



AI and ML Week

Forecasting service demand and planning capacity



Paolo Emilio Barbano
Sr. Data Scientist
Artificial intelligence and machine learning, AWS



Agenda

- Statistical modeling of time series
- Traditional vs. machine learning-based models
- AWS services including Amazon Forecast
- Demo of Amazon Forecast

Applications of forecasting in the public sector



Student demand and drop-outs



Citizen transportation, web, and benefits usage



Budget and staffing loads



Donations and donor activities



Population health data and research results



Financial analysis

Key challenges with statistical modeling of time series

- How do we describe the important features of the time series pattern?
- How is the past behavior affecting the future?
- How accurately can we predict future values of the series?

Three Approaches

Traditional
methods

Deep
learning
(DL)
methods

Combination:
Amazon
Forecast

Traditional time-series models

Advantages

- Independent forecasts
- Strong structural assumptions
- De facto industry standard
- Well-understood,
 >50 yrs. research
- High data efficiency

Algorithm types

- Nonparametric time series model
- Exponential smoothing (ETS)
- (Auto-) ARIMA
- Prophet

Limitations

- Metadata can't be naturally embedded in the analysis
- External factors are difficult to include
- Historical data is required

Keep in mind...

- Data must match the structural assumptions
- You cannot identify patterns across a time series

Deep learning methods


Deep learning outperforms traditional methods in a variety of cases

Deep learning tends to do well on highly complex, interpret-able data

Deep learning may be able to handle cold-start situations

Amazon Forecast

Automated machine learning service for accurate forecasting



Fully managed service

Automatically sets up data pipeline, training and prediction



Highly accurate

50% improvement in accuracy over traditional methods



Easy to use

No deep learning experience required



Your data, your models

Encrypted with customer keys through Amazon Key Management Service

Amazon Forecast – How it works

Historical Data

Sales, Inventory, Pricing, etc.

Related Data

Weather, Competitive Promotions etc.

Mata-Data

Color, City, Country, Author etc.



1. Load Data

2. Inspect Data

3. Identity Features

4. Algorithm Selection

5. Hyper-parameters Selection

6. Model Training

7. Optimize Models

8. Model Deployment and Hosting



Private
Customized
Forecasting
API

Amazon Forecast

There are three types of datasets in Amazon Forecast

Target time-series

Historic time series data of items to forecast

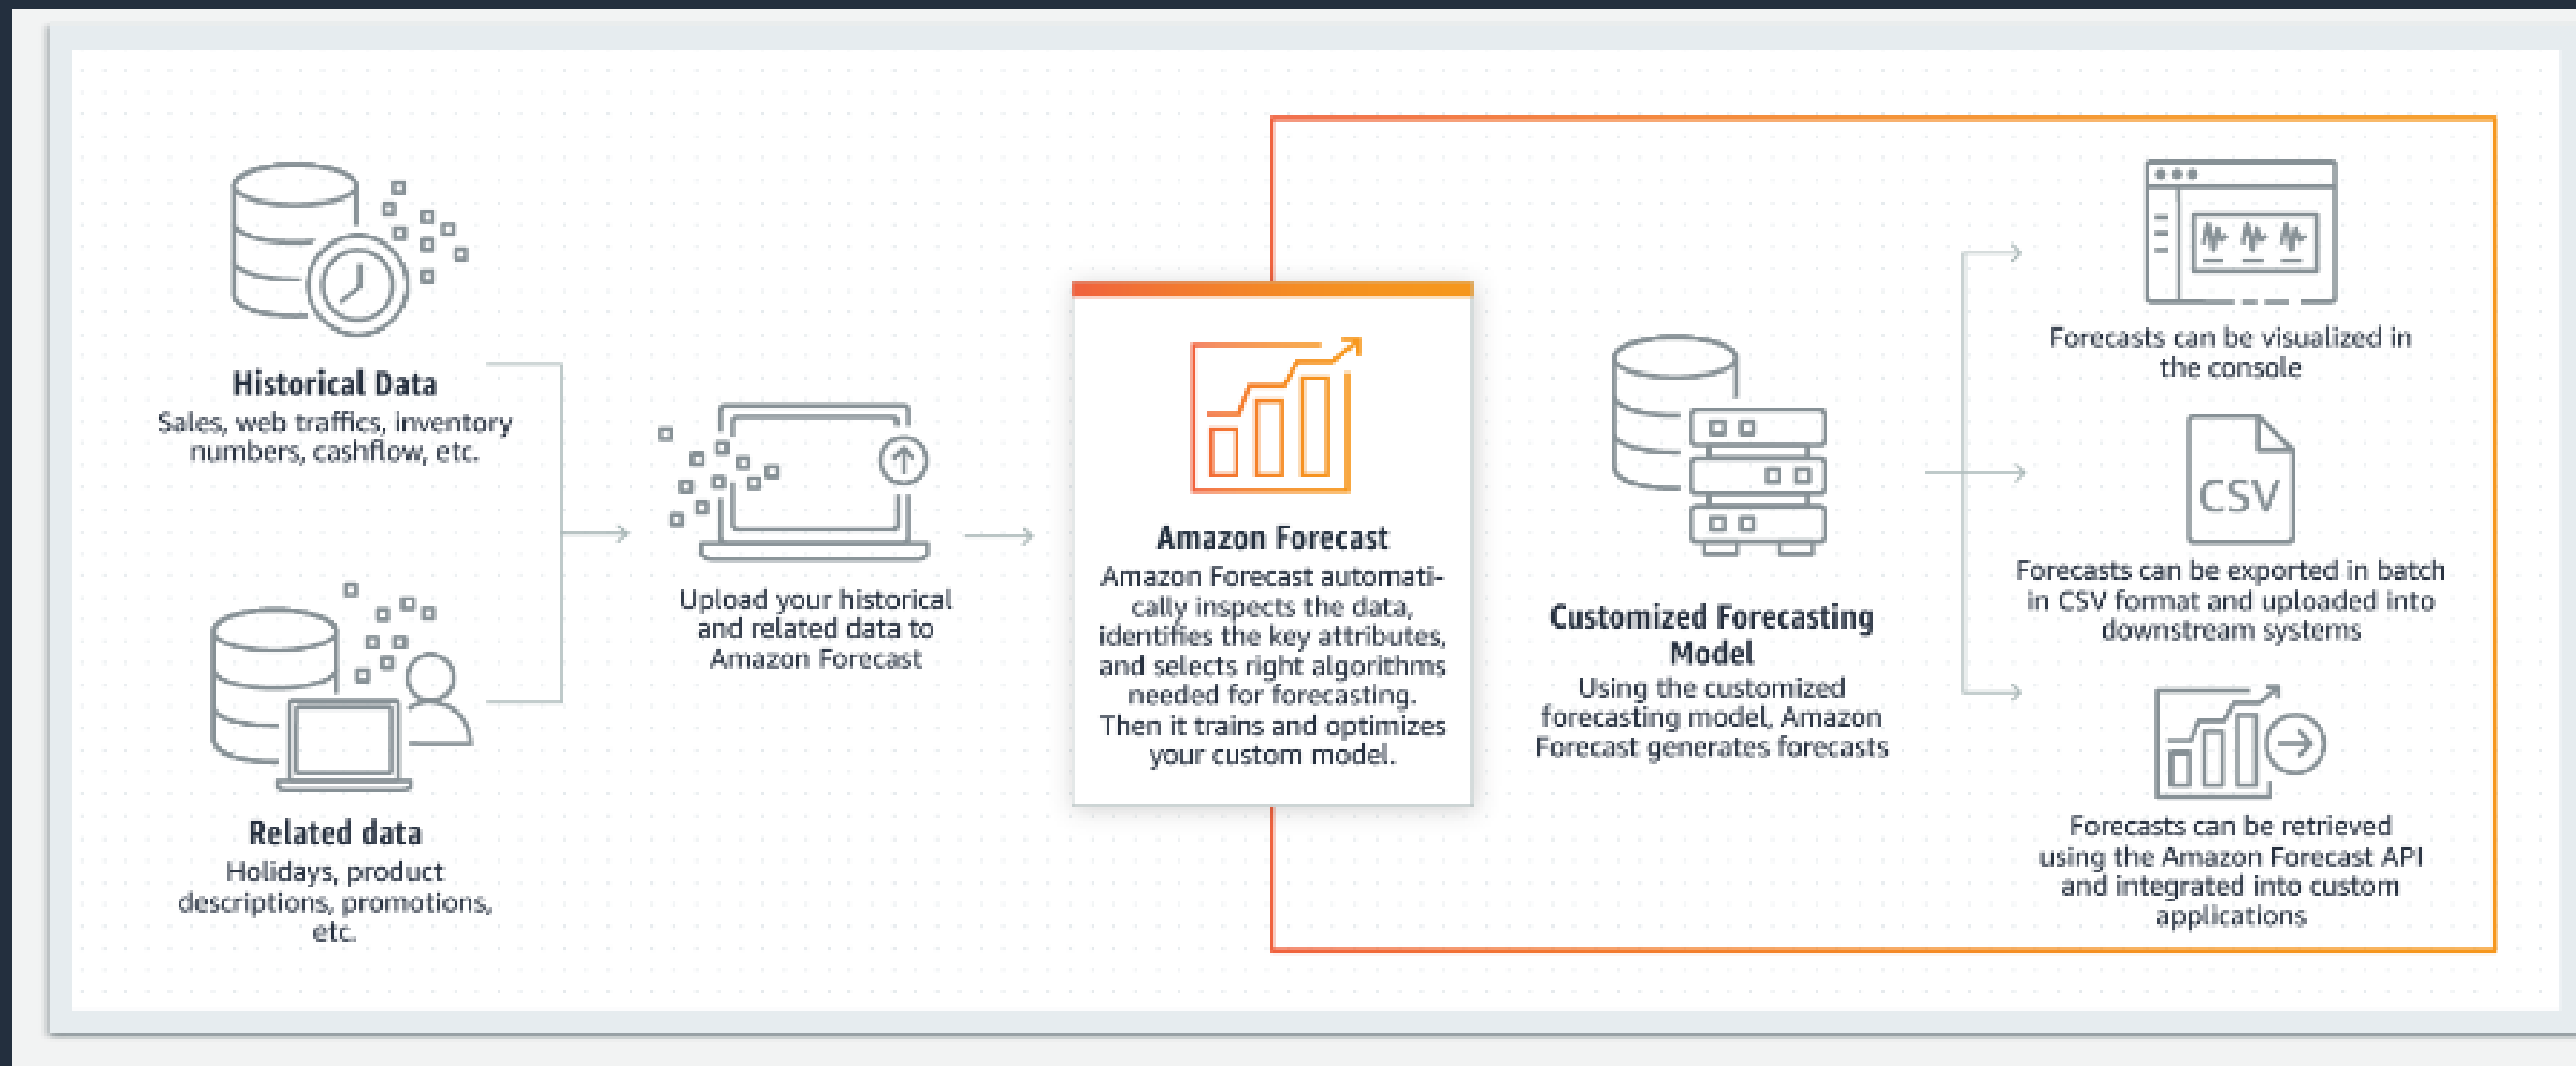
Related time-series

Related time-series such as price, web hits, etc.

Item metadata

Attributes of the item such as category, genre, and brand.


How Amazon Forecast works



Amazon Forecast options

Model Training

- Auto-regressive integrated moving average (ARIMA)
- Exponential smoothing (ETS)
- Non-parametric time series (NPTS)
- Prophet
- Deep auto-regressive plus (DeepAR+)



Metric Comparison

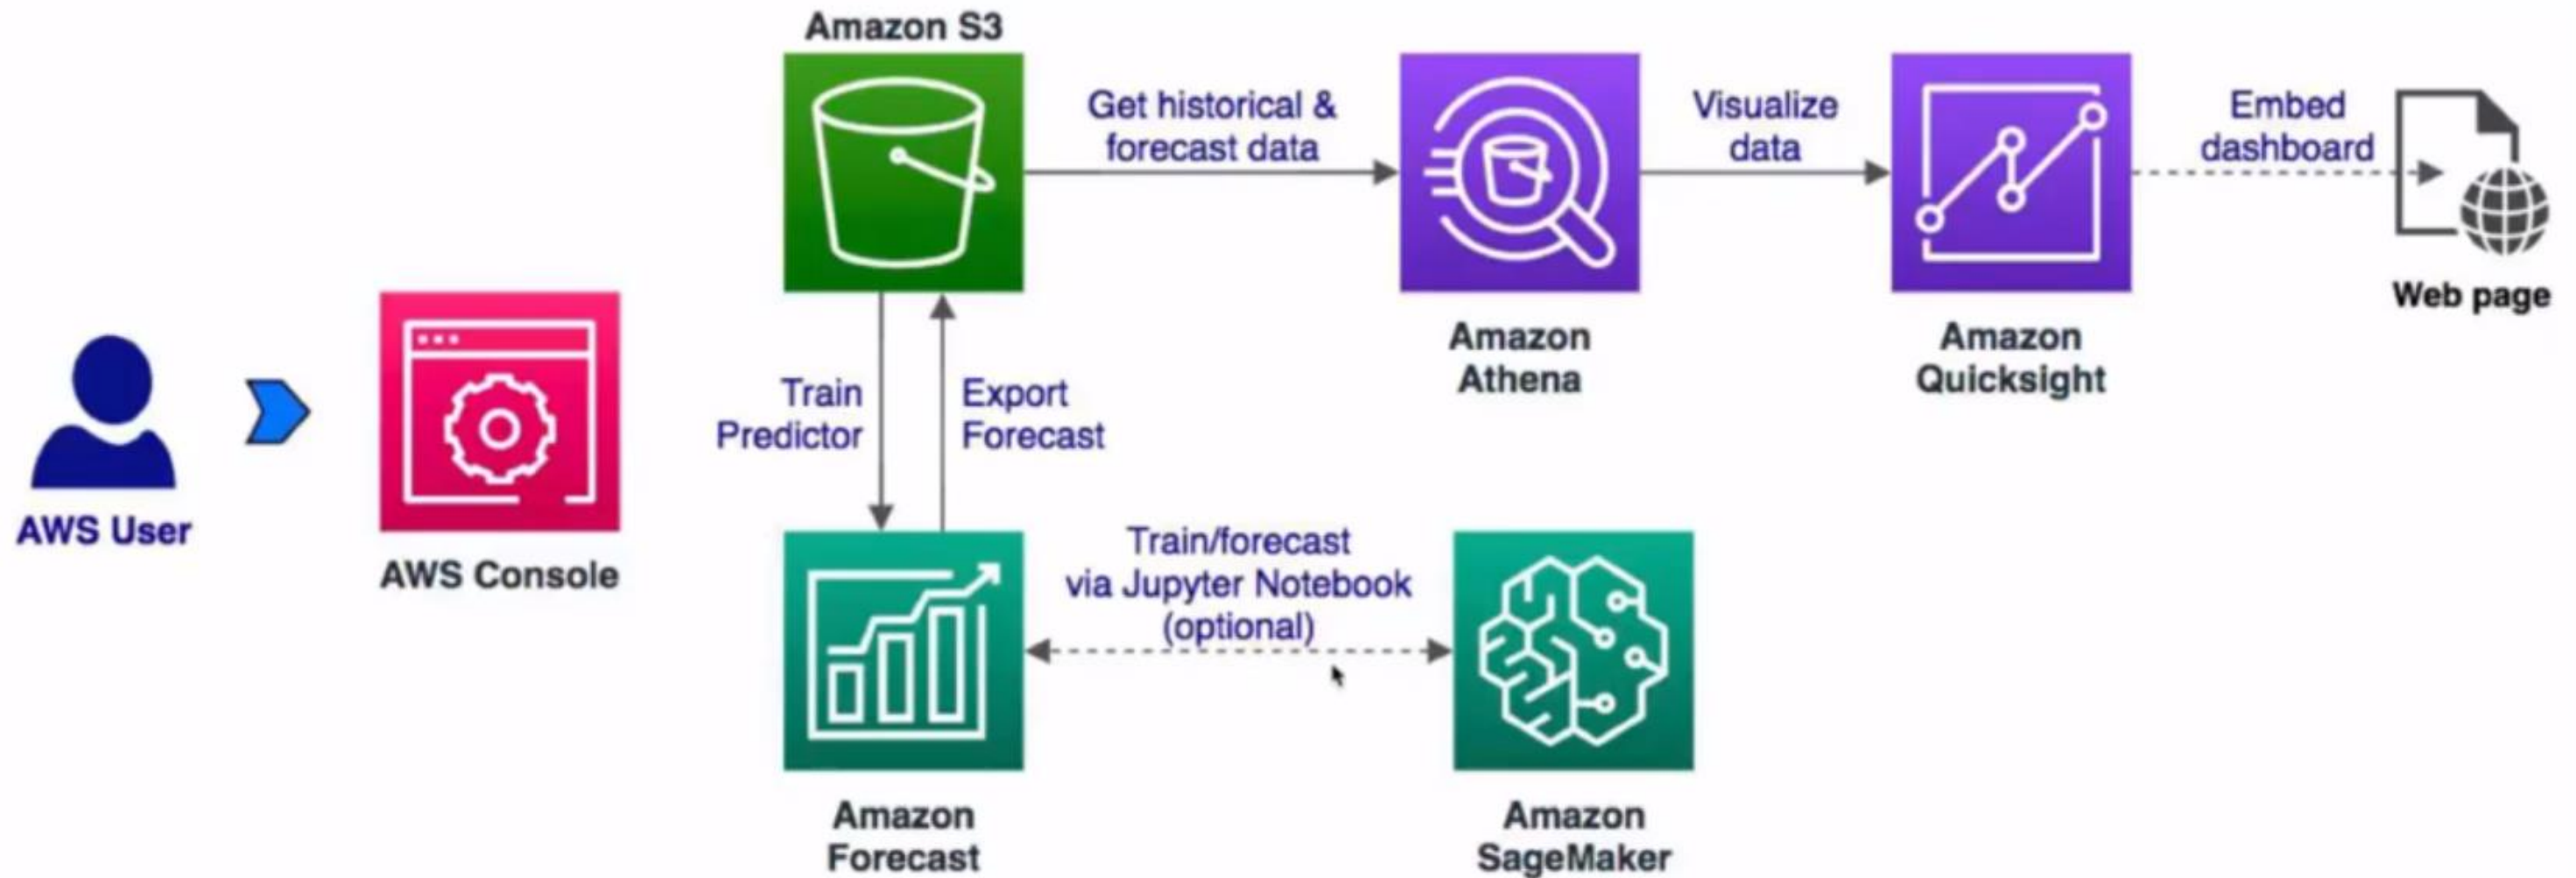
- Choose hyper parameter tuning
- Compare results with different Metrics

Application across multiple Domains

- Set your domain from console or via the API
- Upload datasets with different schemas based on the domain

Amazon Forecast demo

Amazon Forecast demo



Useful Resources for Amazon Forecast

A description of the functionalities and capabilities is here:

<https://docs.aws.amazon.com/forecast/latest/dg/what-is-forecast.html>

Also, you can find a full setup and detailed documentation of Forecast here:

<https://docs.aws.amazon.com/forecast/latest/dg/forecast.dg.pdf>

More advanced users will be able to find fully documented notebooks on this link:

<https://github.com/aws-samples/amazon-forecast-samples>

Thank You!

