

# Predicting and Optimizing Cloud Infrastructure Costs

Using FLAML, AutoML, Prophet, ARIMA & Isolation Forest

**Sai Teja Sunku**

MSDS692 – Data Science Practicum 1 | Regis University



# PROBLEM STATEMENT



**\$ 600B+**

Global cloud  
spend (2025)



**~30%**

Estimated  
wasted spend



**73%**

Orgs say cost is  
top cloud  
challenge

Cloud bills are unpredictable, difficult to forecast and often inefficient

**GOAL:** Build ML models to predict costs, forecast cost and identify anomalies

# Dataset

## AWS EC2 Pricing



**7,260 rows**  
**15 features**

**VCPUs, memory, network, region**  
**OS, Pricing model**  
**18 instance families**  
**5 regions 4 pricing models**  
**Used for: Regression, anomaly**  
**detection, cost optimization**


## AZURE Billing



**93,605 rows**  
**11 features**


**50 services, 89 days**  
**\$69- \$170/day**  
**Dec 2022- Mar 2023**  
**Real billing data**  
**Used for: Time series, cost anomaly**  
**detection**

# Methodology



**FLAML AutoML**

**Automated model  
selection &  
hyperparameter  
tuning**




**Prophet + ARIMA**

**Time series  
forecasting with  
trend detection**



**Isolation Forest**

**Anomaly detection  
for pricing & cost  
spikes**



**Cost Optimization**

**Savings analysis &  
efficiency ranking**

# EDA: Key Findings

## Right Skewed Costs

Media \$0.57/hr vs Mean \$1.43/hr, most instances are small, some expensive outliers increase the average

## Windows Premium

Windows instances cost 15-25% more than Linux due to licensing. Linux is recommended to reduce costs

## Regional Pricing

us-east-1 is cheapest. South America carries 10-15% premium. Region selection affects cost control

## Hardware Drives Cost

vCPUs ( $r=0.90$ ) and Memory ( $r=0.86$ ) are the strongest cost predictors. Clear linear relationships confirmed

# AutoML Regression: FLAML Results

---

Best: CatBoost

---

$R^2$  Score 0.9999

---

MAE \$0.0012/hr

---

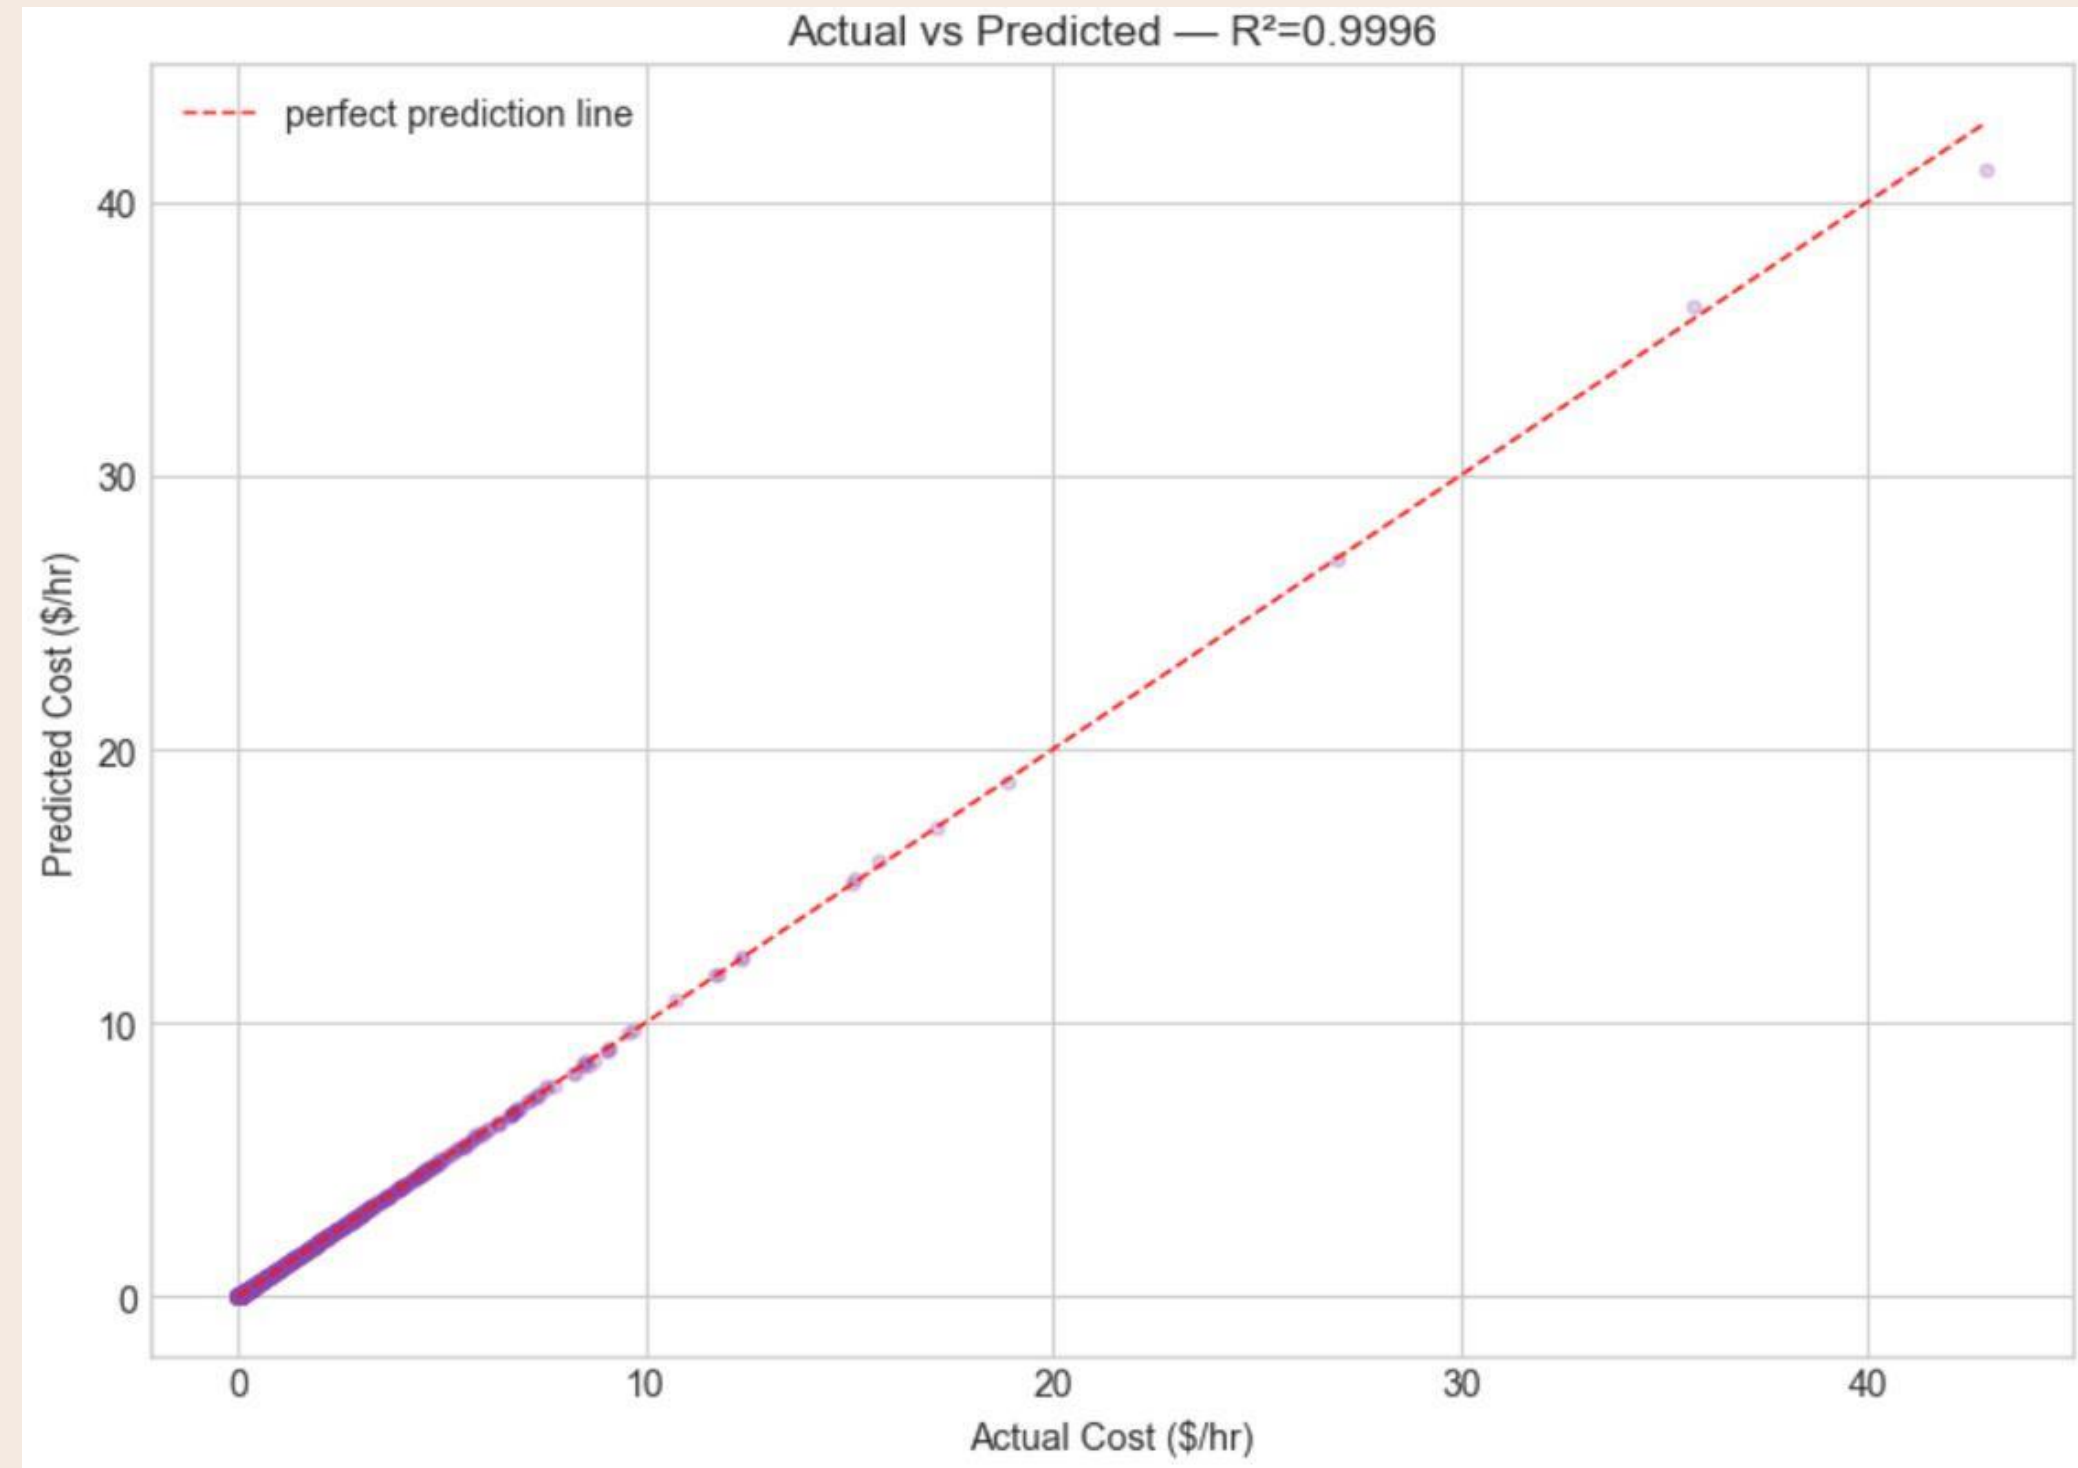
RMSE \$0.0089/hr

---

MAPE <2%

---

Target:  $R^2 > 0.80$

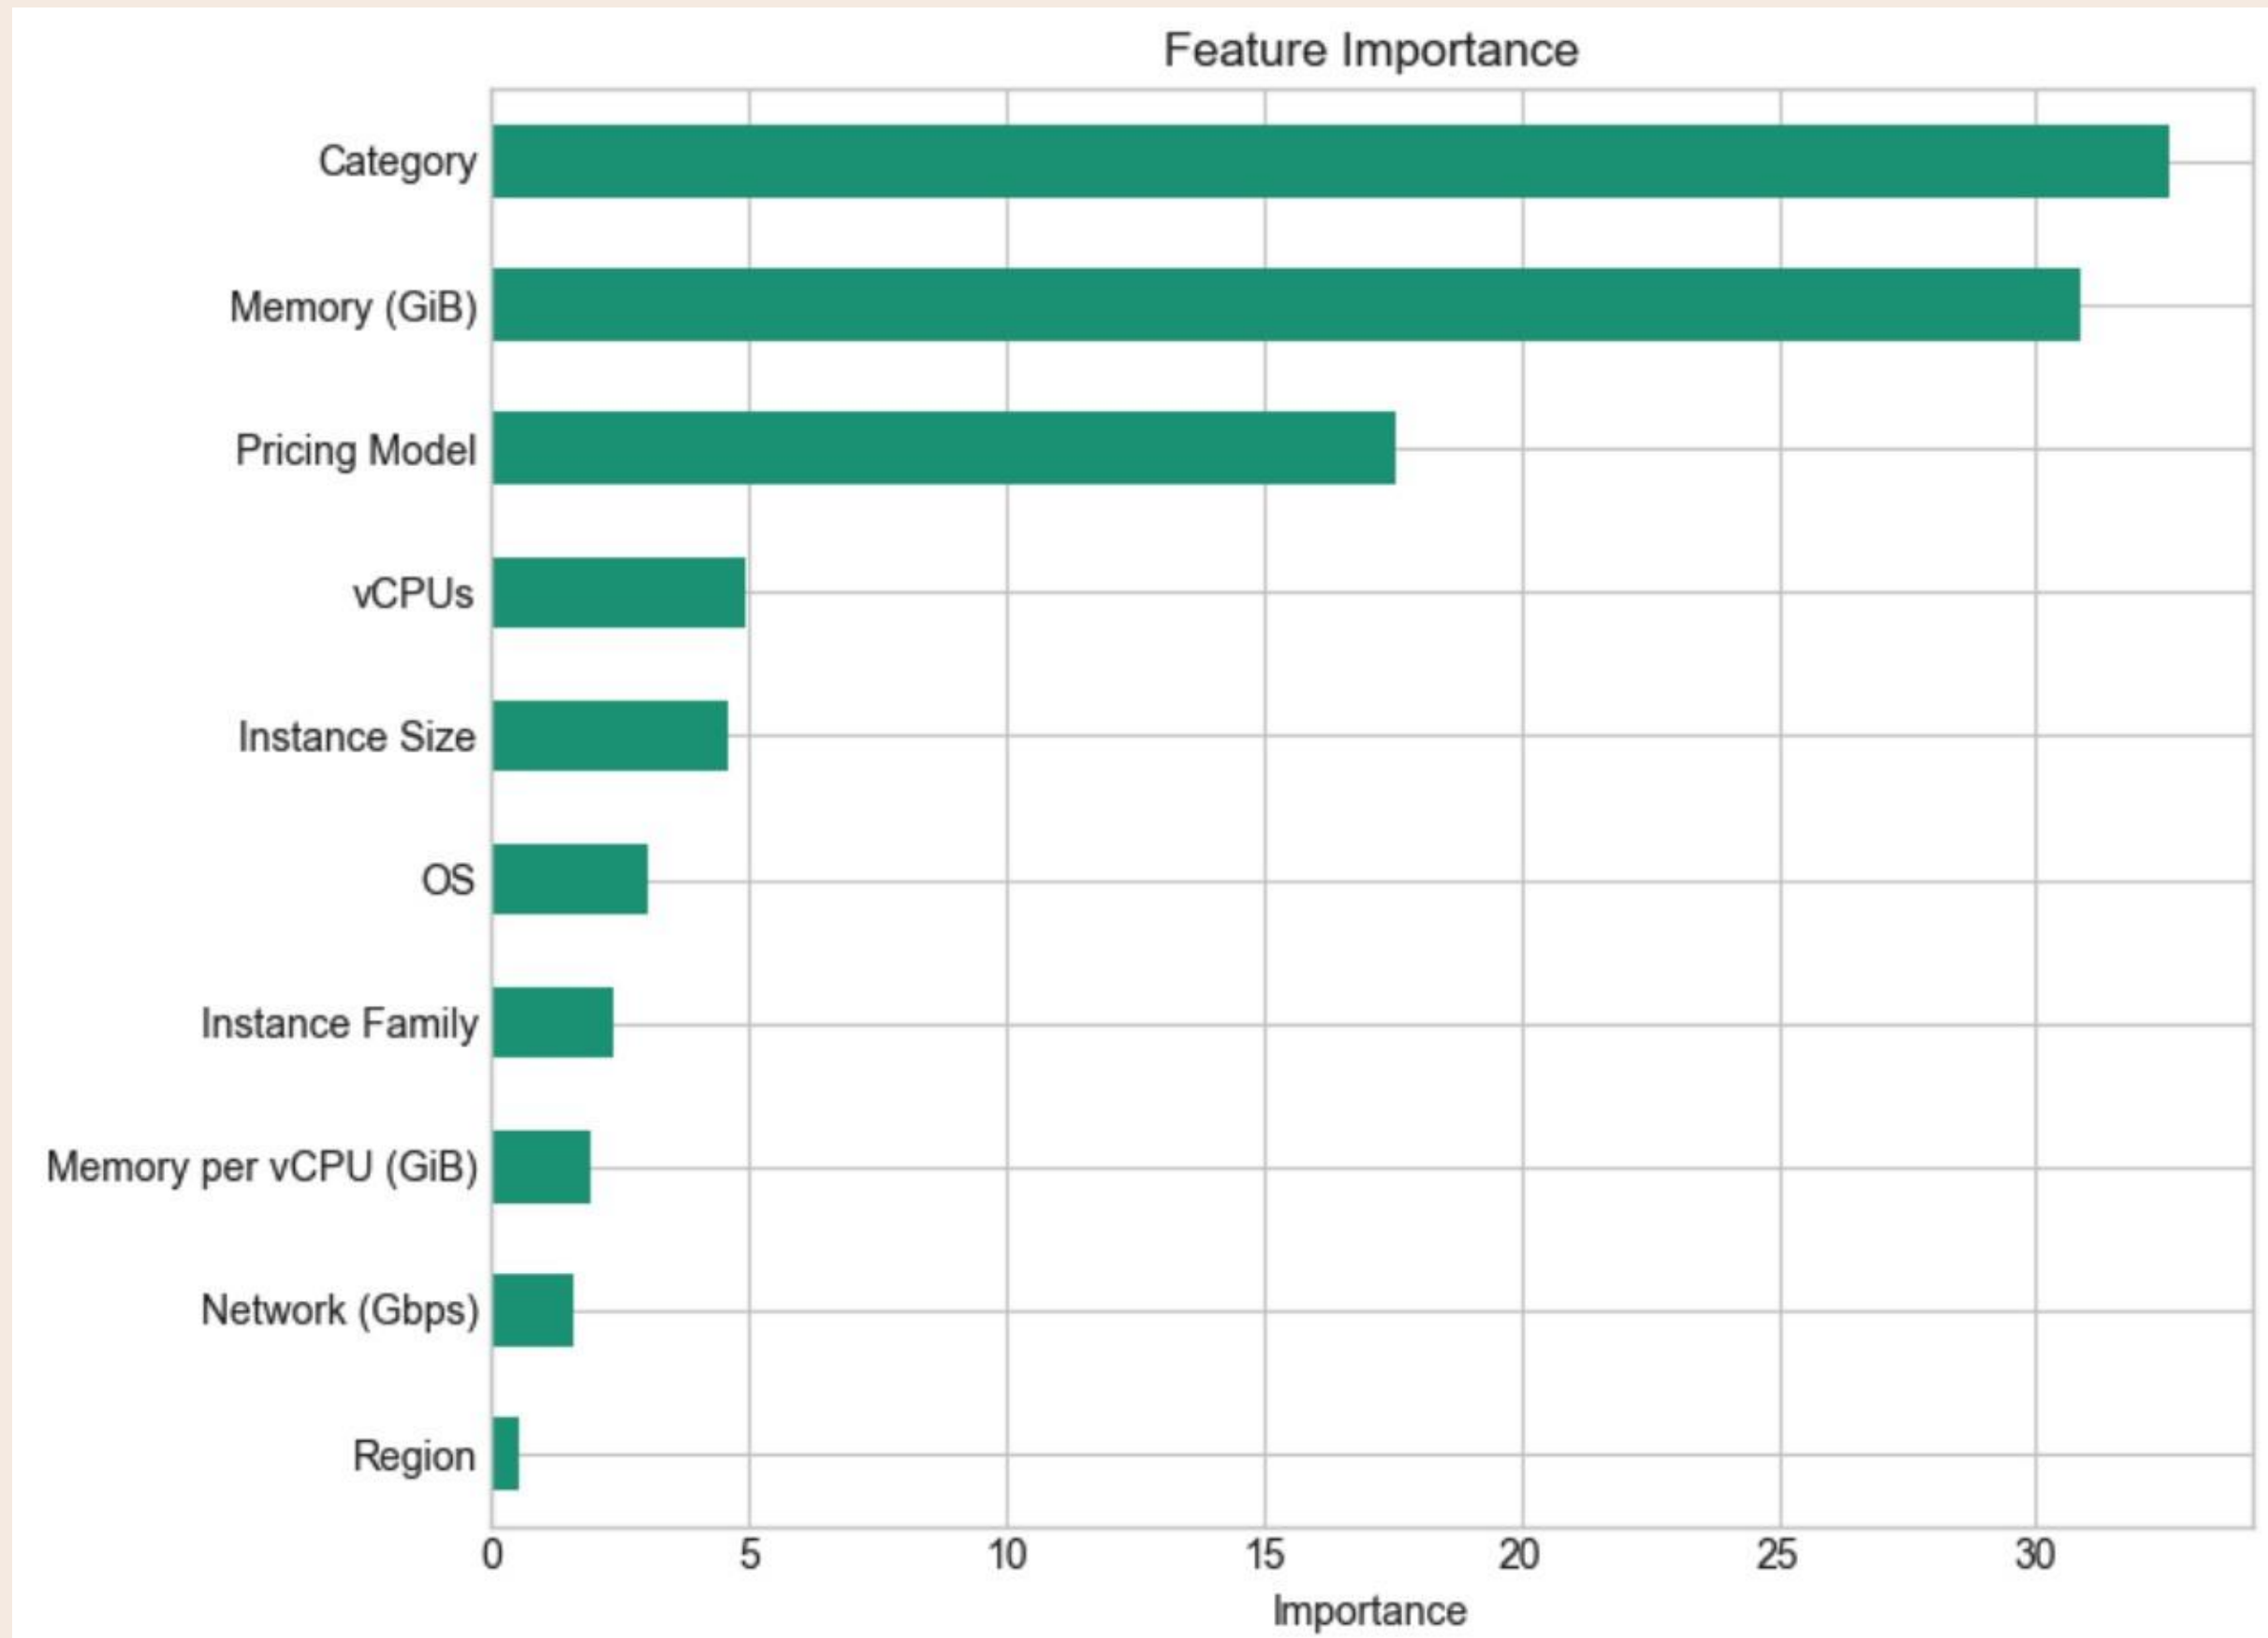


# What Drives Cloud Costs?

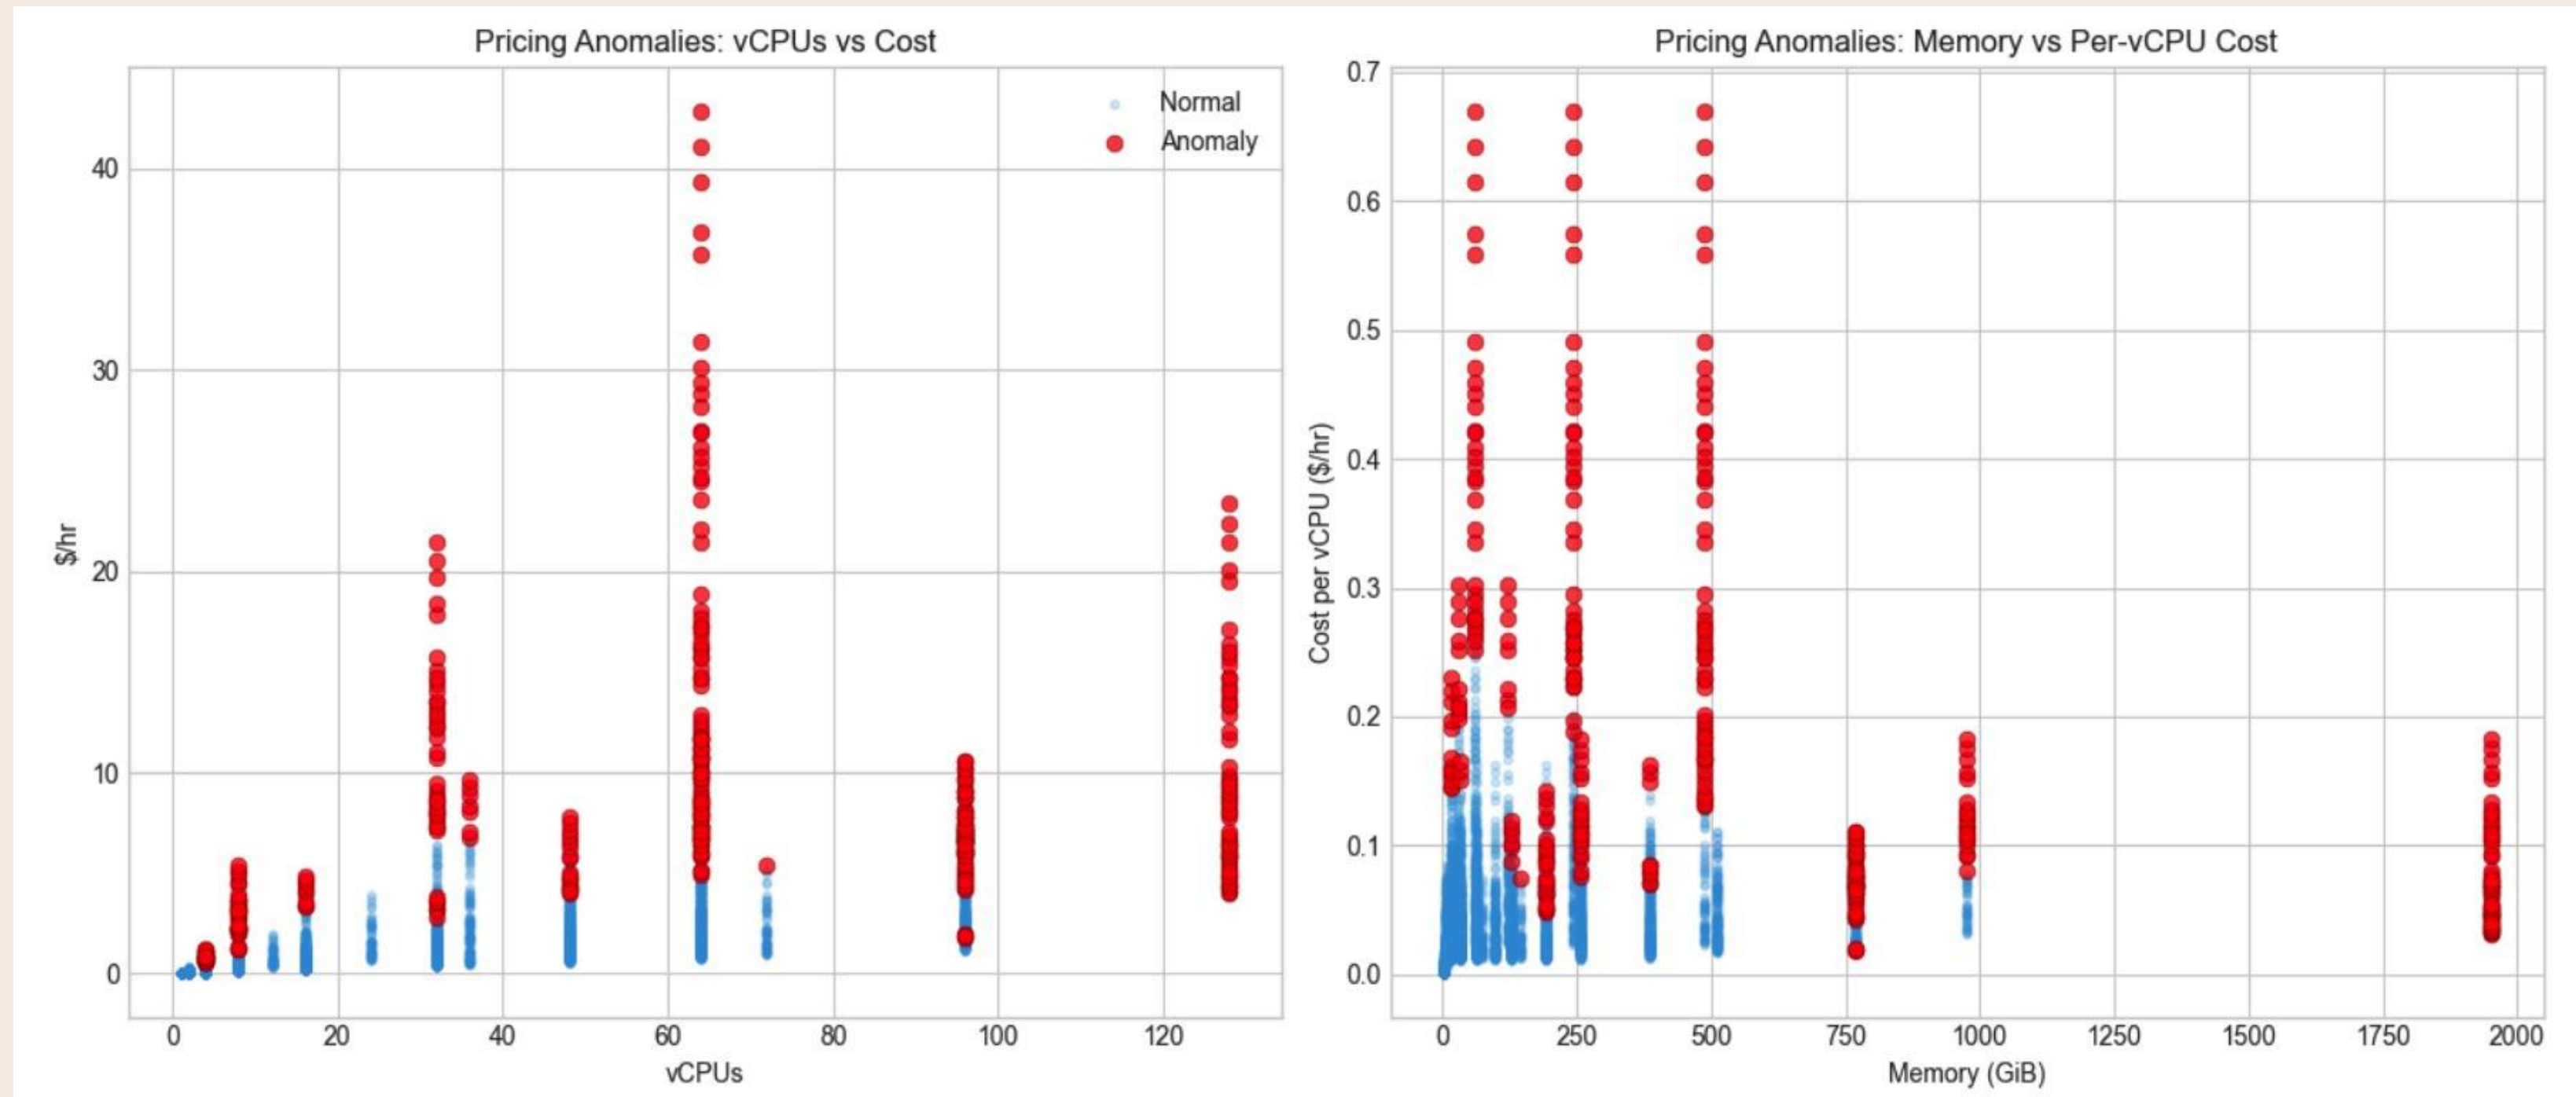
---

## Key Insights:

- Memory and vCPUs account for ~70% of cost prediction.
- Pricing Model and Instance Family contribute another ~20%.
- Region and OS have minimal impact
- Hardware specifications Determine cloud pricing



# Pricing Anomaly Detection



**Isolation Forest identified ~5 % of instances as overpriced, which costs ~9x higher than similar configurations, highlighting cost-saving opportunities**

# Cost Optimization – Savings Potential

**Reserved 1 yr**

**~ 40%**

**Savings vs On-Demand**

**Stable workloads with  
predictable usage**

**Reserved 3 yr**

**~ 60%**

**Savings vs On-demand**

**Long-term committed  
infrastructure**

**Spot Instances**

**~ 70%**

**Savings vs On-demand**

**Batch jobs, fault-tolerant  
workloads**



**For \$50K/month cloud spend, switching to reserved/spot pricing could save \$240K- \$420K annually.**

# Azure Billing: Overview

**93,605**

**Total Records**

**89 days**

**Time period**

**\$69- \$170**

**Daily Range**

**50 Unique**

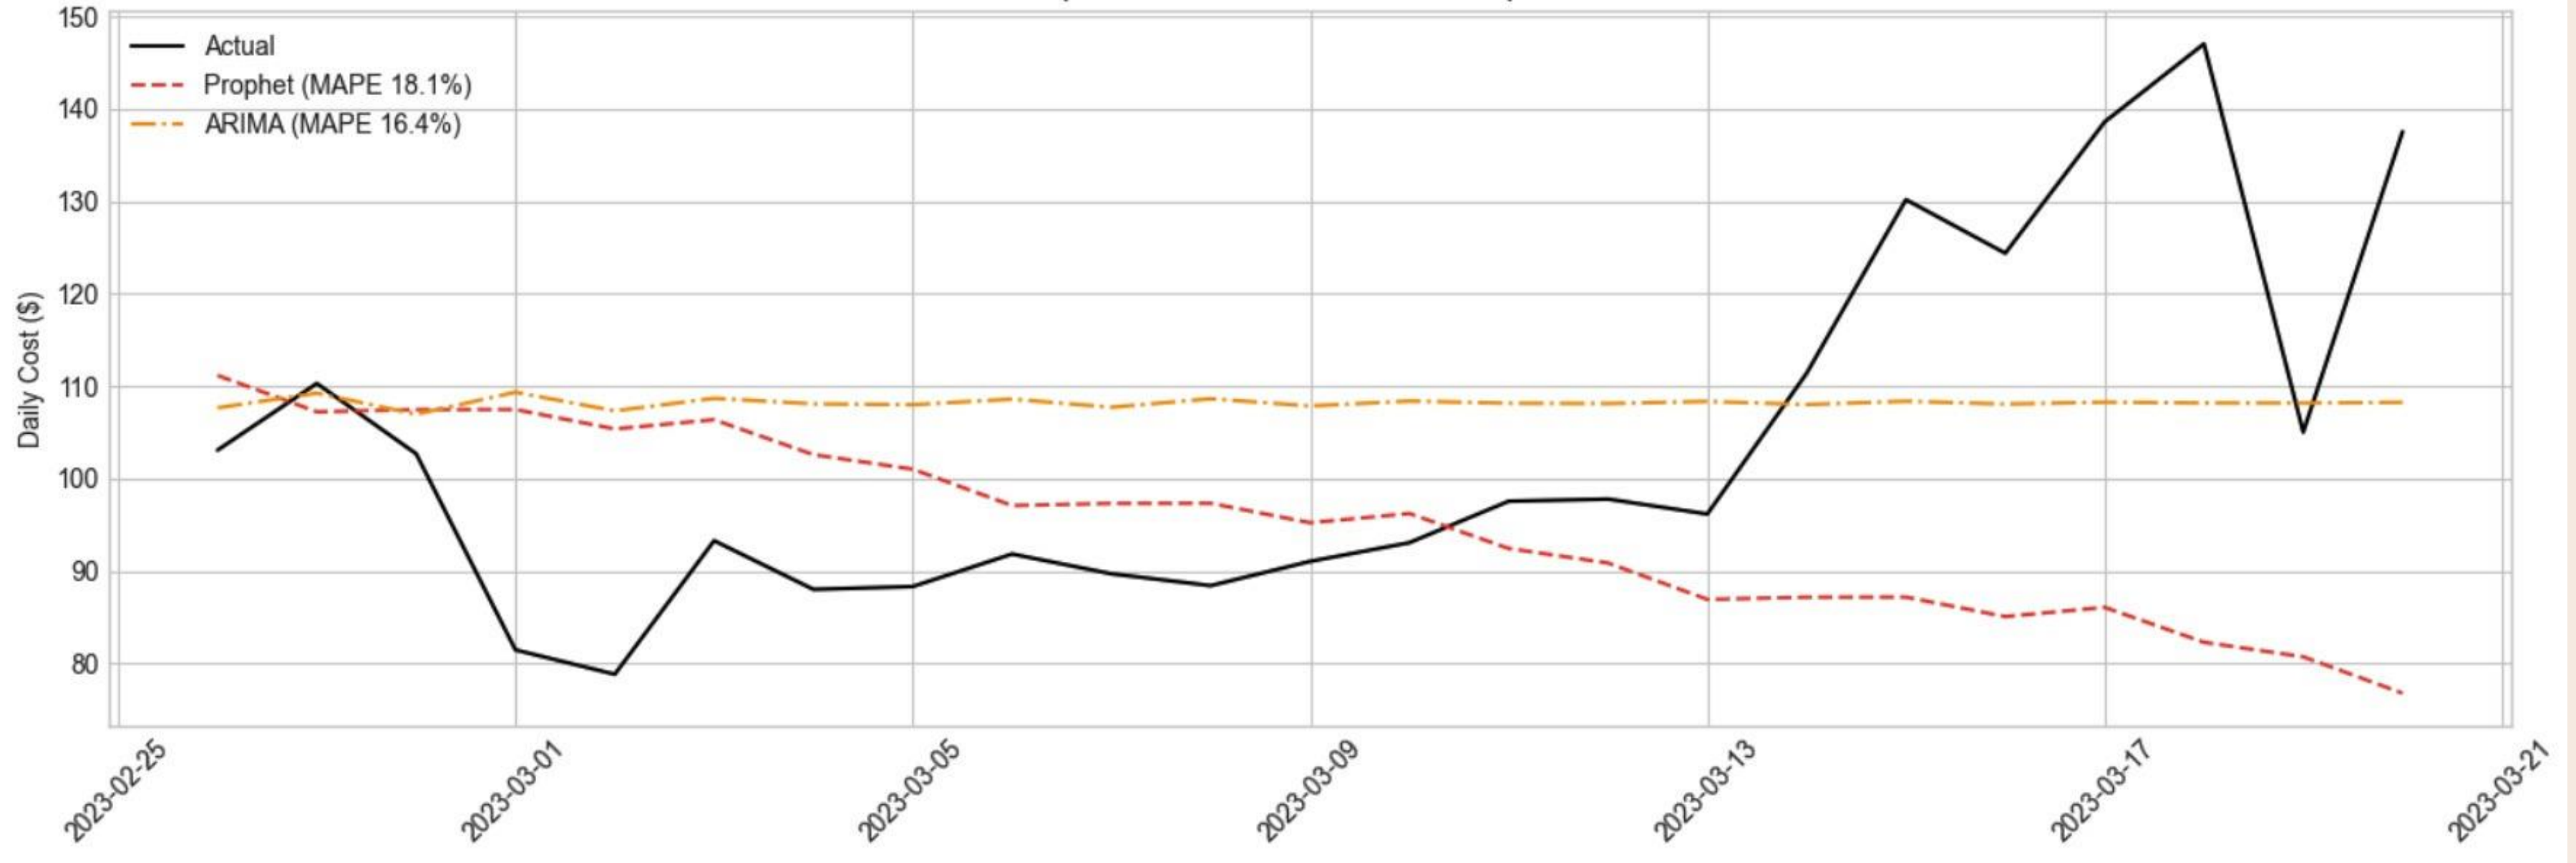
**Services**

## **Key observations:**

- Clear upward trend in spending, rising from \$70/day to \$170/day over 3 month period
- Storage and Bandwidth contribute to total spend (top 2 services)
- Virtual Machines are the third largest cost center
- High volatility in some services indicates anomaly detection opportunities

# Time Series Forecasting: Prophet Vs ARIMA

Prophet vs ARIMA — Forecast Comparison



## Prophet

MAPE: ~ 8-16%,

Captures trend + weekly patterns

## ARIMA (2,1,2)

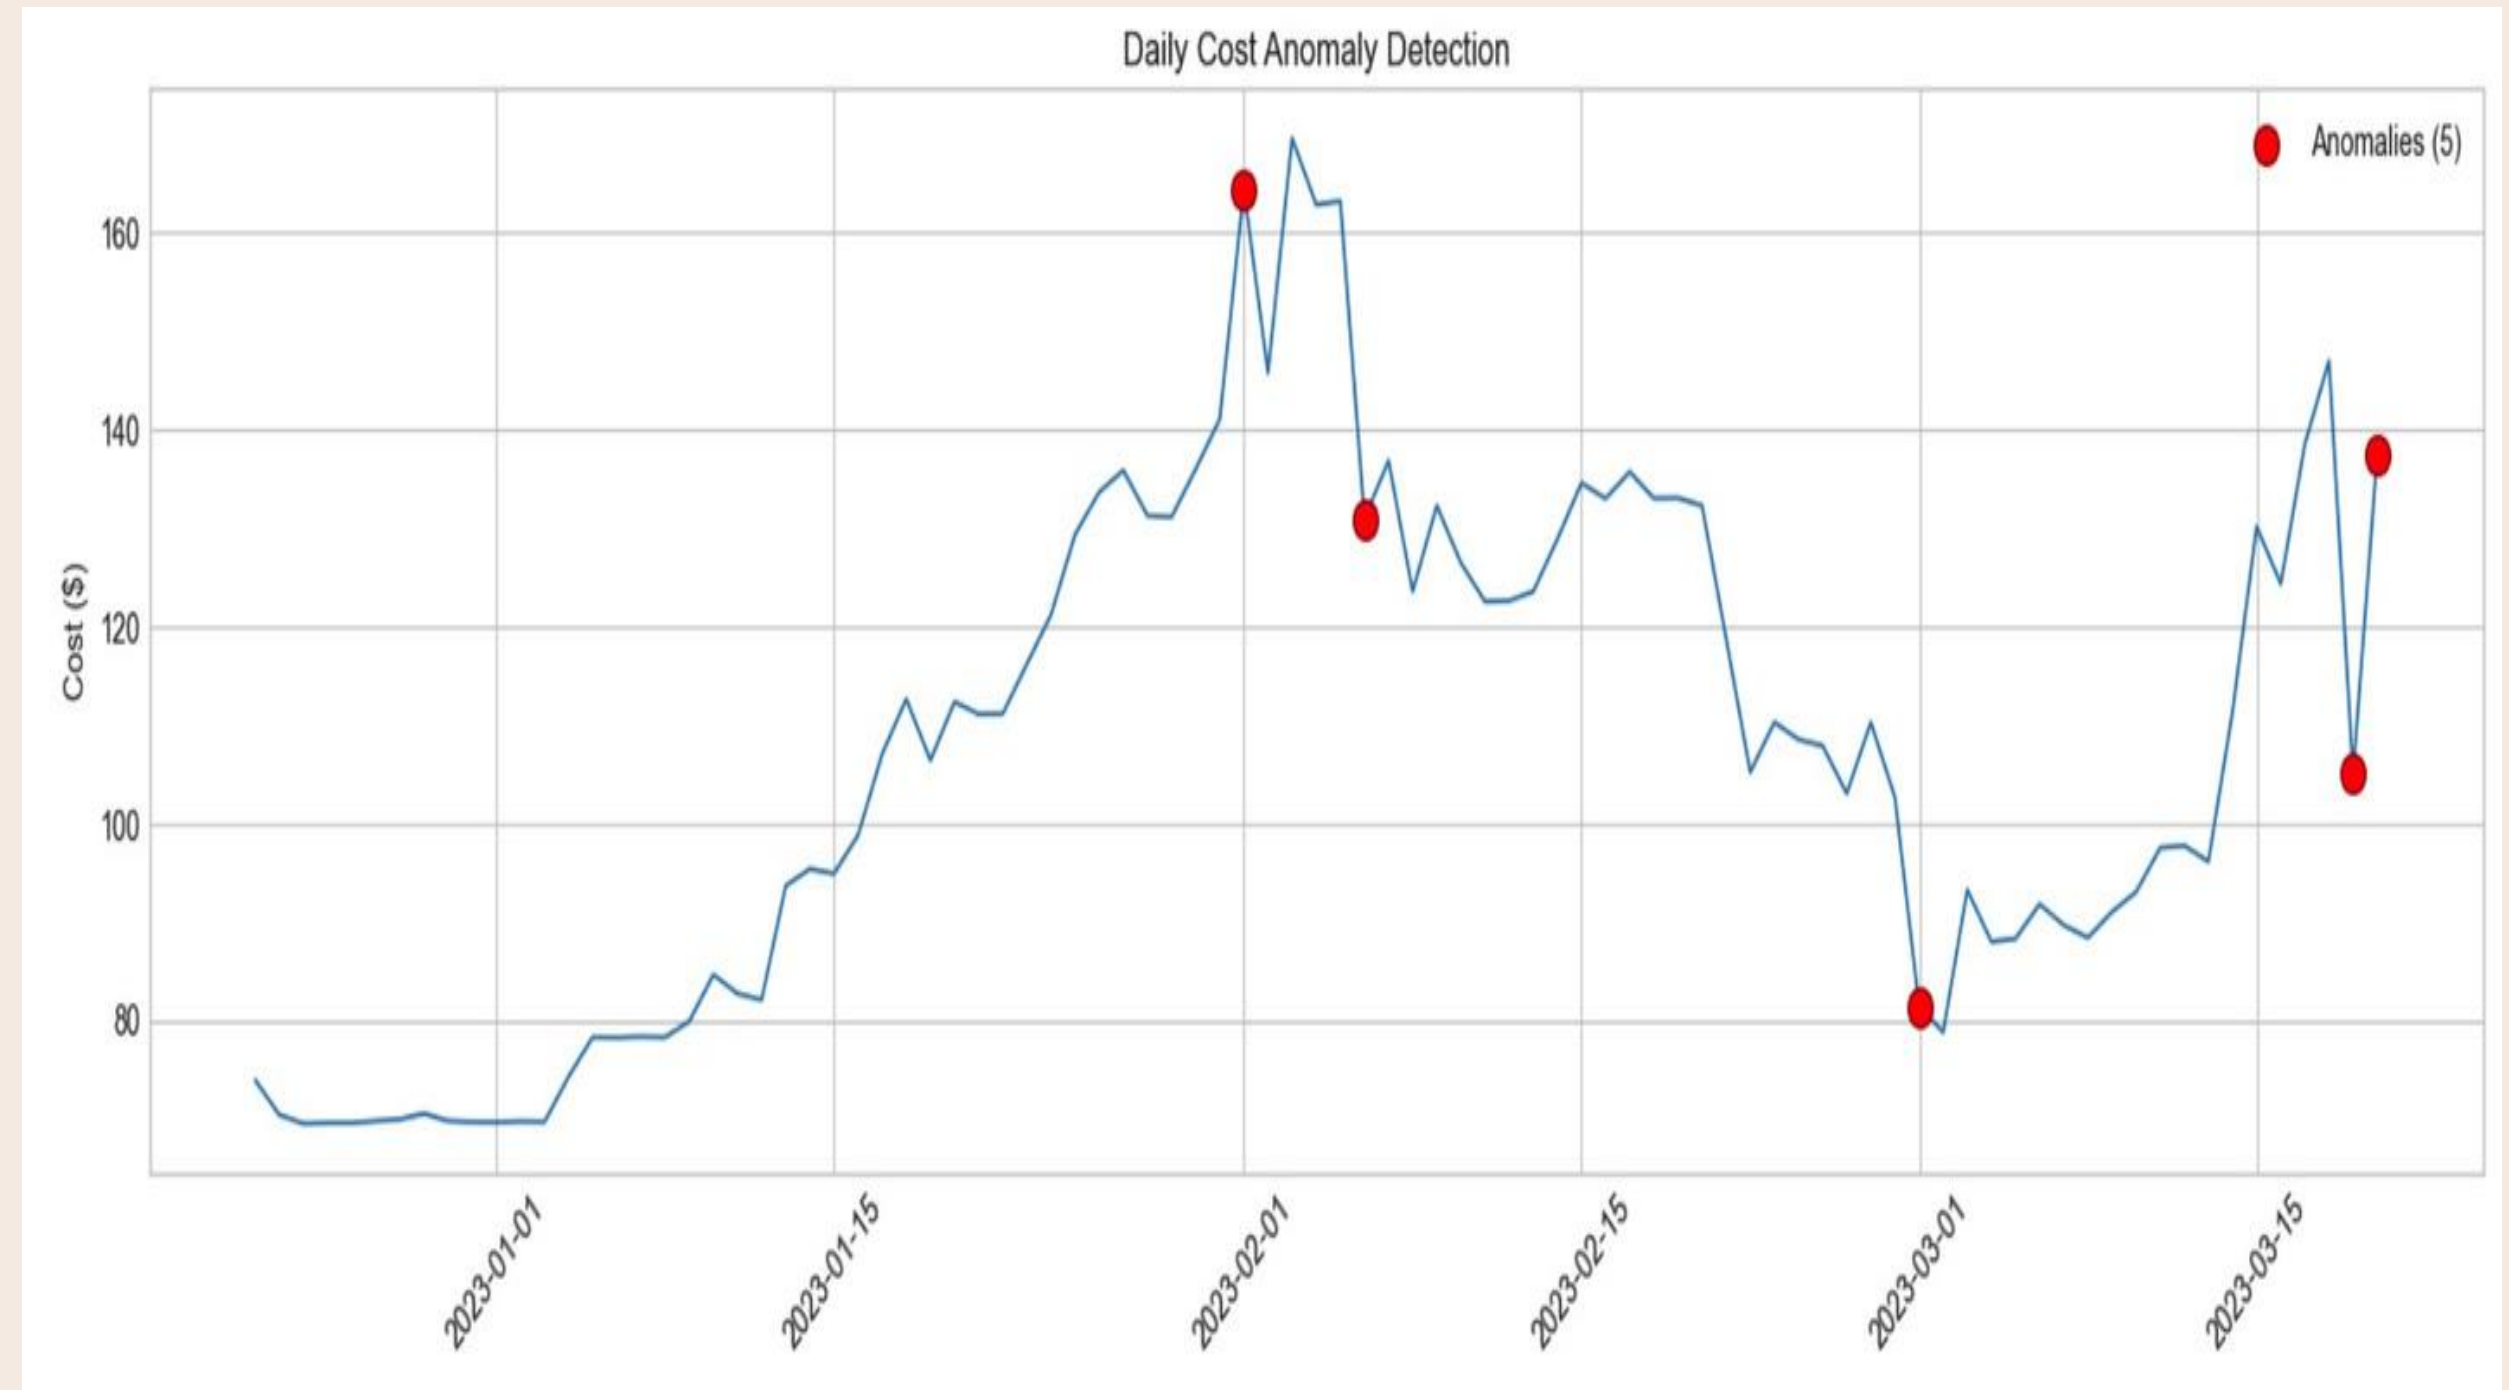
MAPE: ~6-12%

Differencing handles trend

# Cost Anomaly Detection


---

- Isolation Forest flags ~5% of days as anomalous cost spikes, useful for automated alerting.




# Limitations

- 89 days limits seasonal pattern learning
- Azure billing used as proxy for AWS
- Pricing catalog is static, real costs depend on usage
- Unsupervised anomaly detection i.e no labeled ground truth.




# Future Work

- Integrate the AWS Cost Explorer API to enable real-time cost tracking
- Develop an automated cost anomaly alerting and detection pipeline
- Upgrade to multi-cloud unified dashboard
- Leverage 12+ months of data to capture seasonality patterns



# Key Takeaways

- Cloud costs can be accurately predicted from instance specifications, FLAML  $R^2 > 0.99$
- Short-term forecasts perform well even with limited data, ARIMA MAPE ~6-12%
- Anomaly detection identifies real patterns. Also, reveals overprices and unusual spending days
- Switching from On-Demand to Reserved/Spot instances can reduce costs by 40-70%



# References:

<https://pin.it/3OLrvfqjm> -Figure 1

Amazon Web Services LOGO – Figure 2

Microsoft Azure LOGO – Figure 3

[https://dzen.ru/a/Z\\_YrrQevQkW4h75A?utm\\_medium=organic&utm\\_source=yandexsmartcamera](https://dzen.ru/a/Z_YrrQevQkW4h75A?utm_medium=organic&utm_source=yandexsmartcamera)

- Figure 4

<https://pin.it/Ro4Ur7QMj> - Figure 5

<https://pin.it/jTpDyBQEU> - Figure 6

**Thank You**