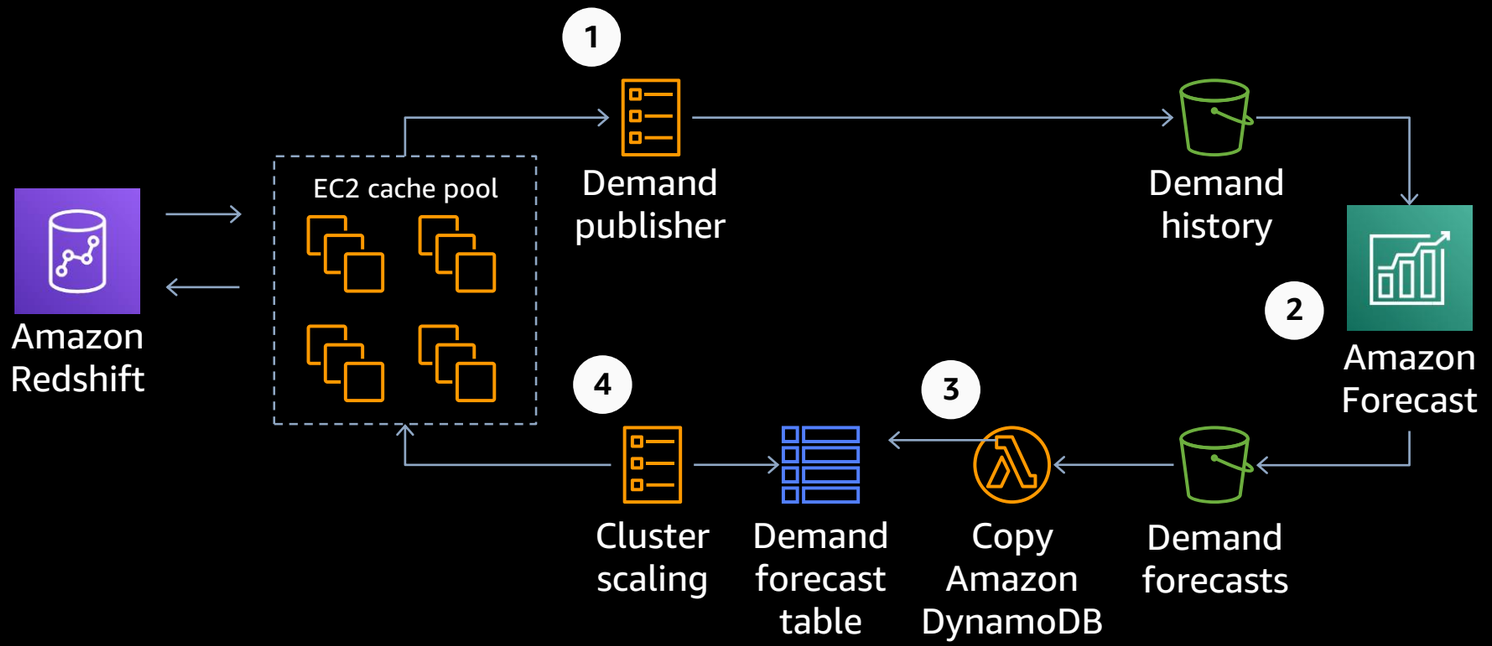

Next steps

Amazon Forecast tutorials

github.com/aws-samples/amazon-forecast-samples/tree/master/notebooks

www.youtube.com/watch?v=Ed4j8Olf_E8

Amazon Forecast documentation

docs.aws.amazon.com/forecast/latest/dg/what-is-forecast.html

docs.aws.amazon.com/forecast/latest/dg/API_Operations.html

Time series with Amazon SageMaker

github.com/aws-samples/amazon-sagemaker-stock-prediction

www.youtube.com/watch?v=g8UYGh0tlK0

Time series with GluonTS

gluon-ts.mxnet.io/

github.com/awslabs/gluon-ts

Related sessions

AIM312 – Predict future business outcomes using Amazon Forecast

AIM423-R – Integrate forecasting with your retail POS

GPSTEC318 – Reinforcement learning: Using AI/ML to boost your software development on AWS

RET309 – Use Amazon Forecast to more accurately predict future retail sales

GPSBUS205 – Building ML practices to address the top four use cases

Learn ML with AWS Training and Certification

The same training that our own developers use, now available on demand

Role-based ML learning paths for developers, data scientists, data platform engineers, and business decision makers

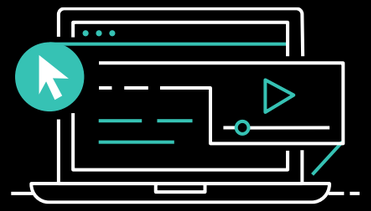

70+ free digital ML courses from AWS experts let you learn from real-world challenges tackled at AWS

Validate expertise with the
AWS Certified Machine Learning - Specialty exam

Visit https://aws.training/machinelearning

For Partners and ISVs

Amazon Forecast – Partnering for success with forecasting powered by Machine Learning.

In this session, partner and ISV decision makers will have the opportunity to learn about how our customers and existing partners are working together to prepare data, drive results and analyze opportunities with the help of Amazon Forecast.

This Thursday from 4 PM to 4:45 PM at the Encore at Wynn,
Chopin Rooms 3-4

Thank you!

AWS re:Invent

aws

Please complete the session survey in the mobile app.