TEAM 25

Electricity Demand Forecasting with News Sentiment Analysis and Outlier Detection

Arsh Arora, Nikunj Agarwal, Soumya Pandey

MLPR Final Presentation

What is Electricity Bidding?

A market mechanism where participants, submit bids indicating how much electricity they are willing to supply or demand, and at what price.

Suppliers and consumers place their bids, specifying the quantity of electricity and the price they are willing to offer or accept.

Electricity bidding is intended to promote efficiency in electricity markets by allowing prices to reflect supply and demand conditions.

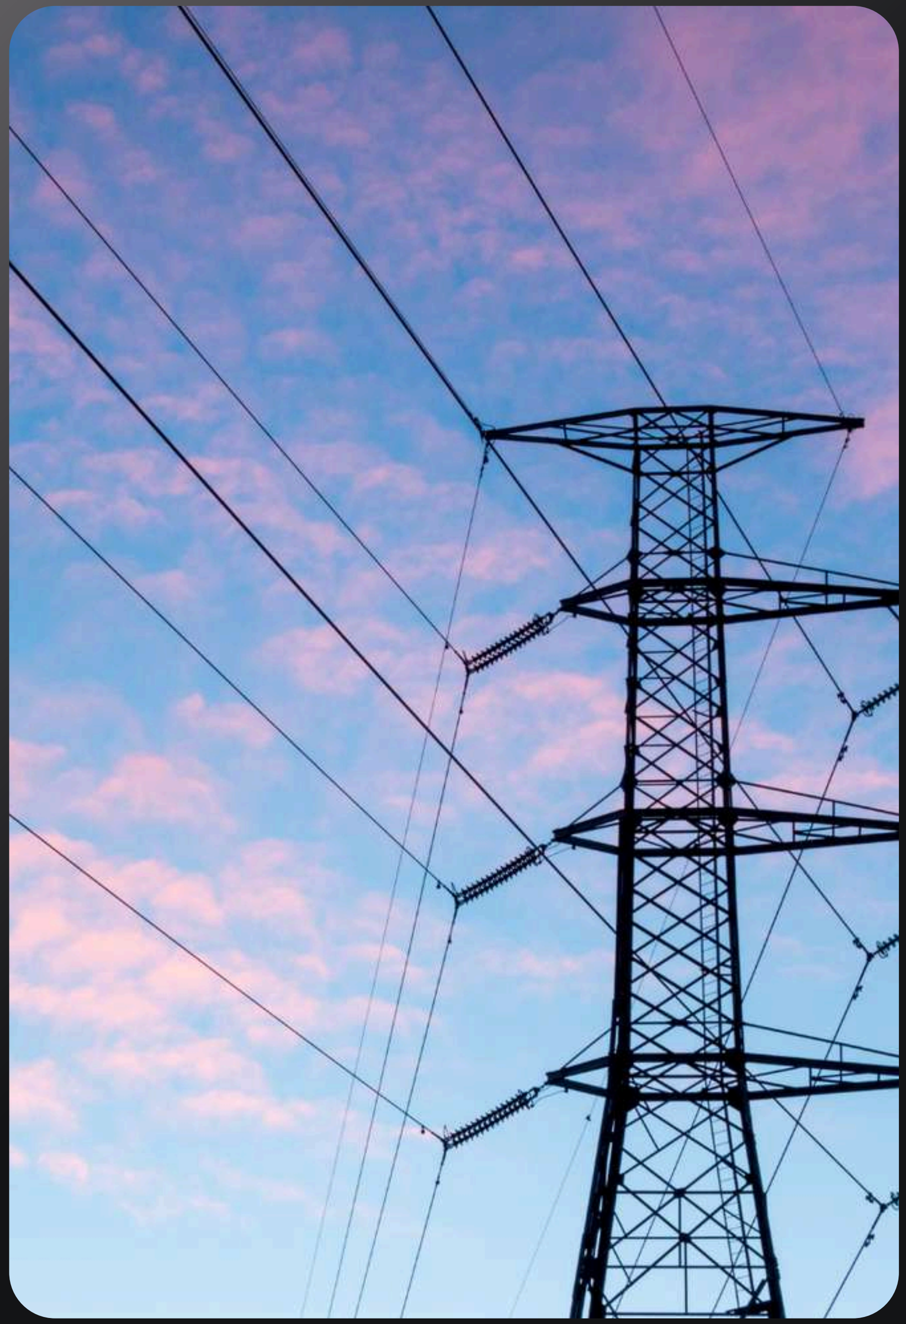

Overview

1 Data Collection
- Weather & Energy Data from Kaggle
- GDELT GKG data extracted through API key

2 Data Split + Clean
- 15 minute time intervals
- Month, Seasonal, Random Splittings

3 Model Training boosted Optuna
- LSTM
- XGBoost

4 Validation & Outcome Evaluation
- Running Evaluation
- R Square
- RMSE

01

Problem Statement

To what extent can integrating news sentiment analysis and outlier detection improve electricity demand forecasting for better bidding strategies.

02 Relevance

For normal Day-to-Day events, the current model works well.
But for such events, the electricity demand fluctuates a lot, leading to immense wastage/spikes, for which the model is not prepared.
What if we had an already trained model ready for such shifts?

03 Deliverables

Predicted Energy that will be utilised by a city in the next 15 minutes.

Why Load Forecasting Still Fails in 2025

Soaring renewable penetration → higher volatility

Weather-only models miss demand shocks triggered by policy & media

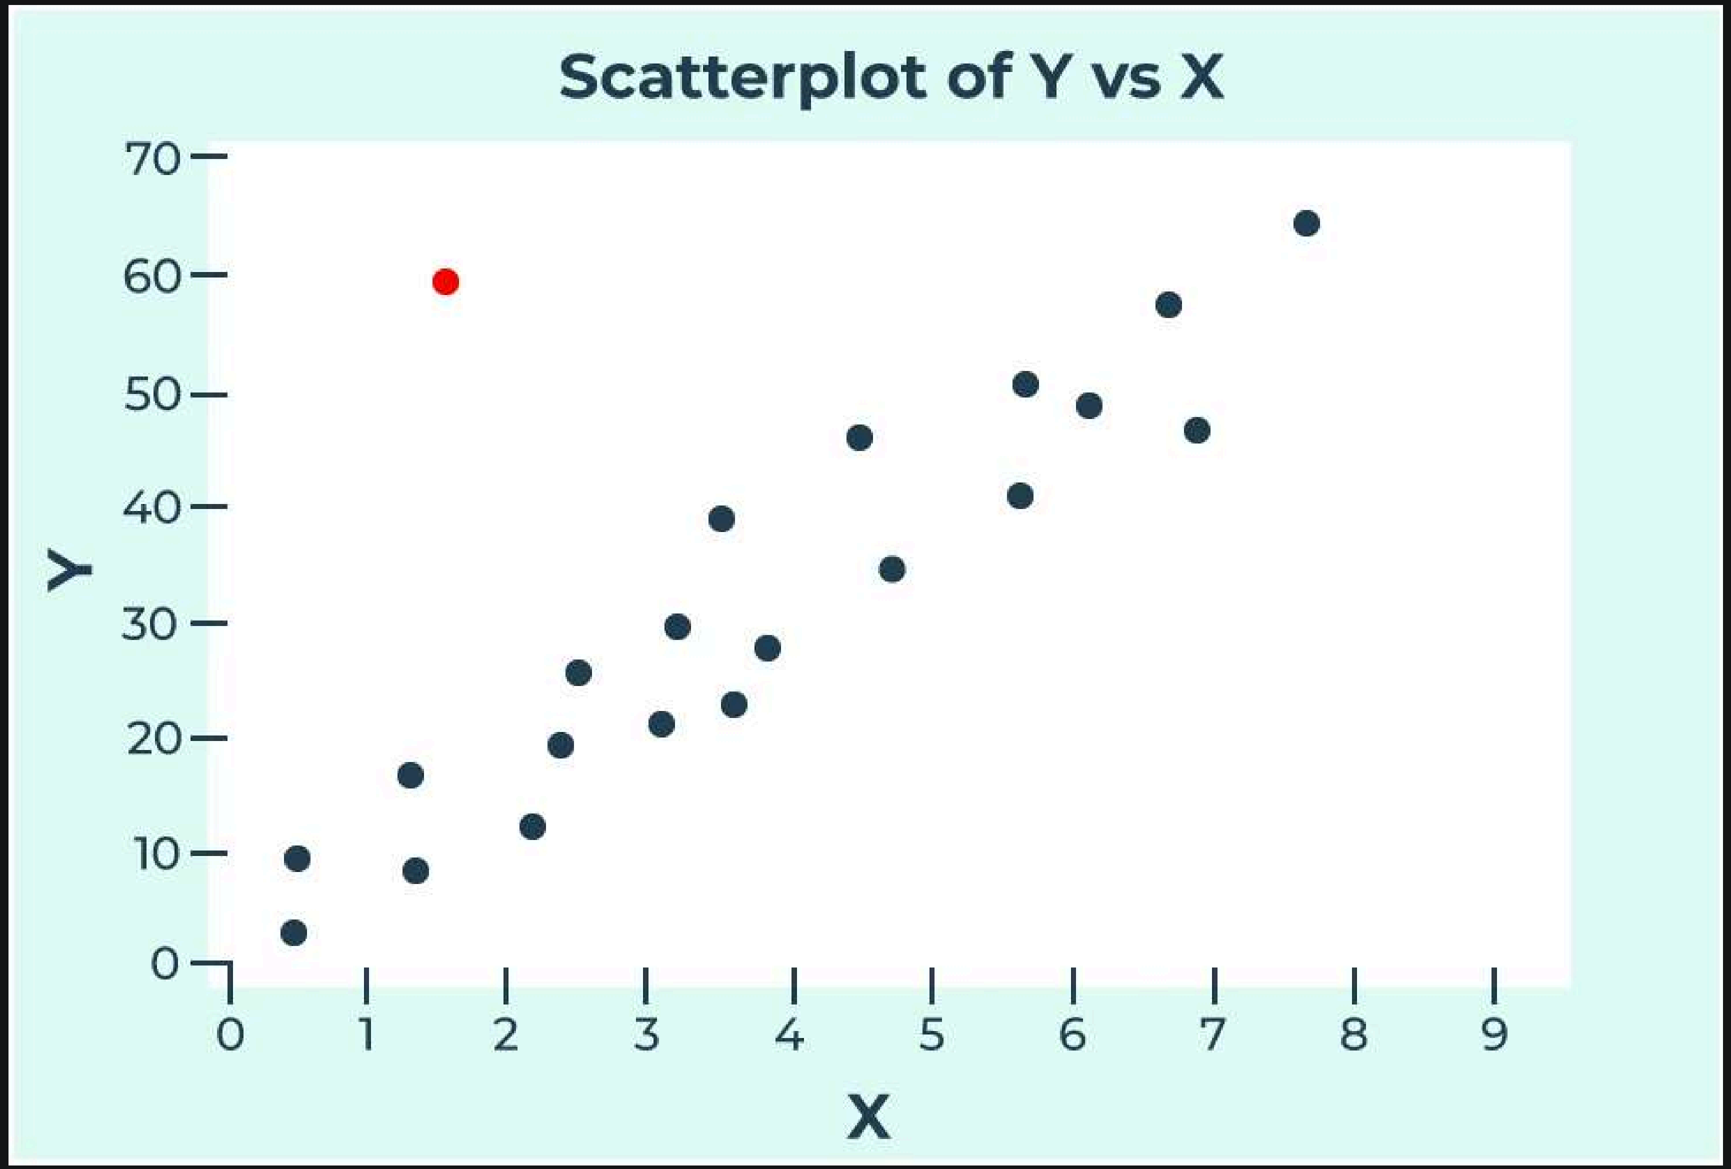

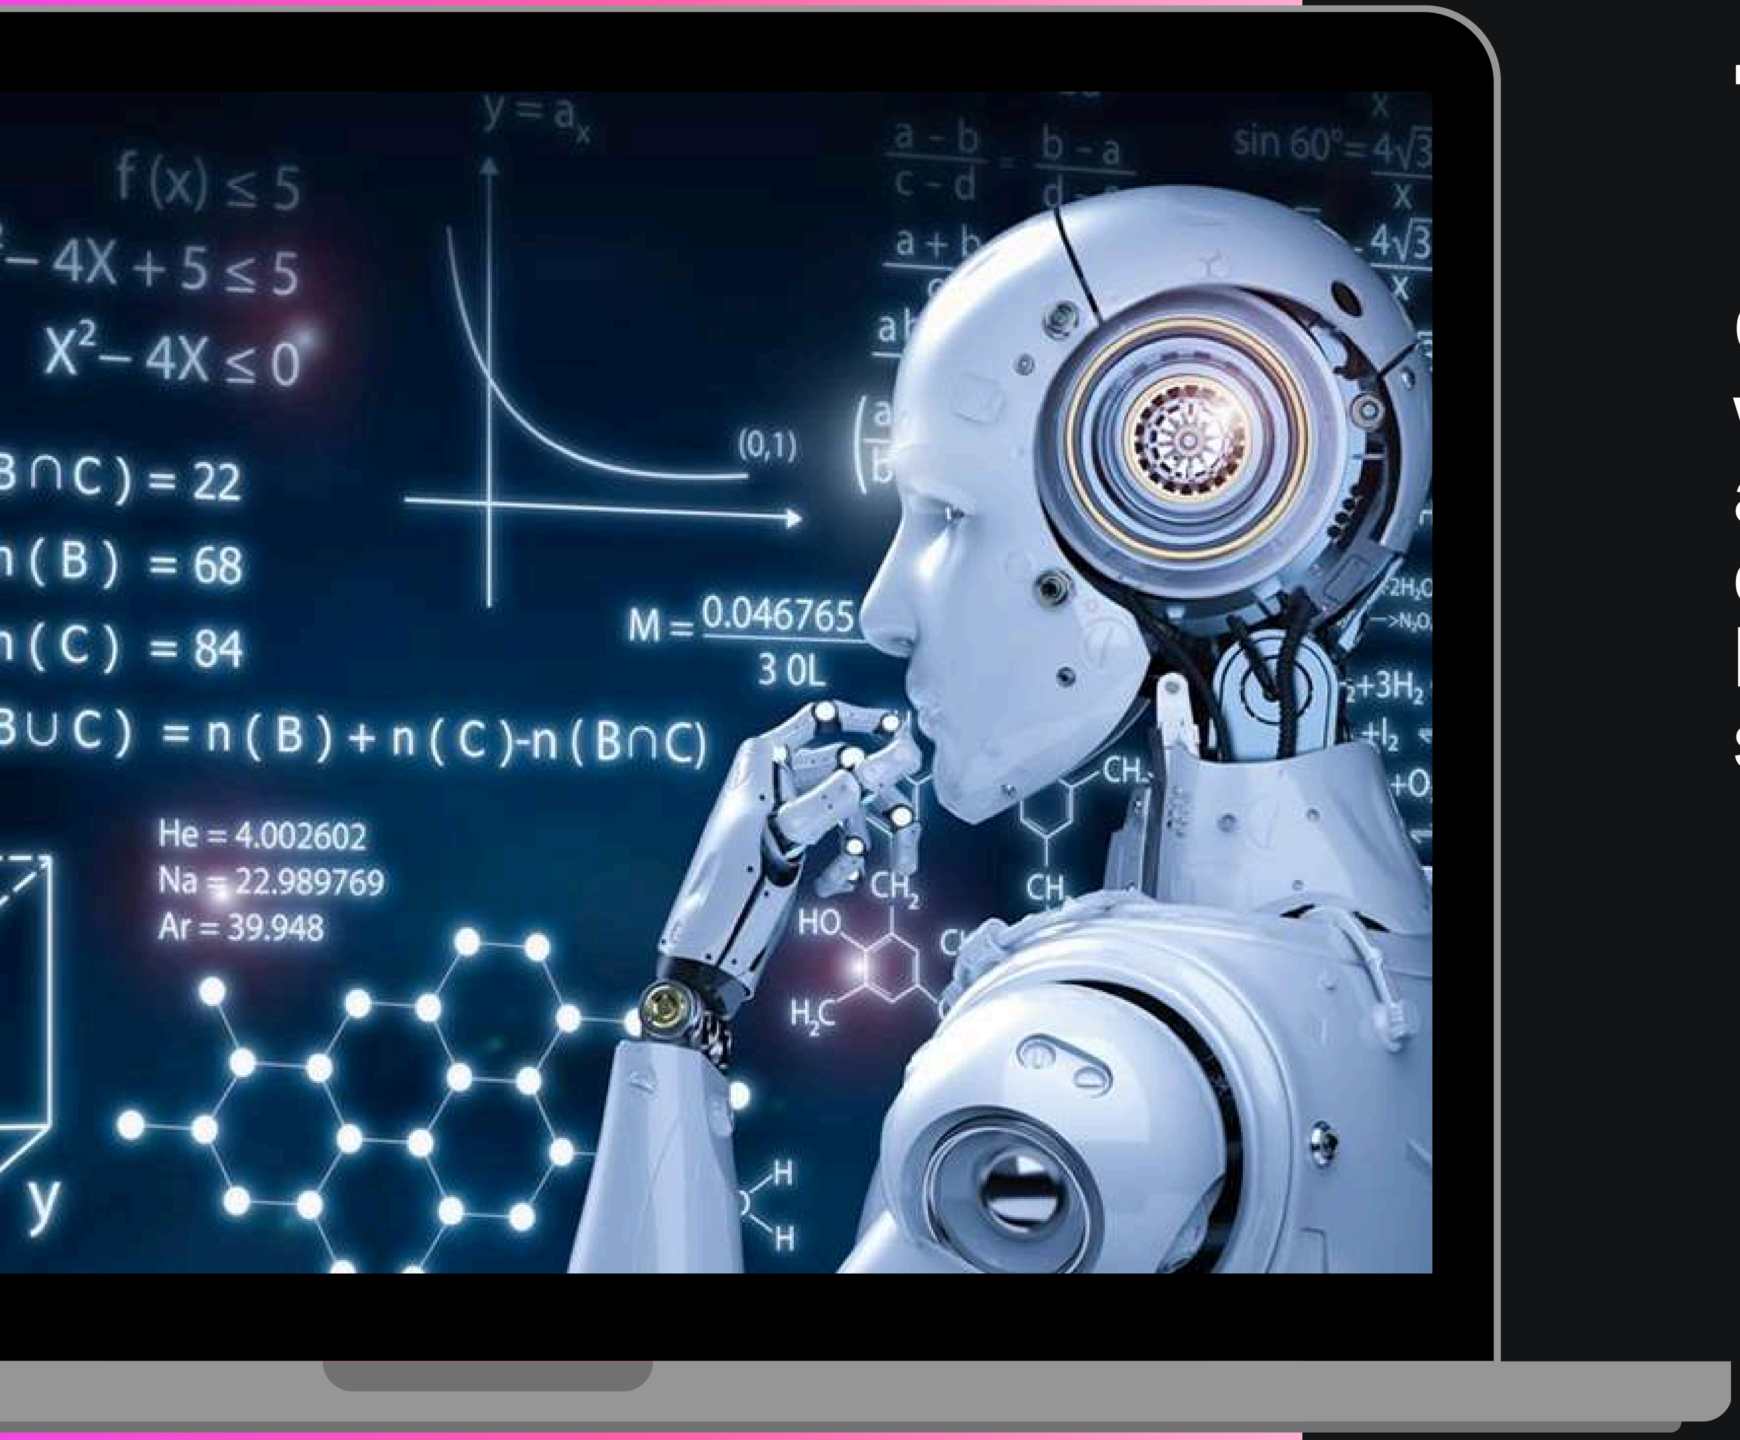

Literature Review: Traditional Methods

Combining real-time internet sentiment with weather data using advanced ML algorithms (like LSTMs and transformers) can enhance forecasts. This adds a behavioral layer to account for mood-driven spikes in energy usage.

- **What goes well**

  work well for energy prediction when combined with relevant features

- **What doesnt**

  overlook real-time public sentiment and extreme weather trends.

https://ieeexplore.ieee.org/document/10372216

https://www.nature.com/articles/s41598-024-56602-4

The data

Amount Of Features & Extraction

Weather
9 Features
Source: Kaggle

GDELT
30 Features
Source: GDELT GKG API

Energy & Time
11 Features
Source: Kaggle (GDELT alignment done)

PREPROCESSING PIPELINE

Align Time-Stamps → Resample → Merge on timestamp, forward-fill 15 min

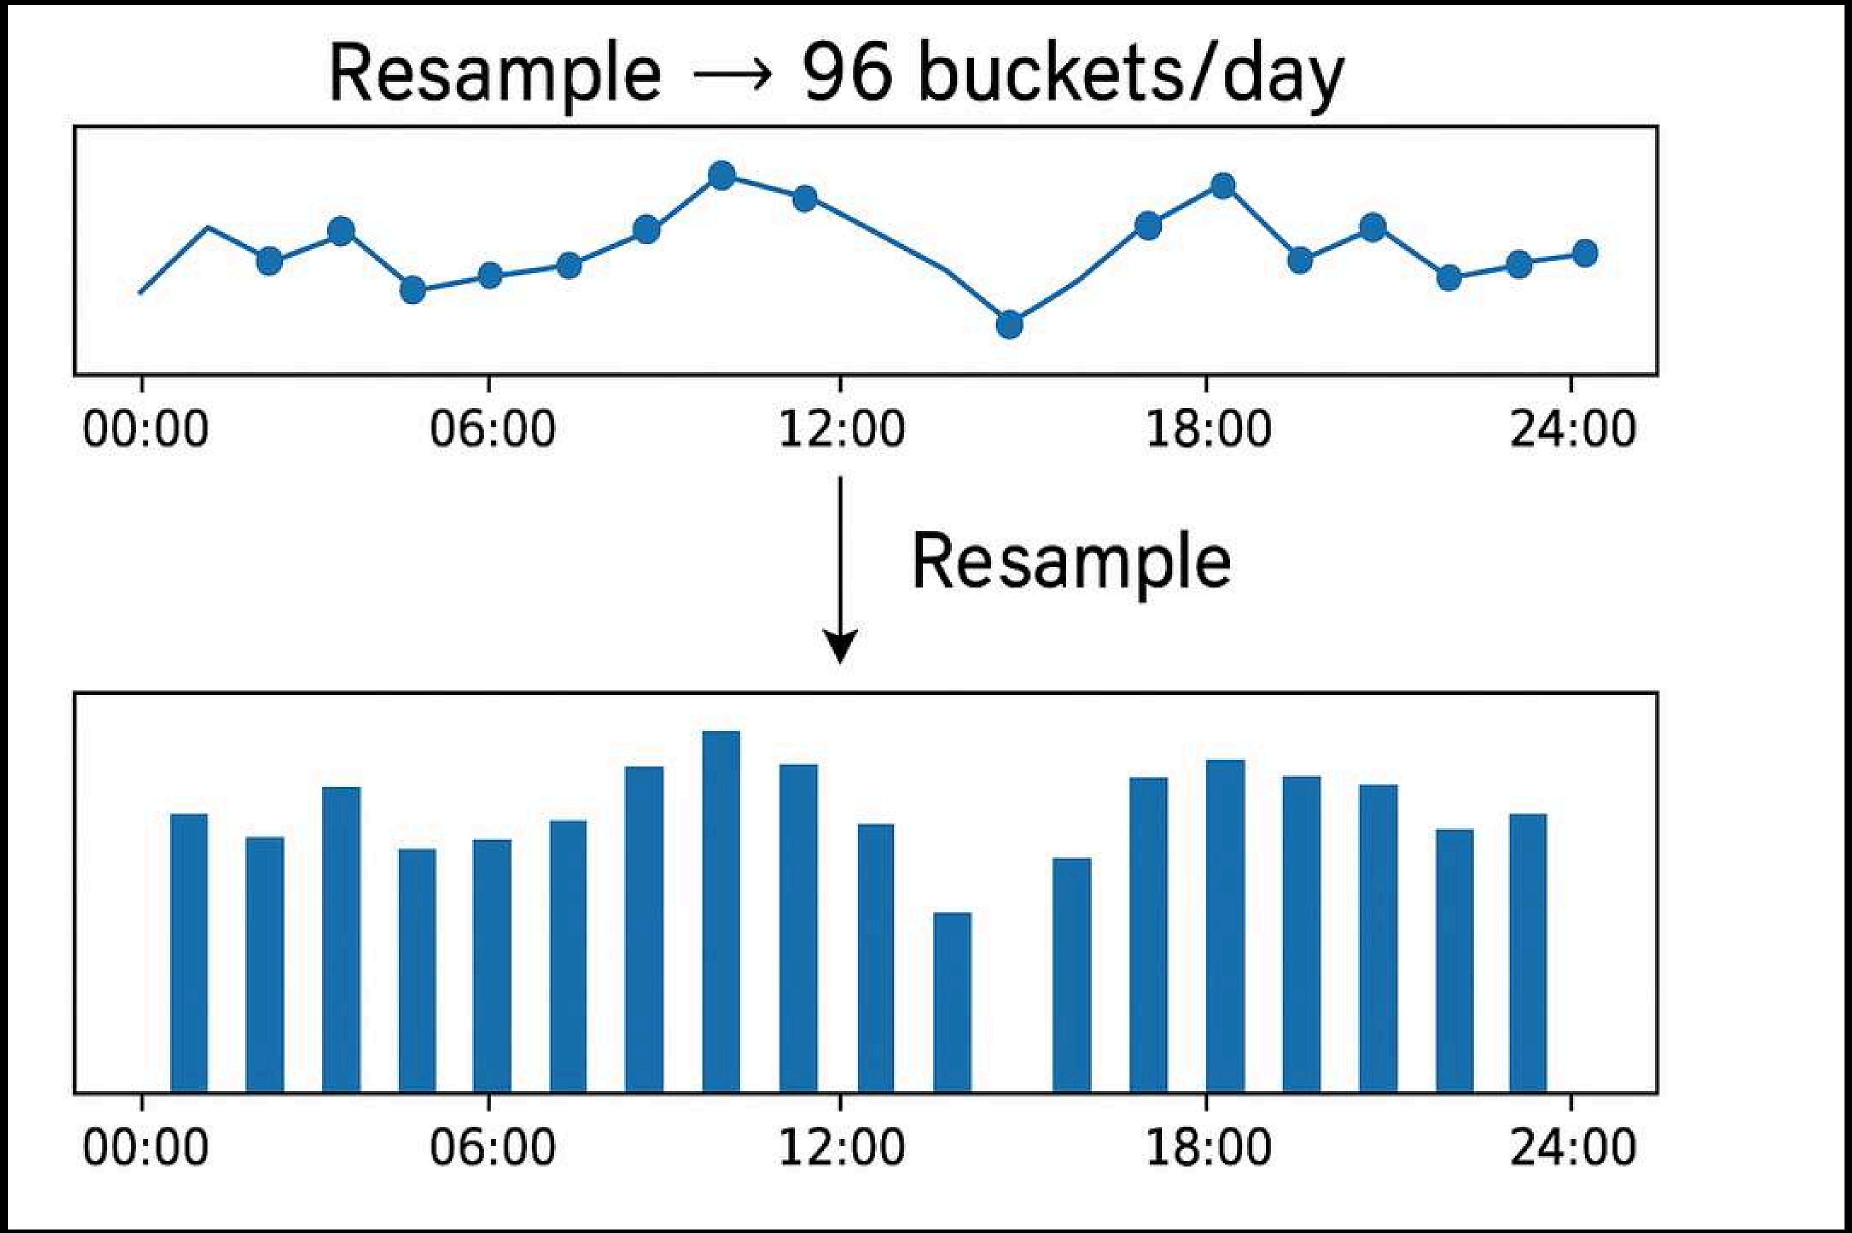

Feature Engineering

Creating interaction terms, ratios, volatility, momentum, novelty, persistence, seasonality, and decay features.

Preprocessing

Aggregation

Summing, averaging, max/min, and standard deviation of theme and tone features over 15-minute intervals.

Winsorization

Capping extreme values to reduce outlier impact.

Correlation Filtering

Removing highly correlated features to reduce multicollinearity.

Temporal Features

Adding time-of-day, day-of-week, and business hour indicators.

Relative Scaling

Calculating relative importance of each theme (e.g., dividing by total theme mentions).

Missing Value Handling

Filling missing intervals and values with 0 or forward/backward fill.

Oops!

Why Our First Scores Were Illusions

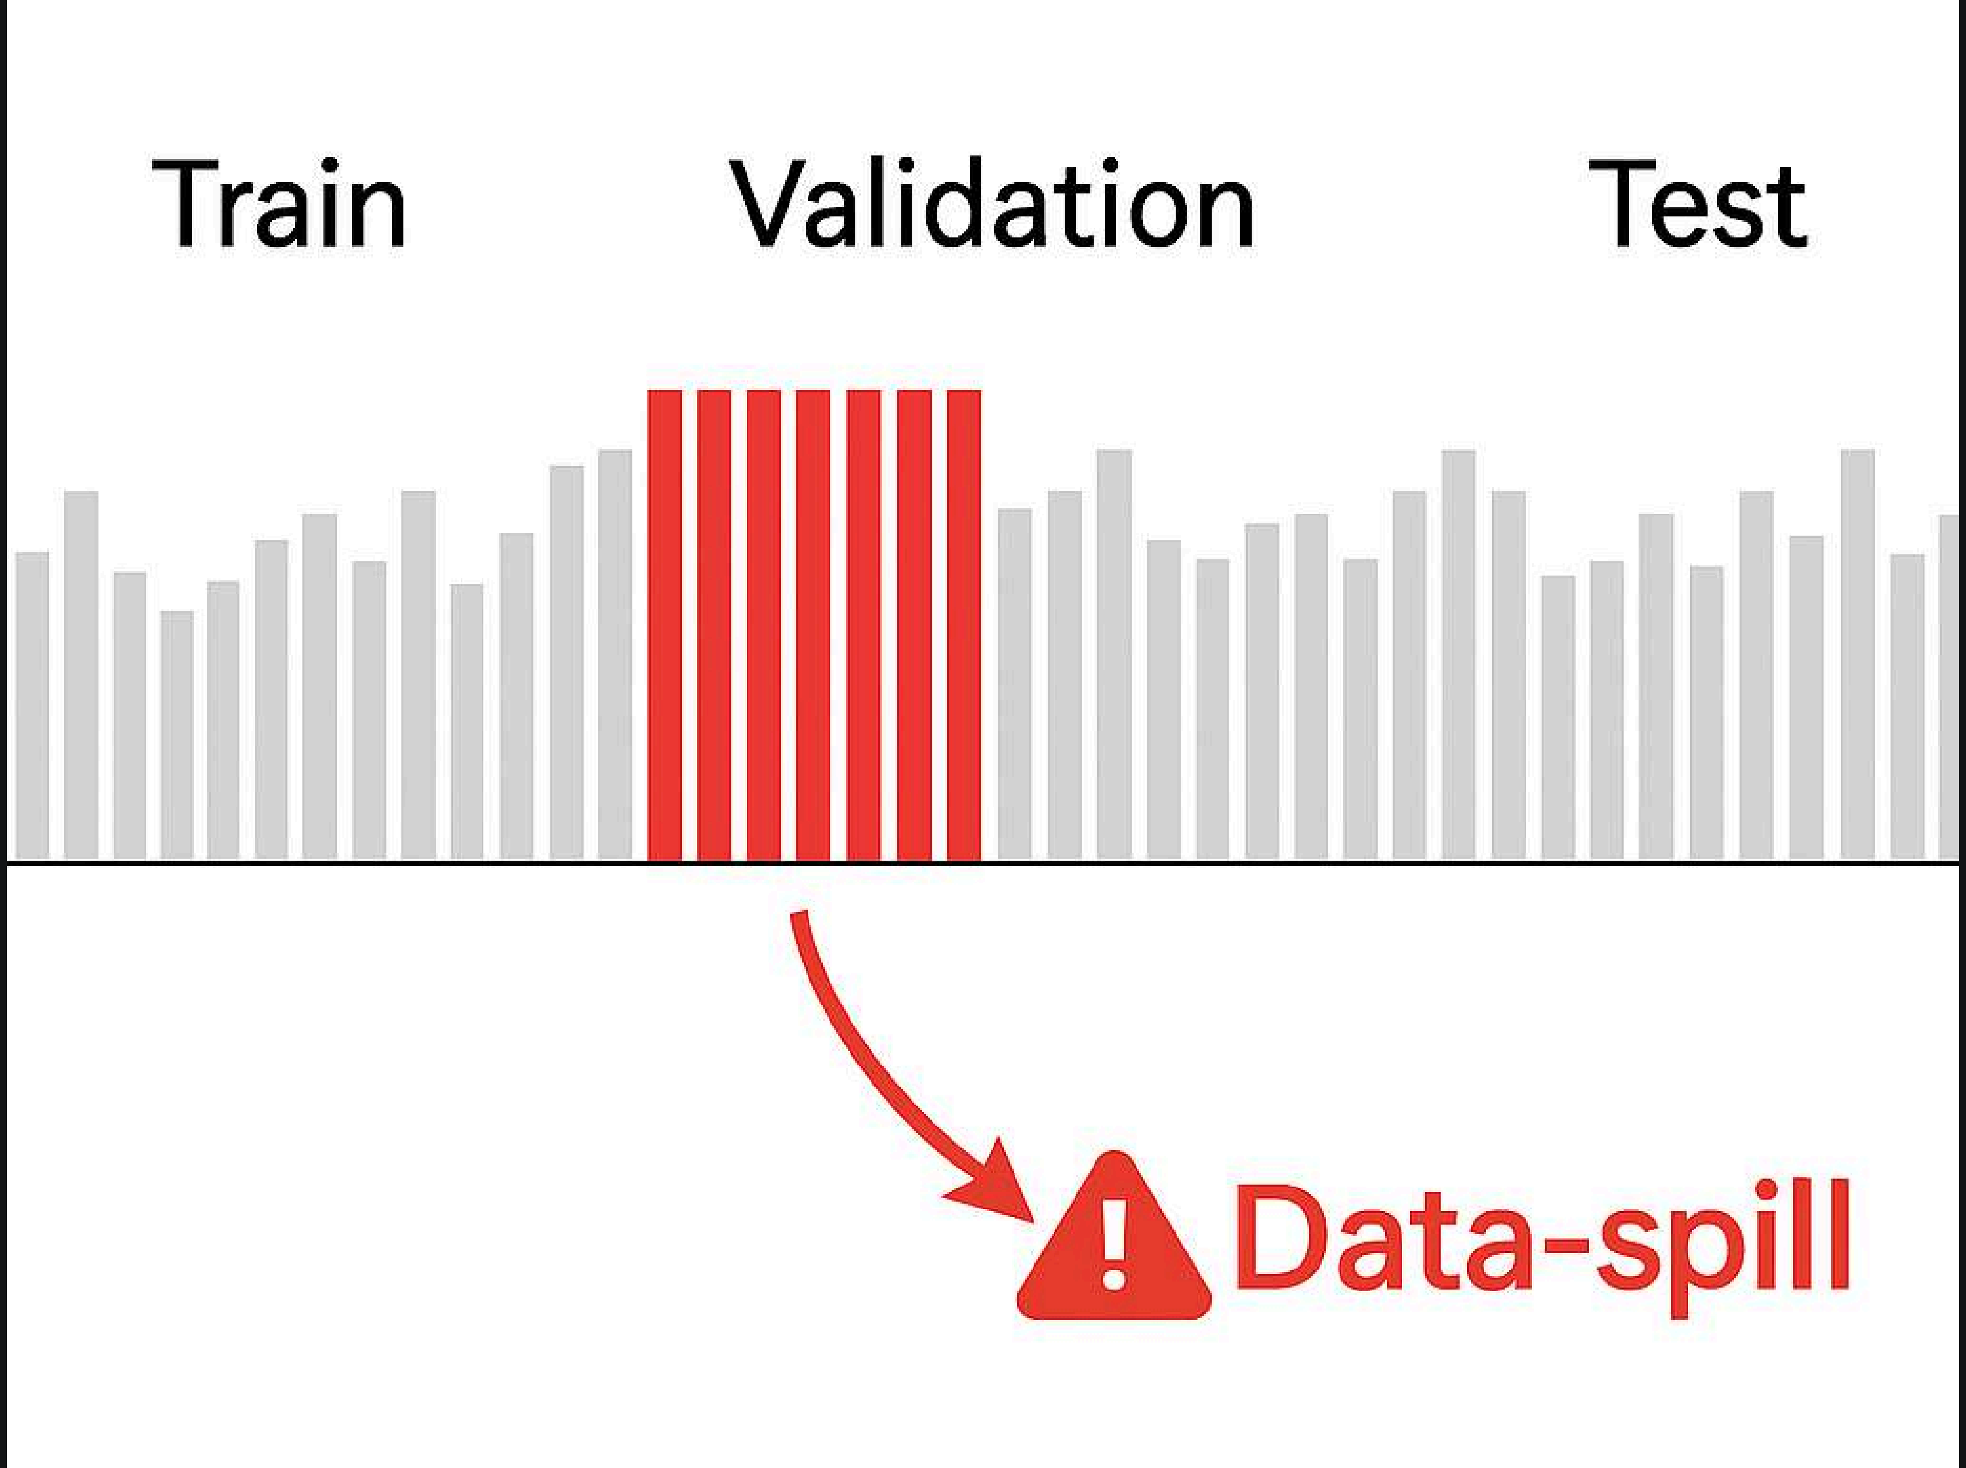

ML Methodology

Regressors we tried:

- LSTM
- XGBoost

Why?

- Effective at Handling Complex Temporal Relationship
- Strong Memory Retention for Long-Term Dependencies
- Robustness in Handling Noise & Anomalies

Parameters?

Optuna!

Performance Metrics

1 **RMSE**
- < 10% of mean target

2 **R^2**
- > 0.90 : Range-: 0 - 1

3 **MAE**
- < 10% of mean target

Running Evaluation

train : 0–80 %
predict : next 2 %
shift → train 0–82 %
predict : next 2 %
Optuna run every 50 trials

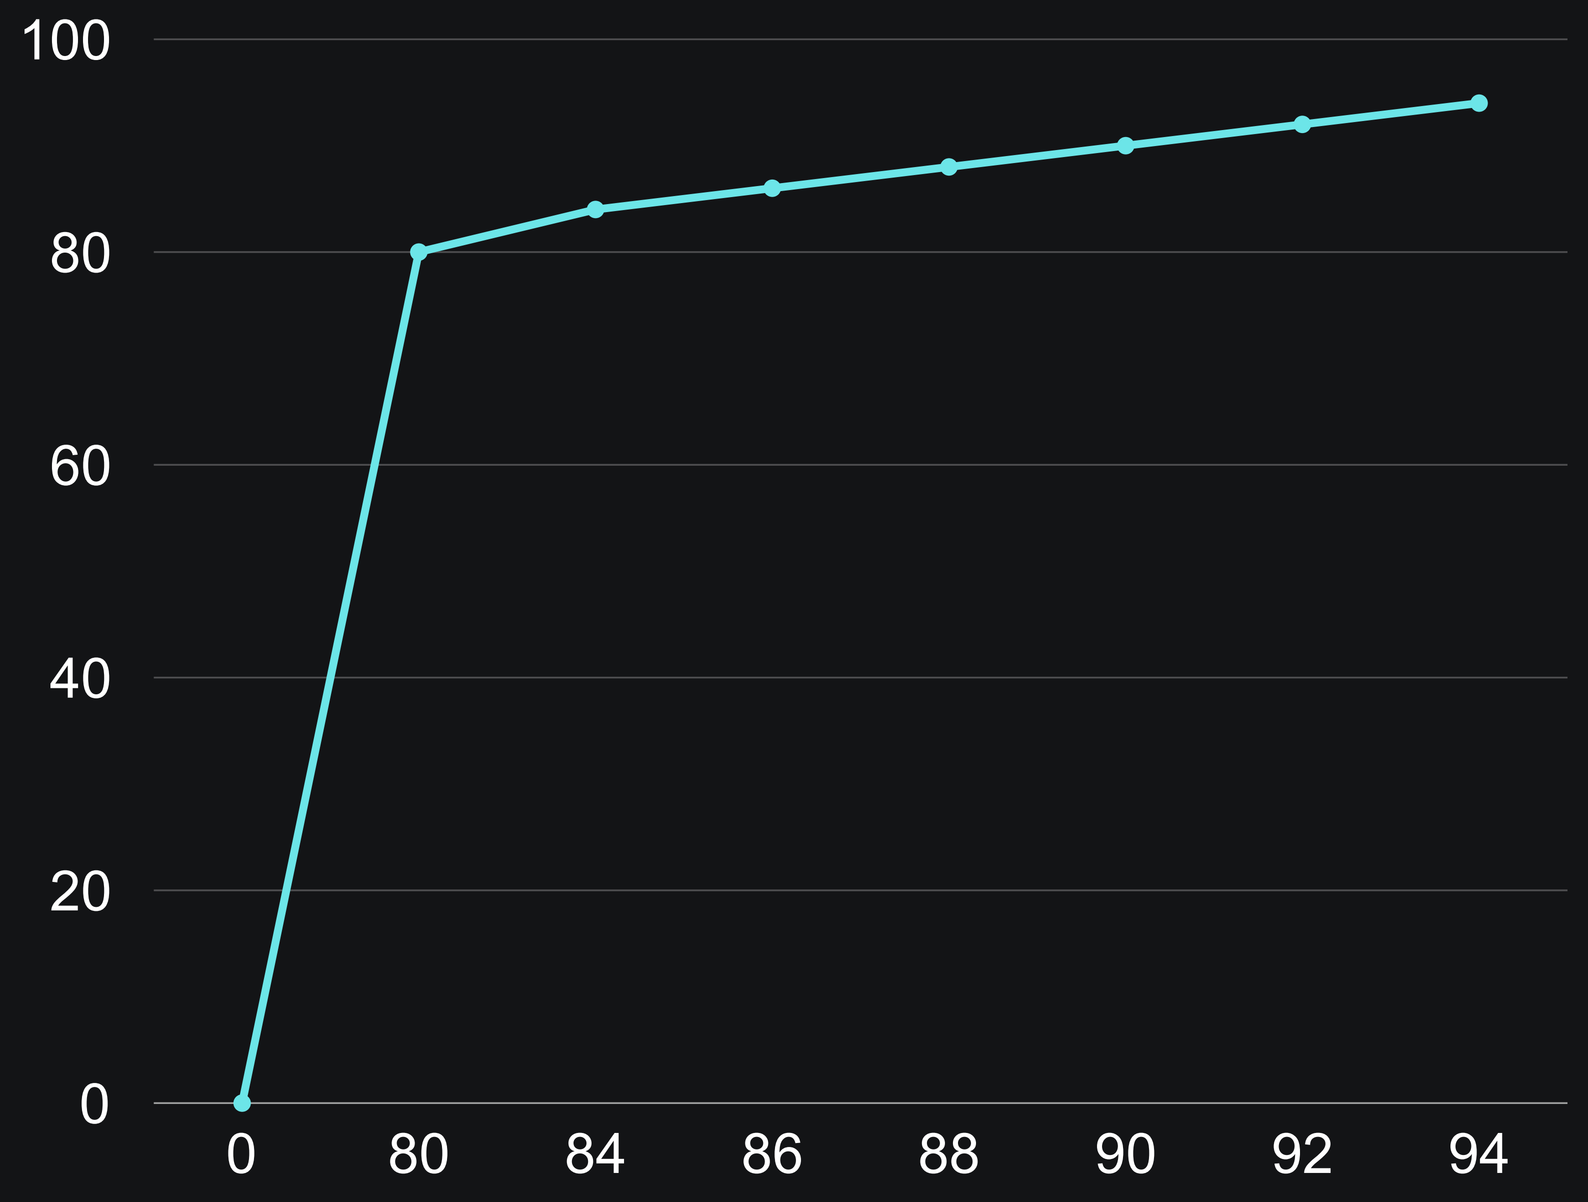

Checking Outlier Prediction

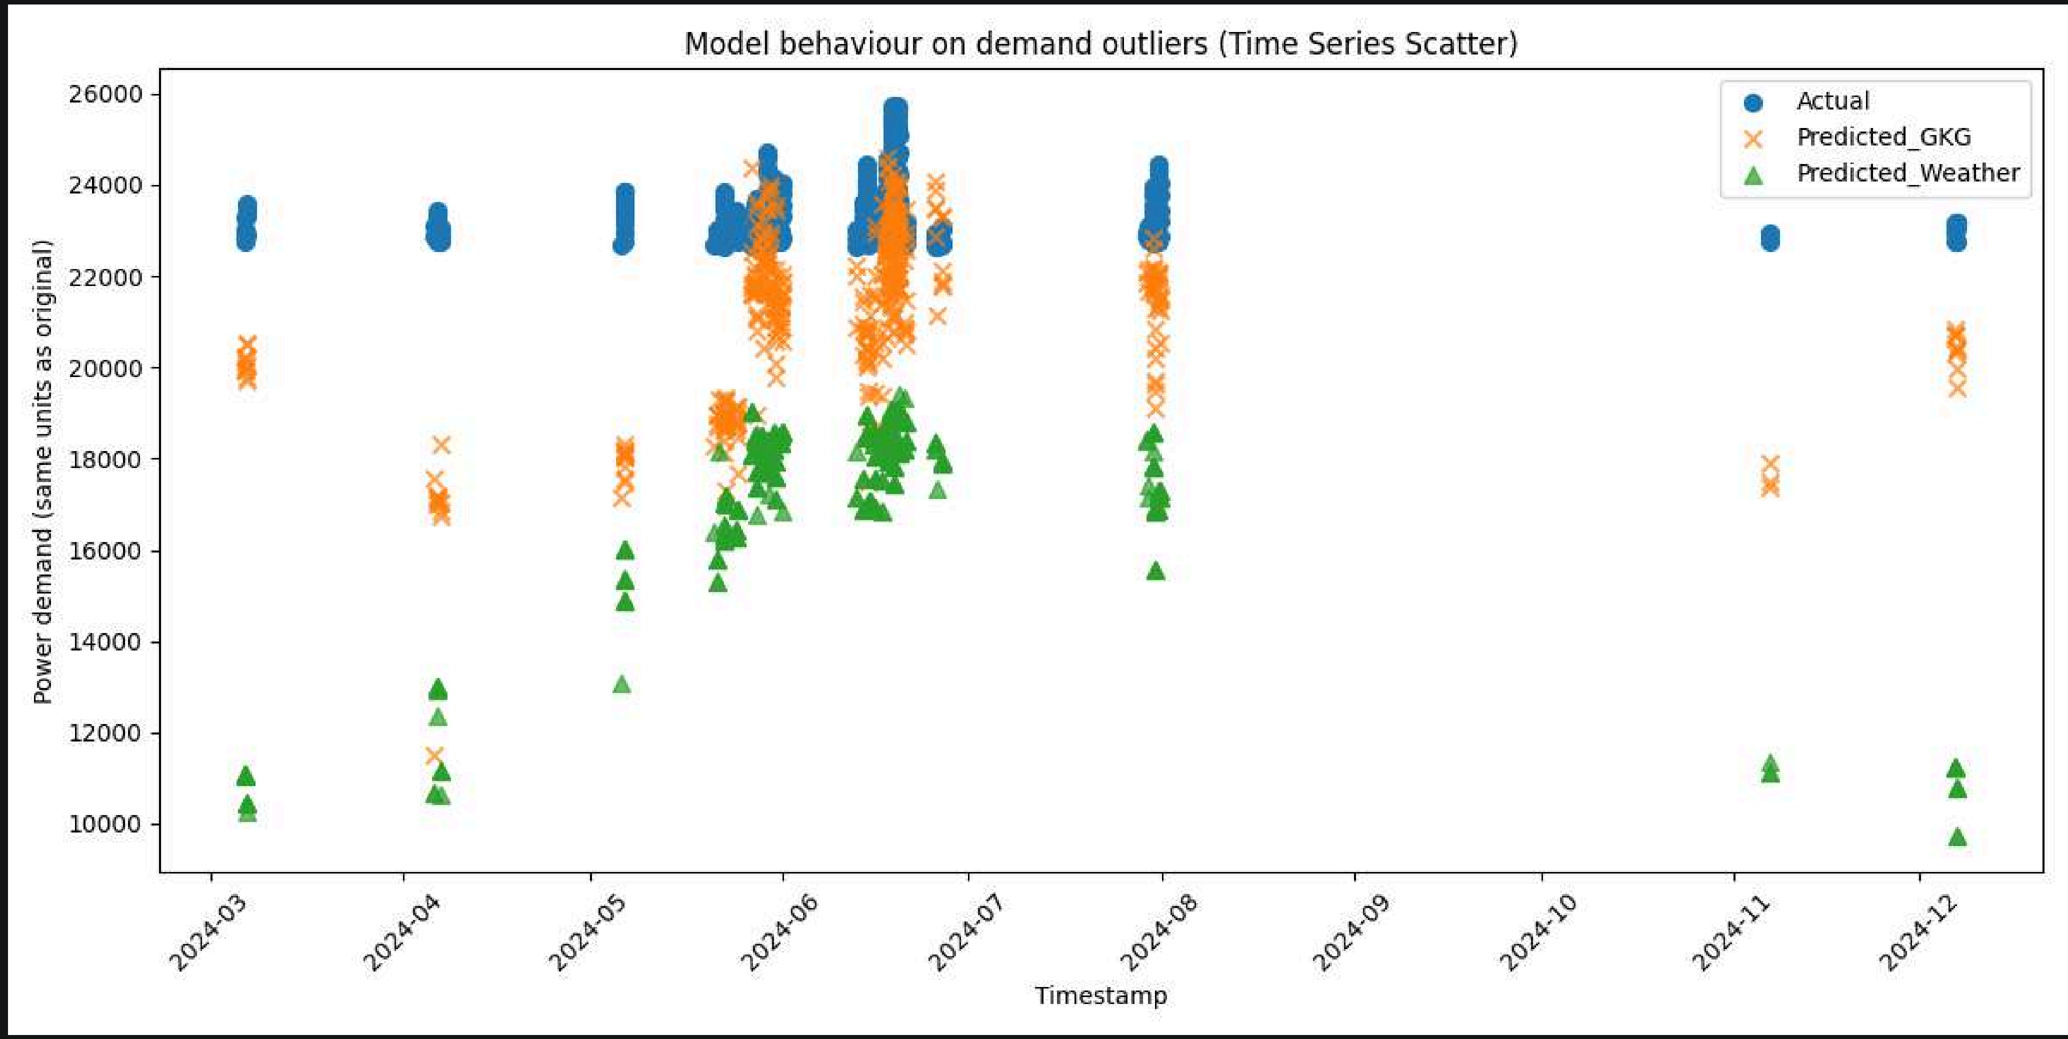

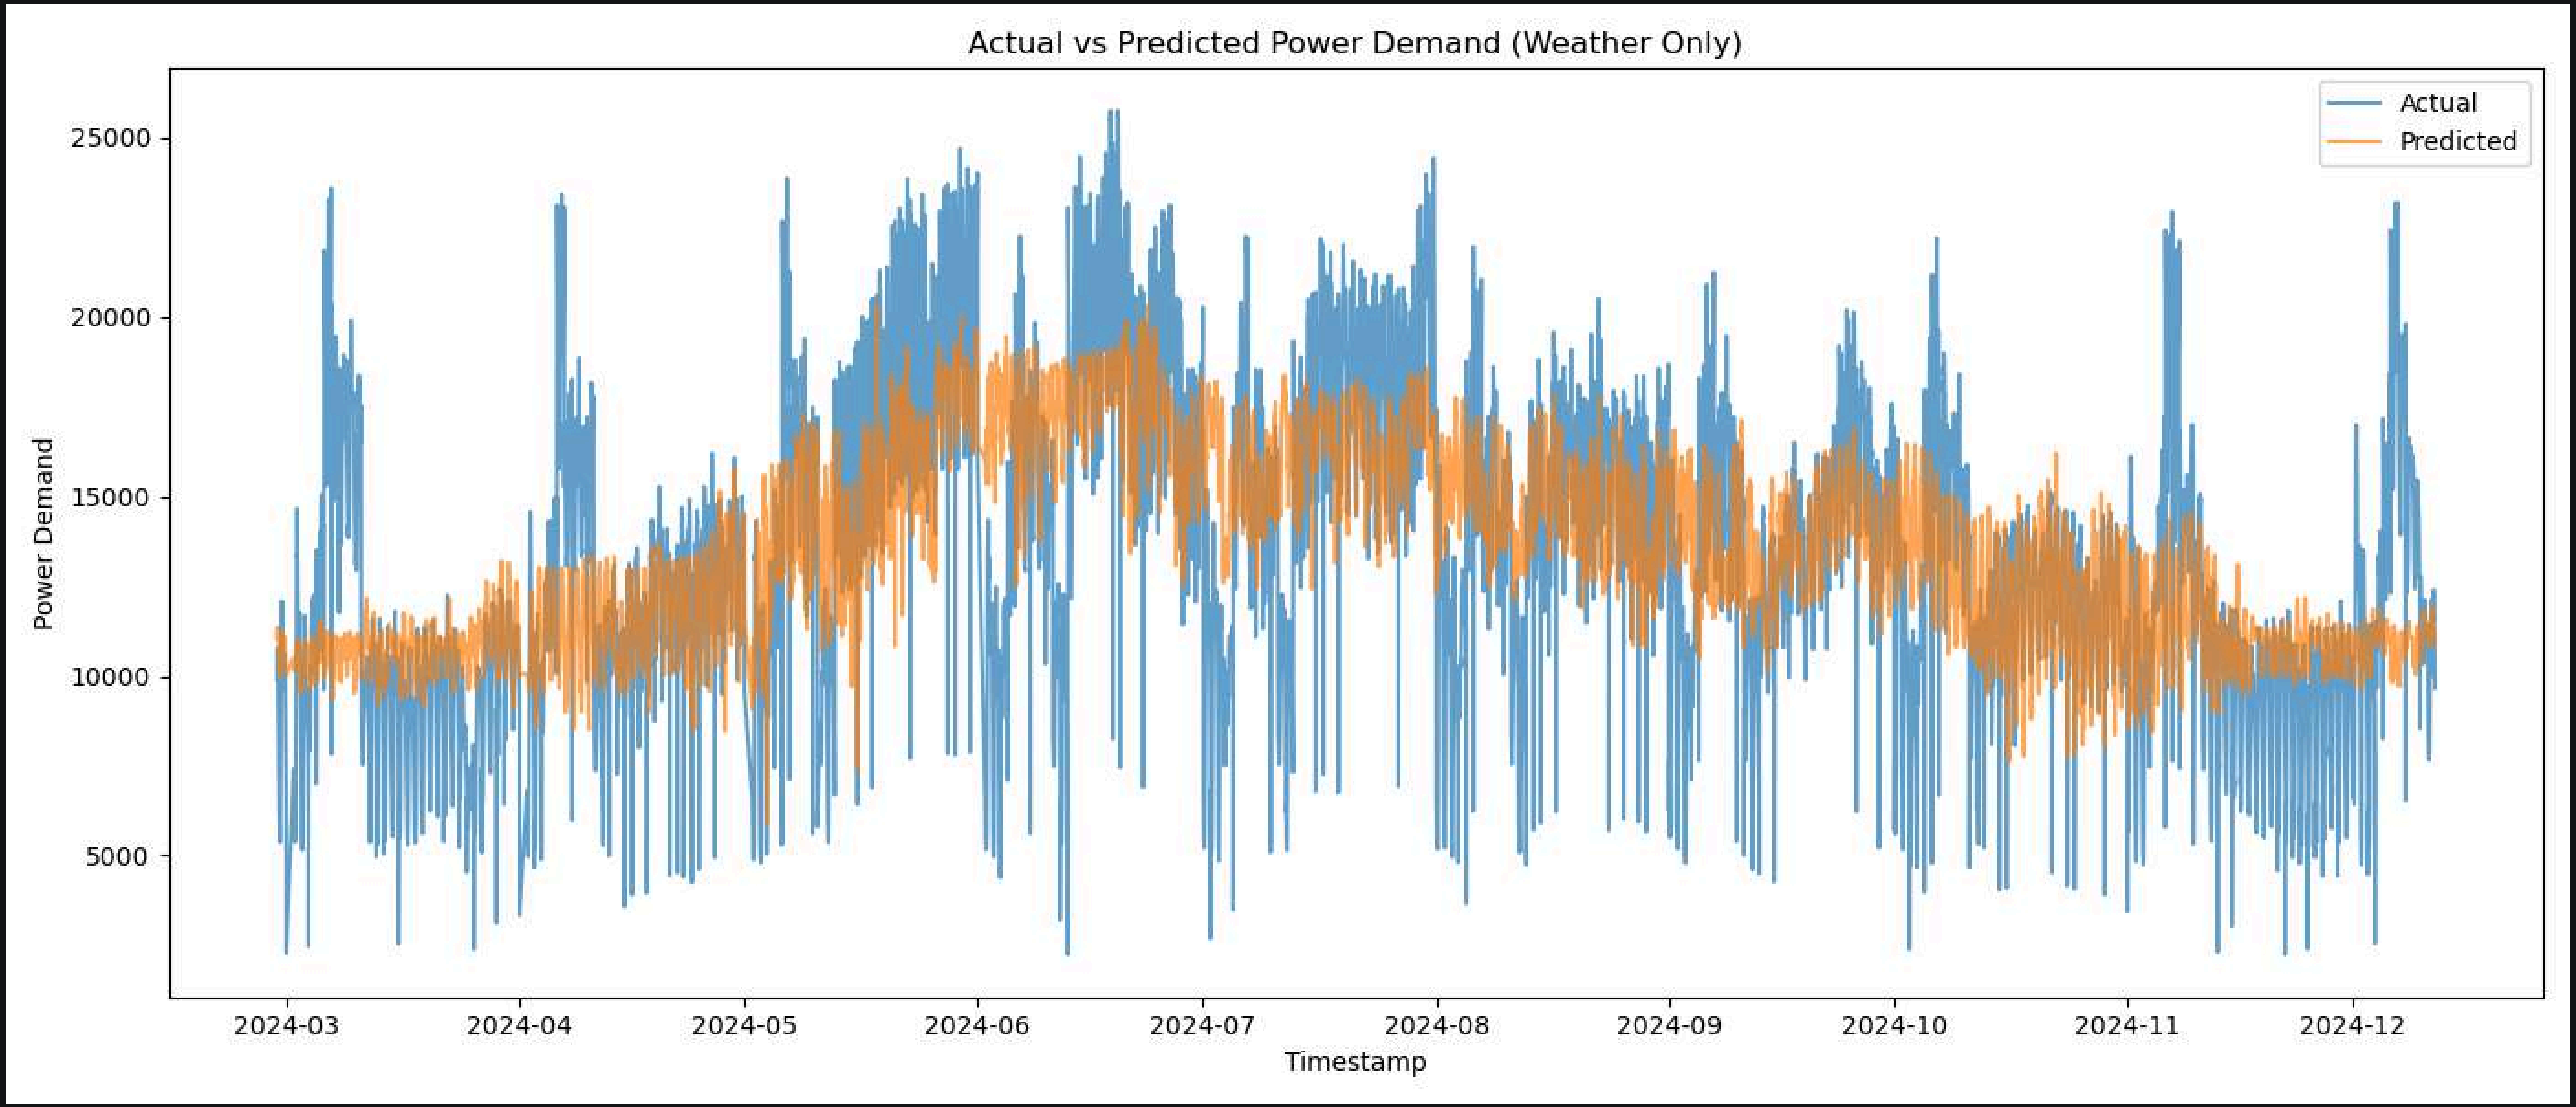

Actual vs Predicted Power Demand (Weather Only)
Actual
Predicted
Power Demand
25000
20000
15000
10000
5000
2024-03
2024-04
2024-05
2024-06
2024-07
2024-08
2024-09
2024-10
2024-11
2024-12
Timestamp

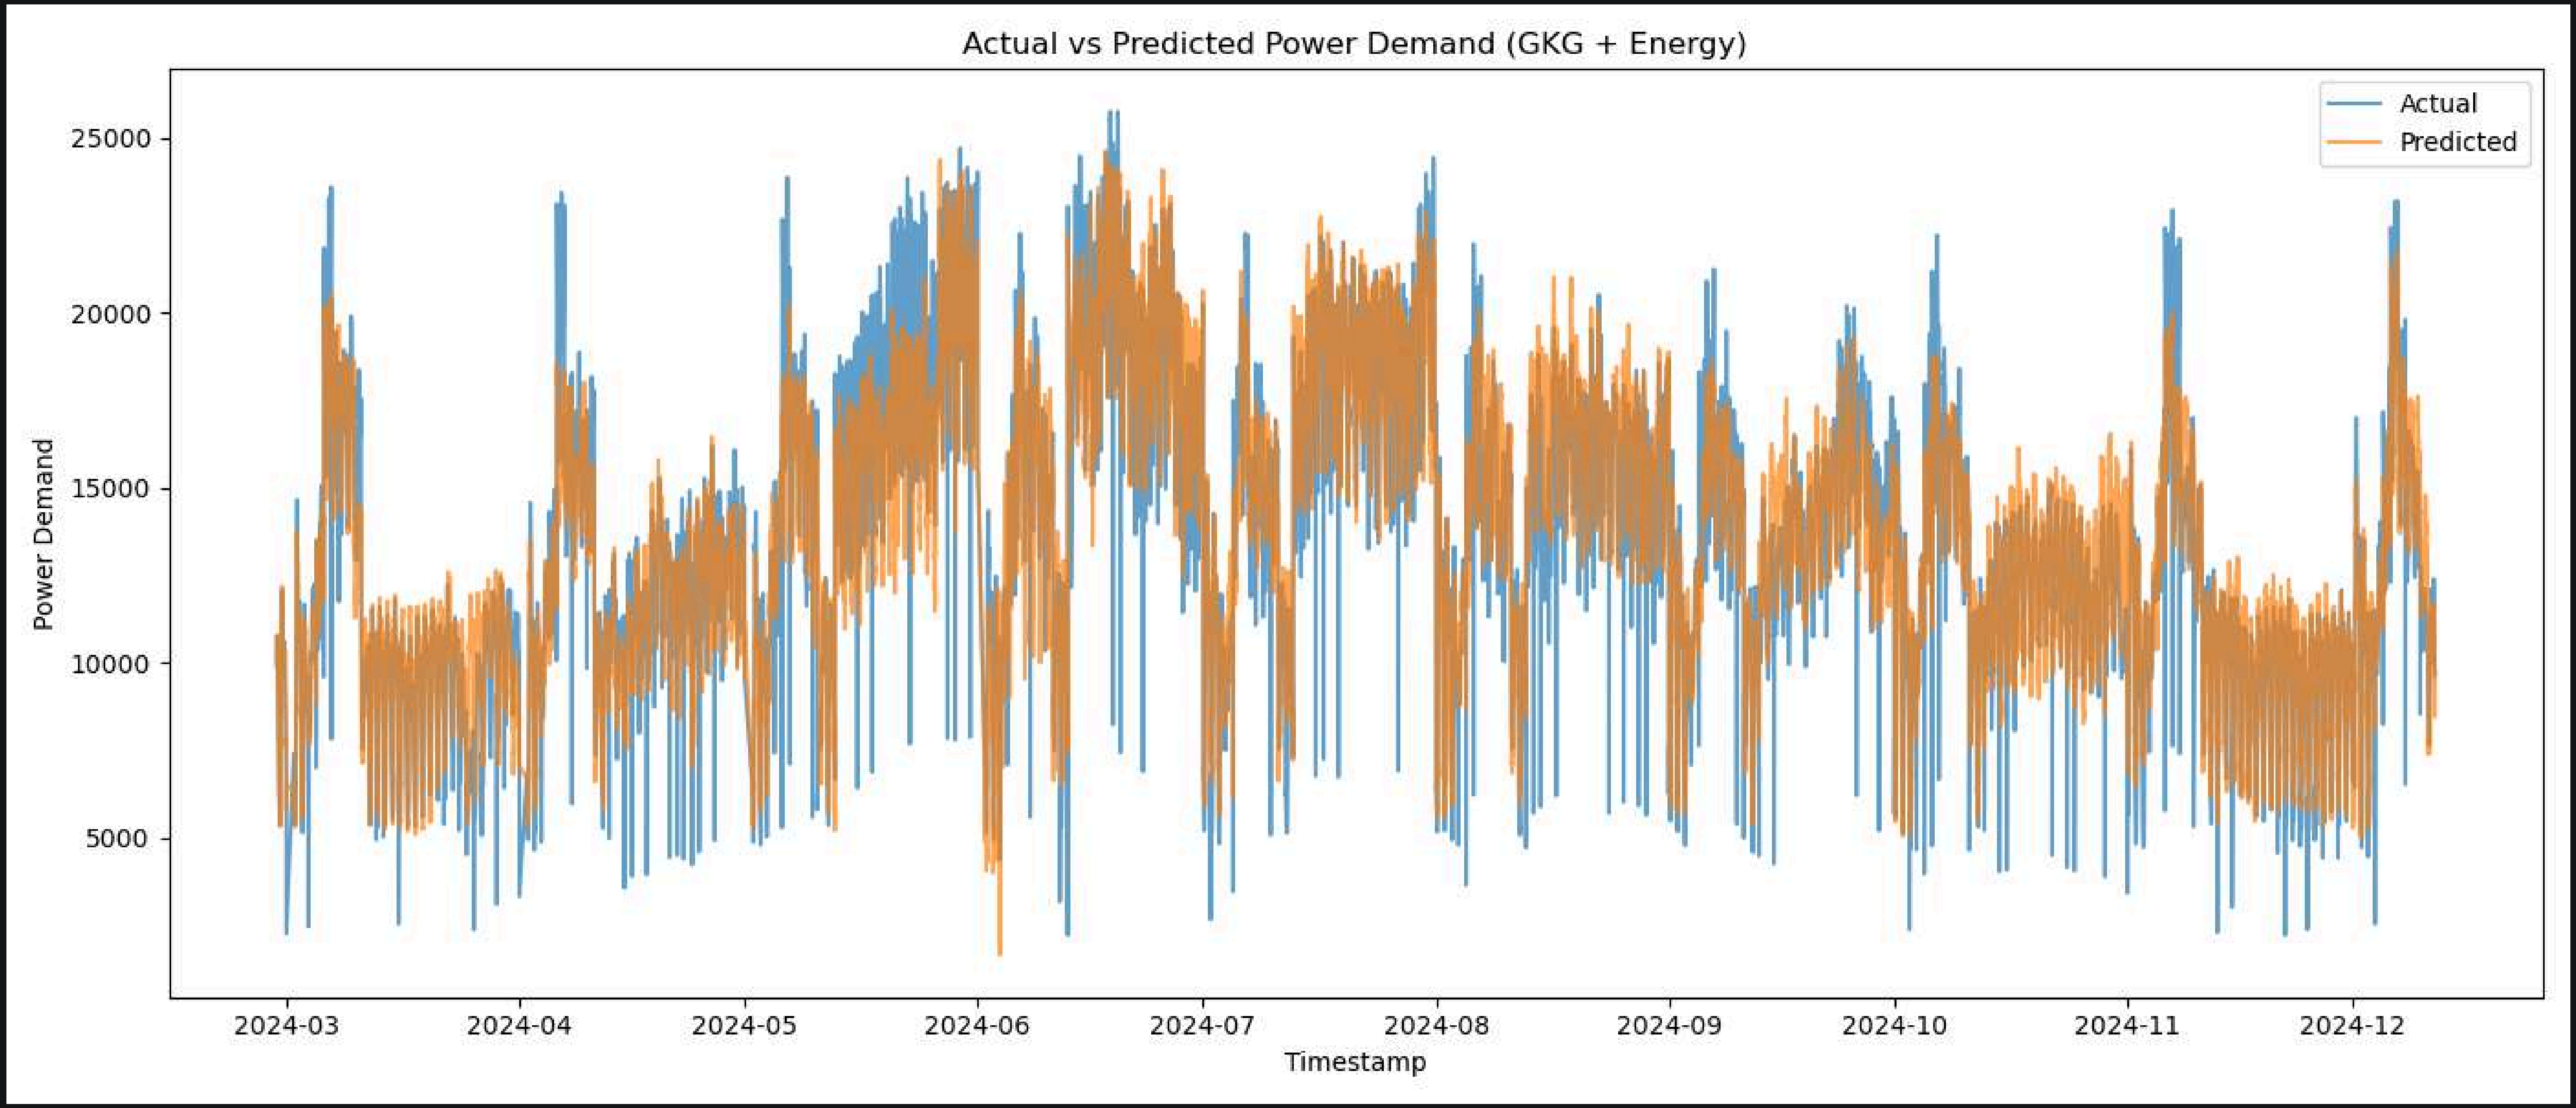
Actual vs Predicted Power Demand (GKG + Energy)
Actual
Predicted
Power Demand
25000
20000
15000
10000
5000
2024-03
2024-04
2024-05
2024-06
2024-07
2024-08
2024-09
2024-10
2024-11
2024-12
Timestamp

	R^2	RMSE	MAE
LSTM	89.4	1221.65	808.27
XGBoost	91.3	1025.377	742.04

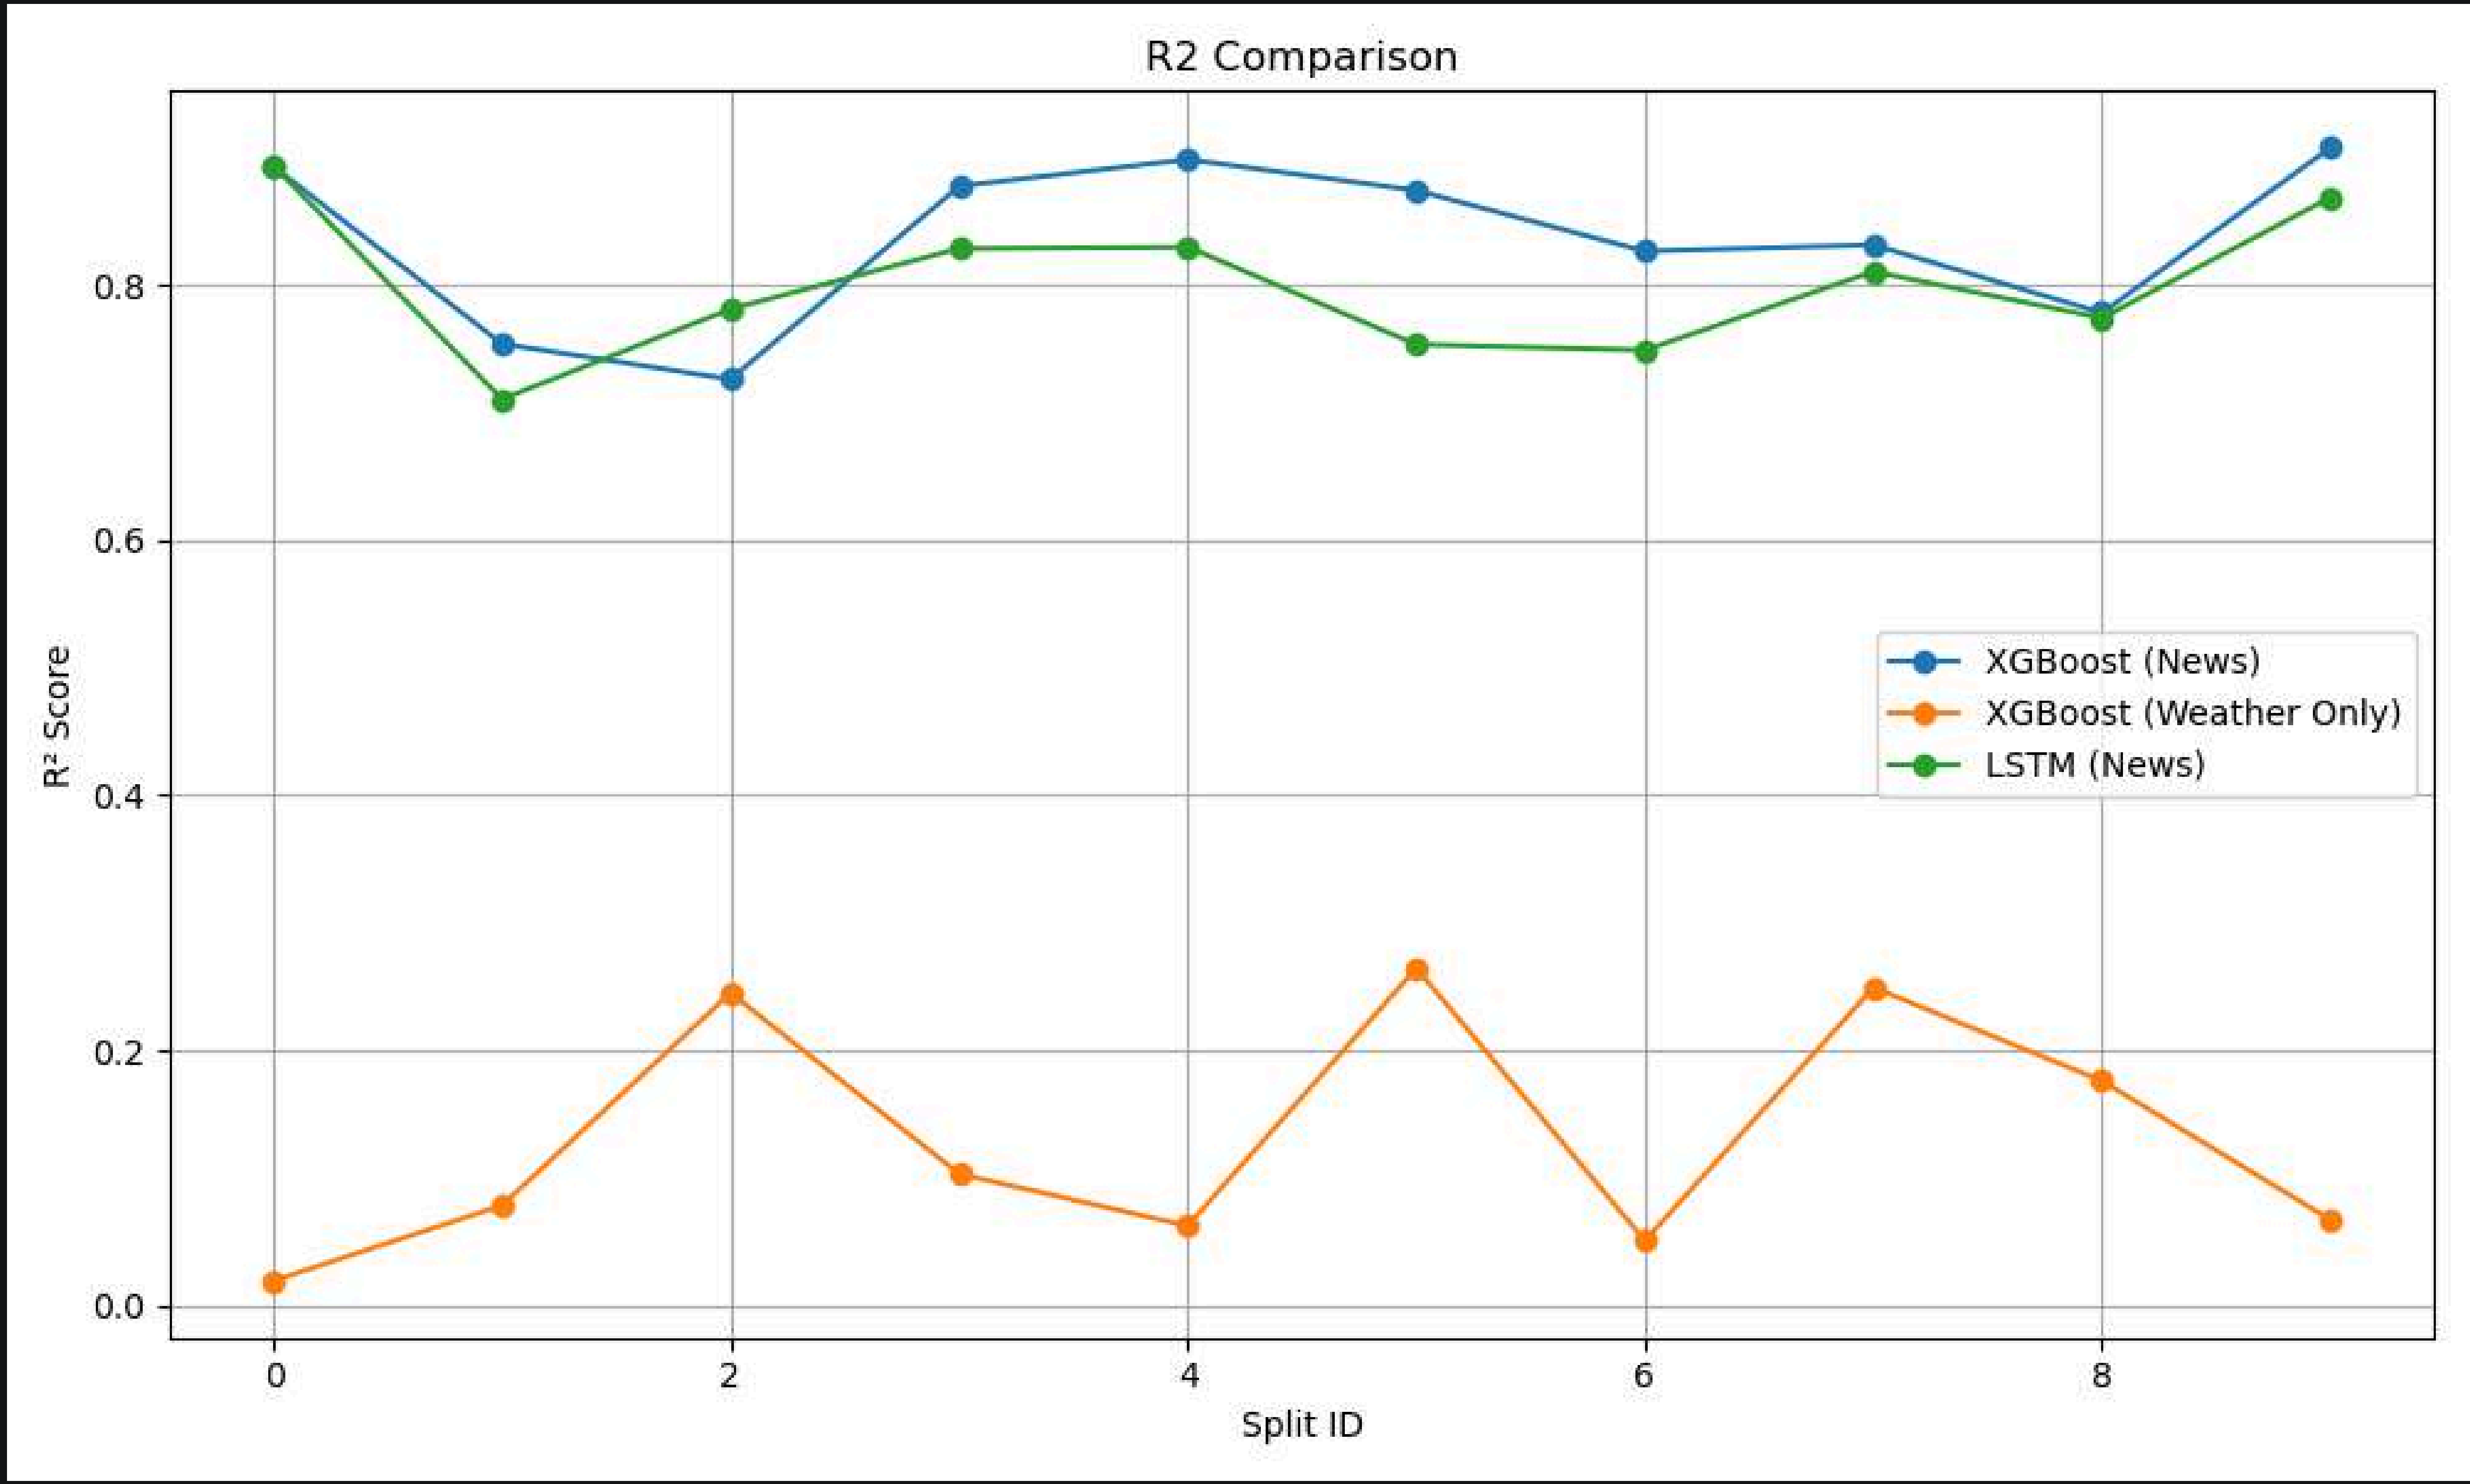
R2 Comparison
R² Score
0.0
0.2
0.4
0.6
0.8
0
2
4
6
8
Split ID
XGBoost (News)
XGBoost (Weather Only)
LSTM (News)

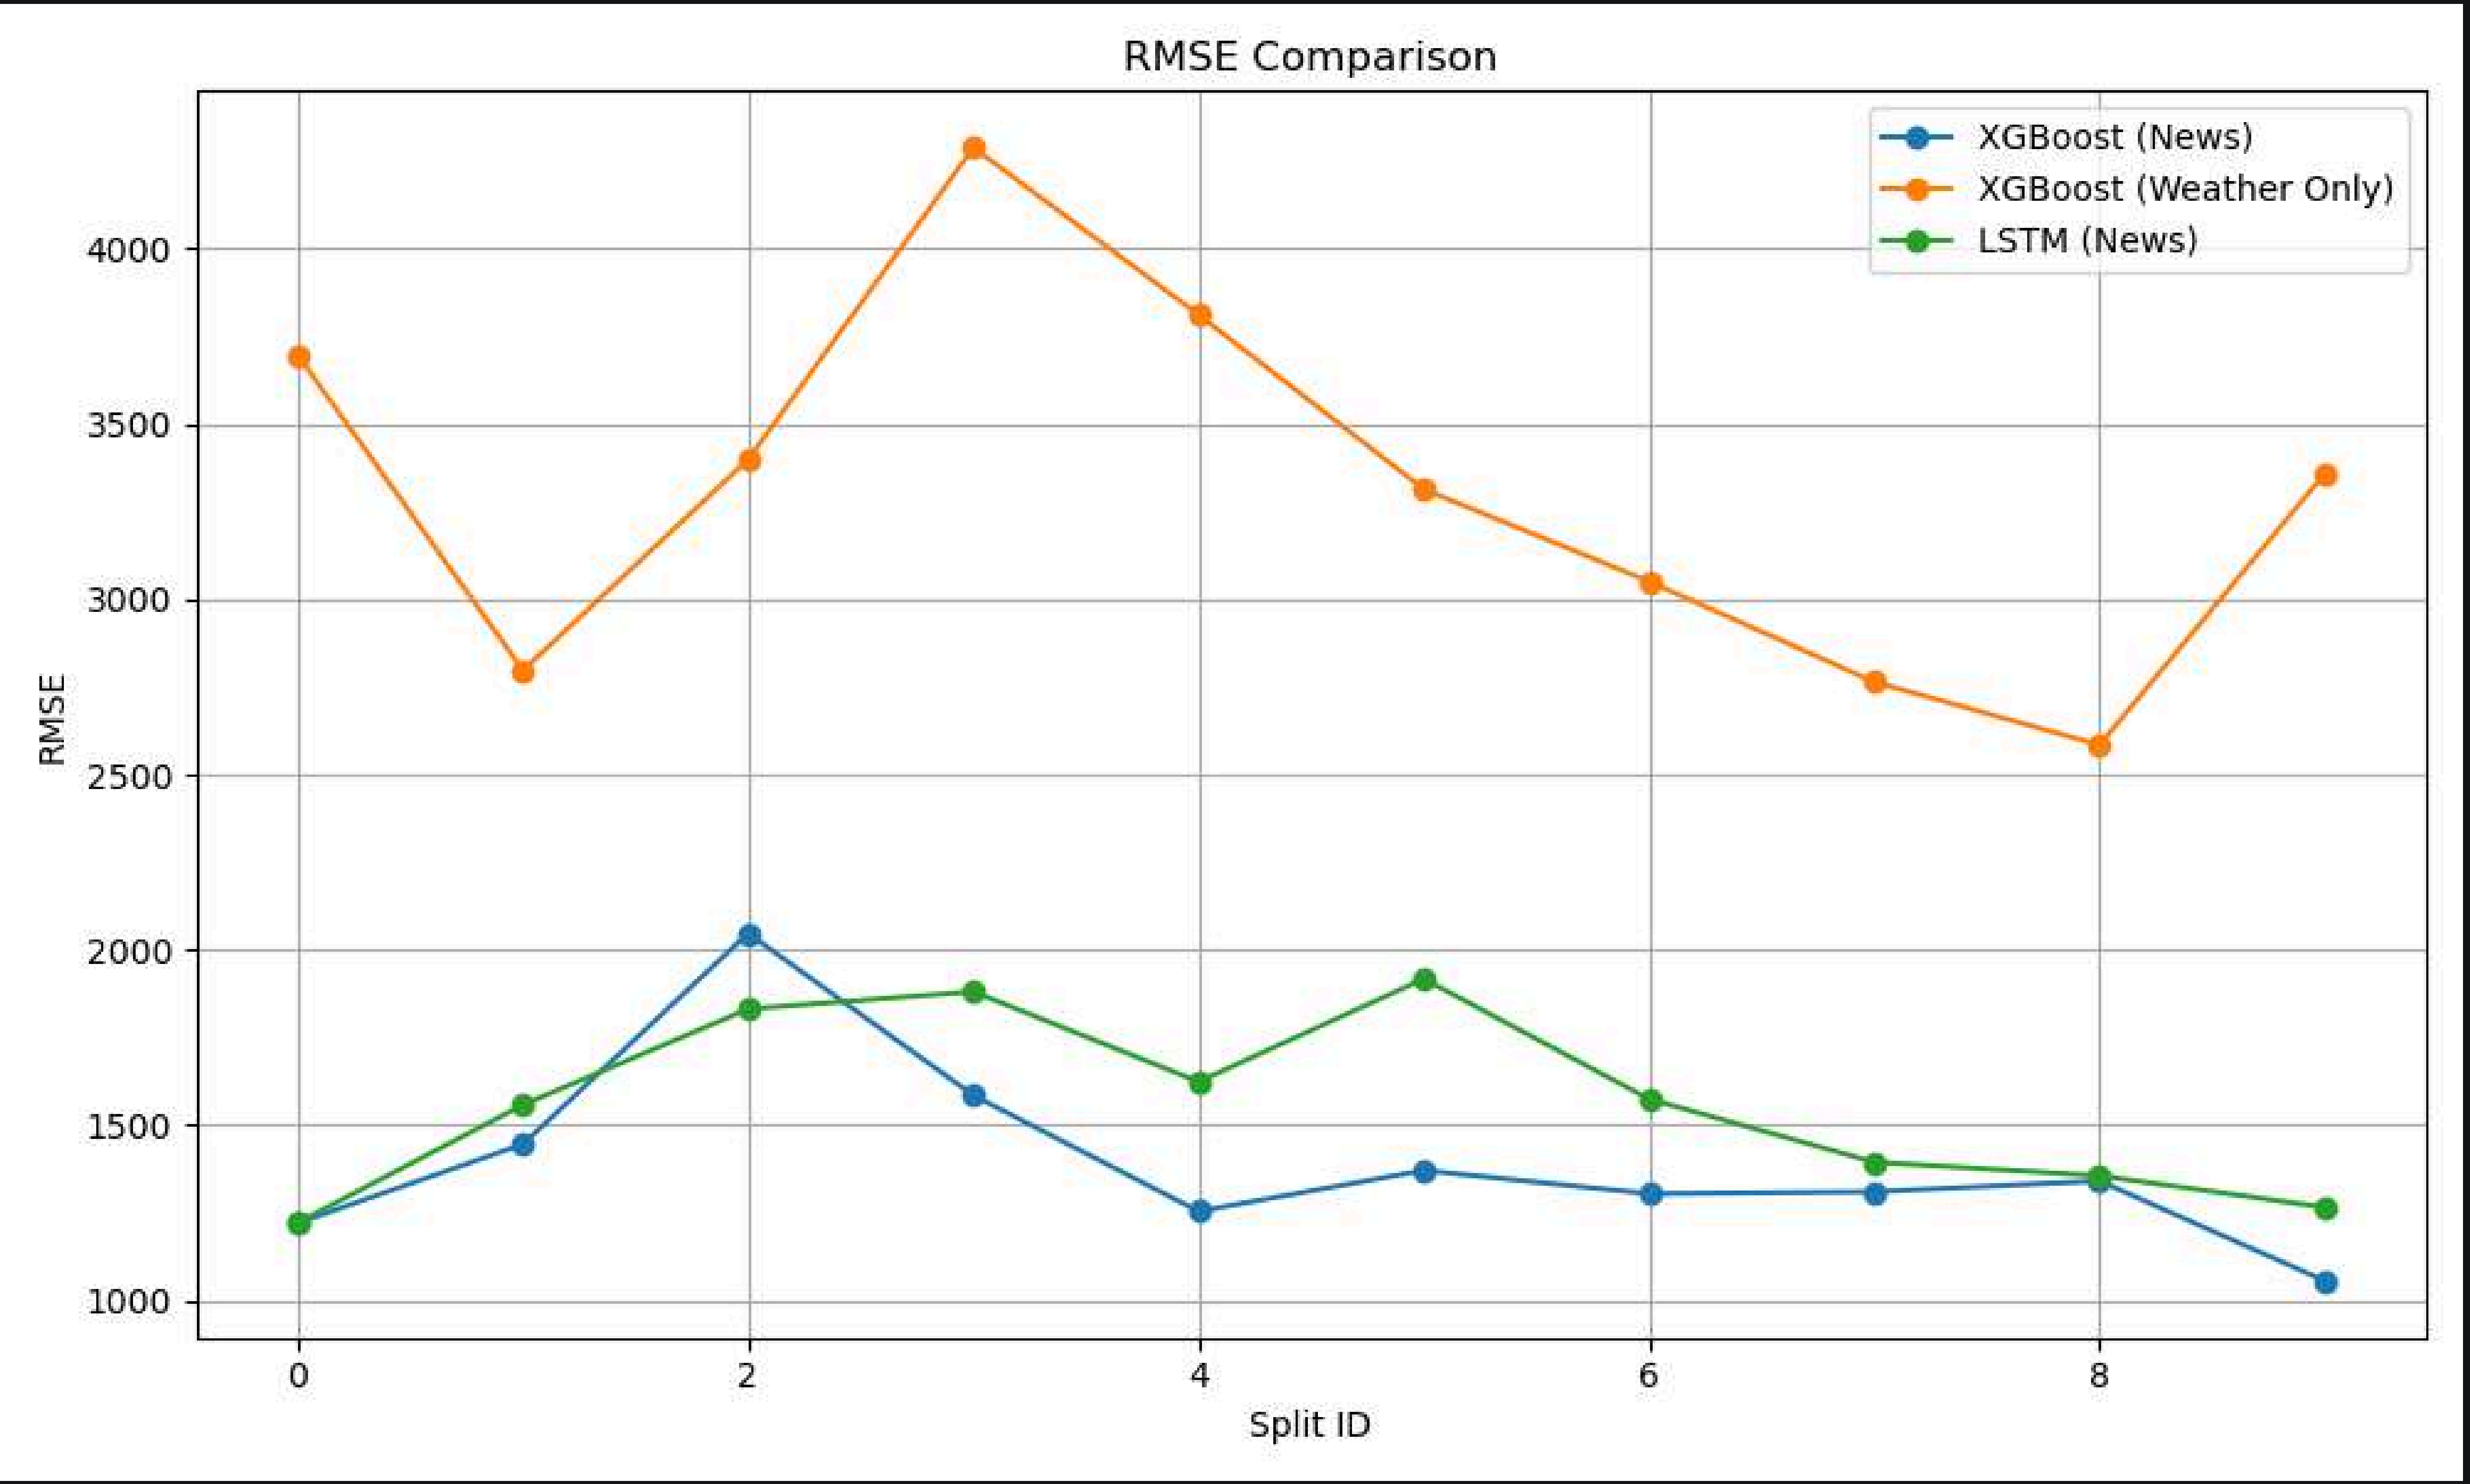
RMSE Comparison
XGBoost (News)
XGBoost (Weather Only)
LSTM (News)
RMSE
1000
1500
2000
2500
3000
3500
4000
0
2
4
6
8
Split ID

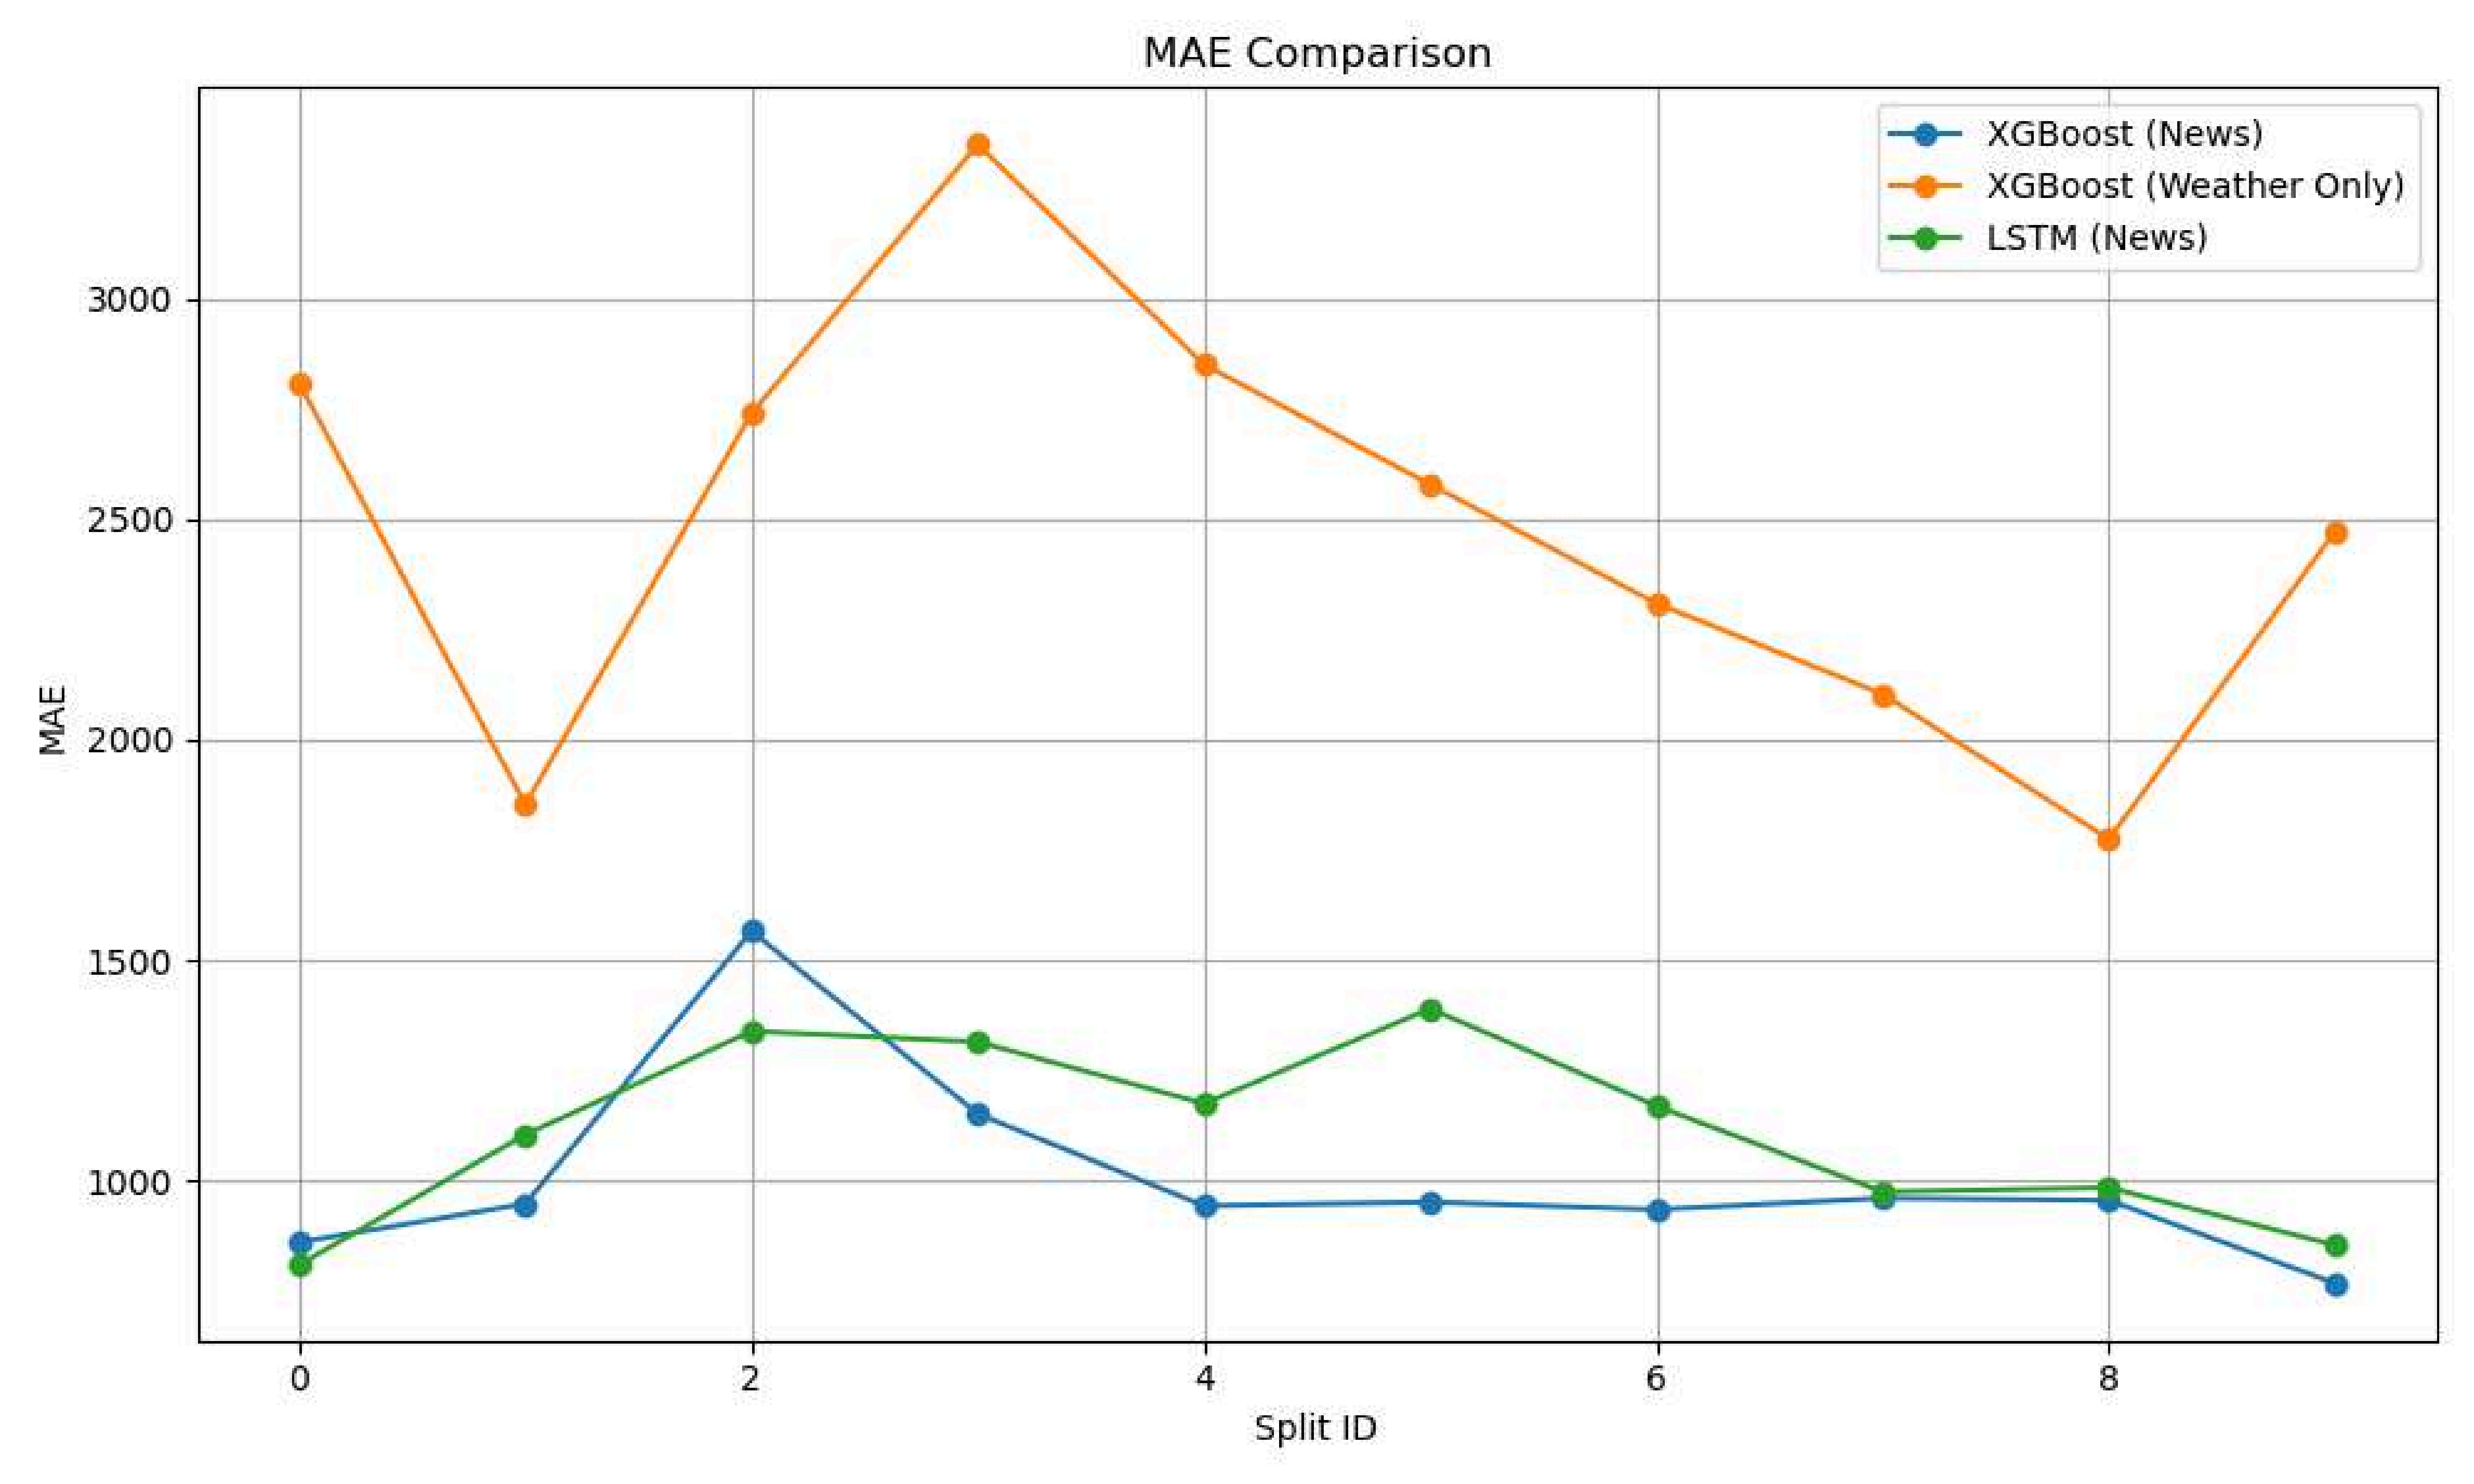
MAE Comparison
XGBoost (News)
XGBoost (Weather Only)
LSTM (News)
MAE
3000
2500
2000
1500
1000
0
2
4
6
8
Split ID

Conclusion

XGBoost Won!

1. Short-term dependencies dominate (load highly autoregressive)
2. Gradient boosting excels with mixed dense (weather) + sparse (news-topics) features
3. Faster Optuna loops → better hyper-tuning frequency
4. LSTM struggled with small effective sample per bucket (vanishing grads)

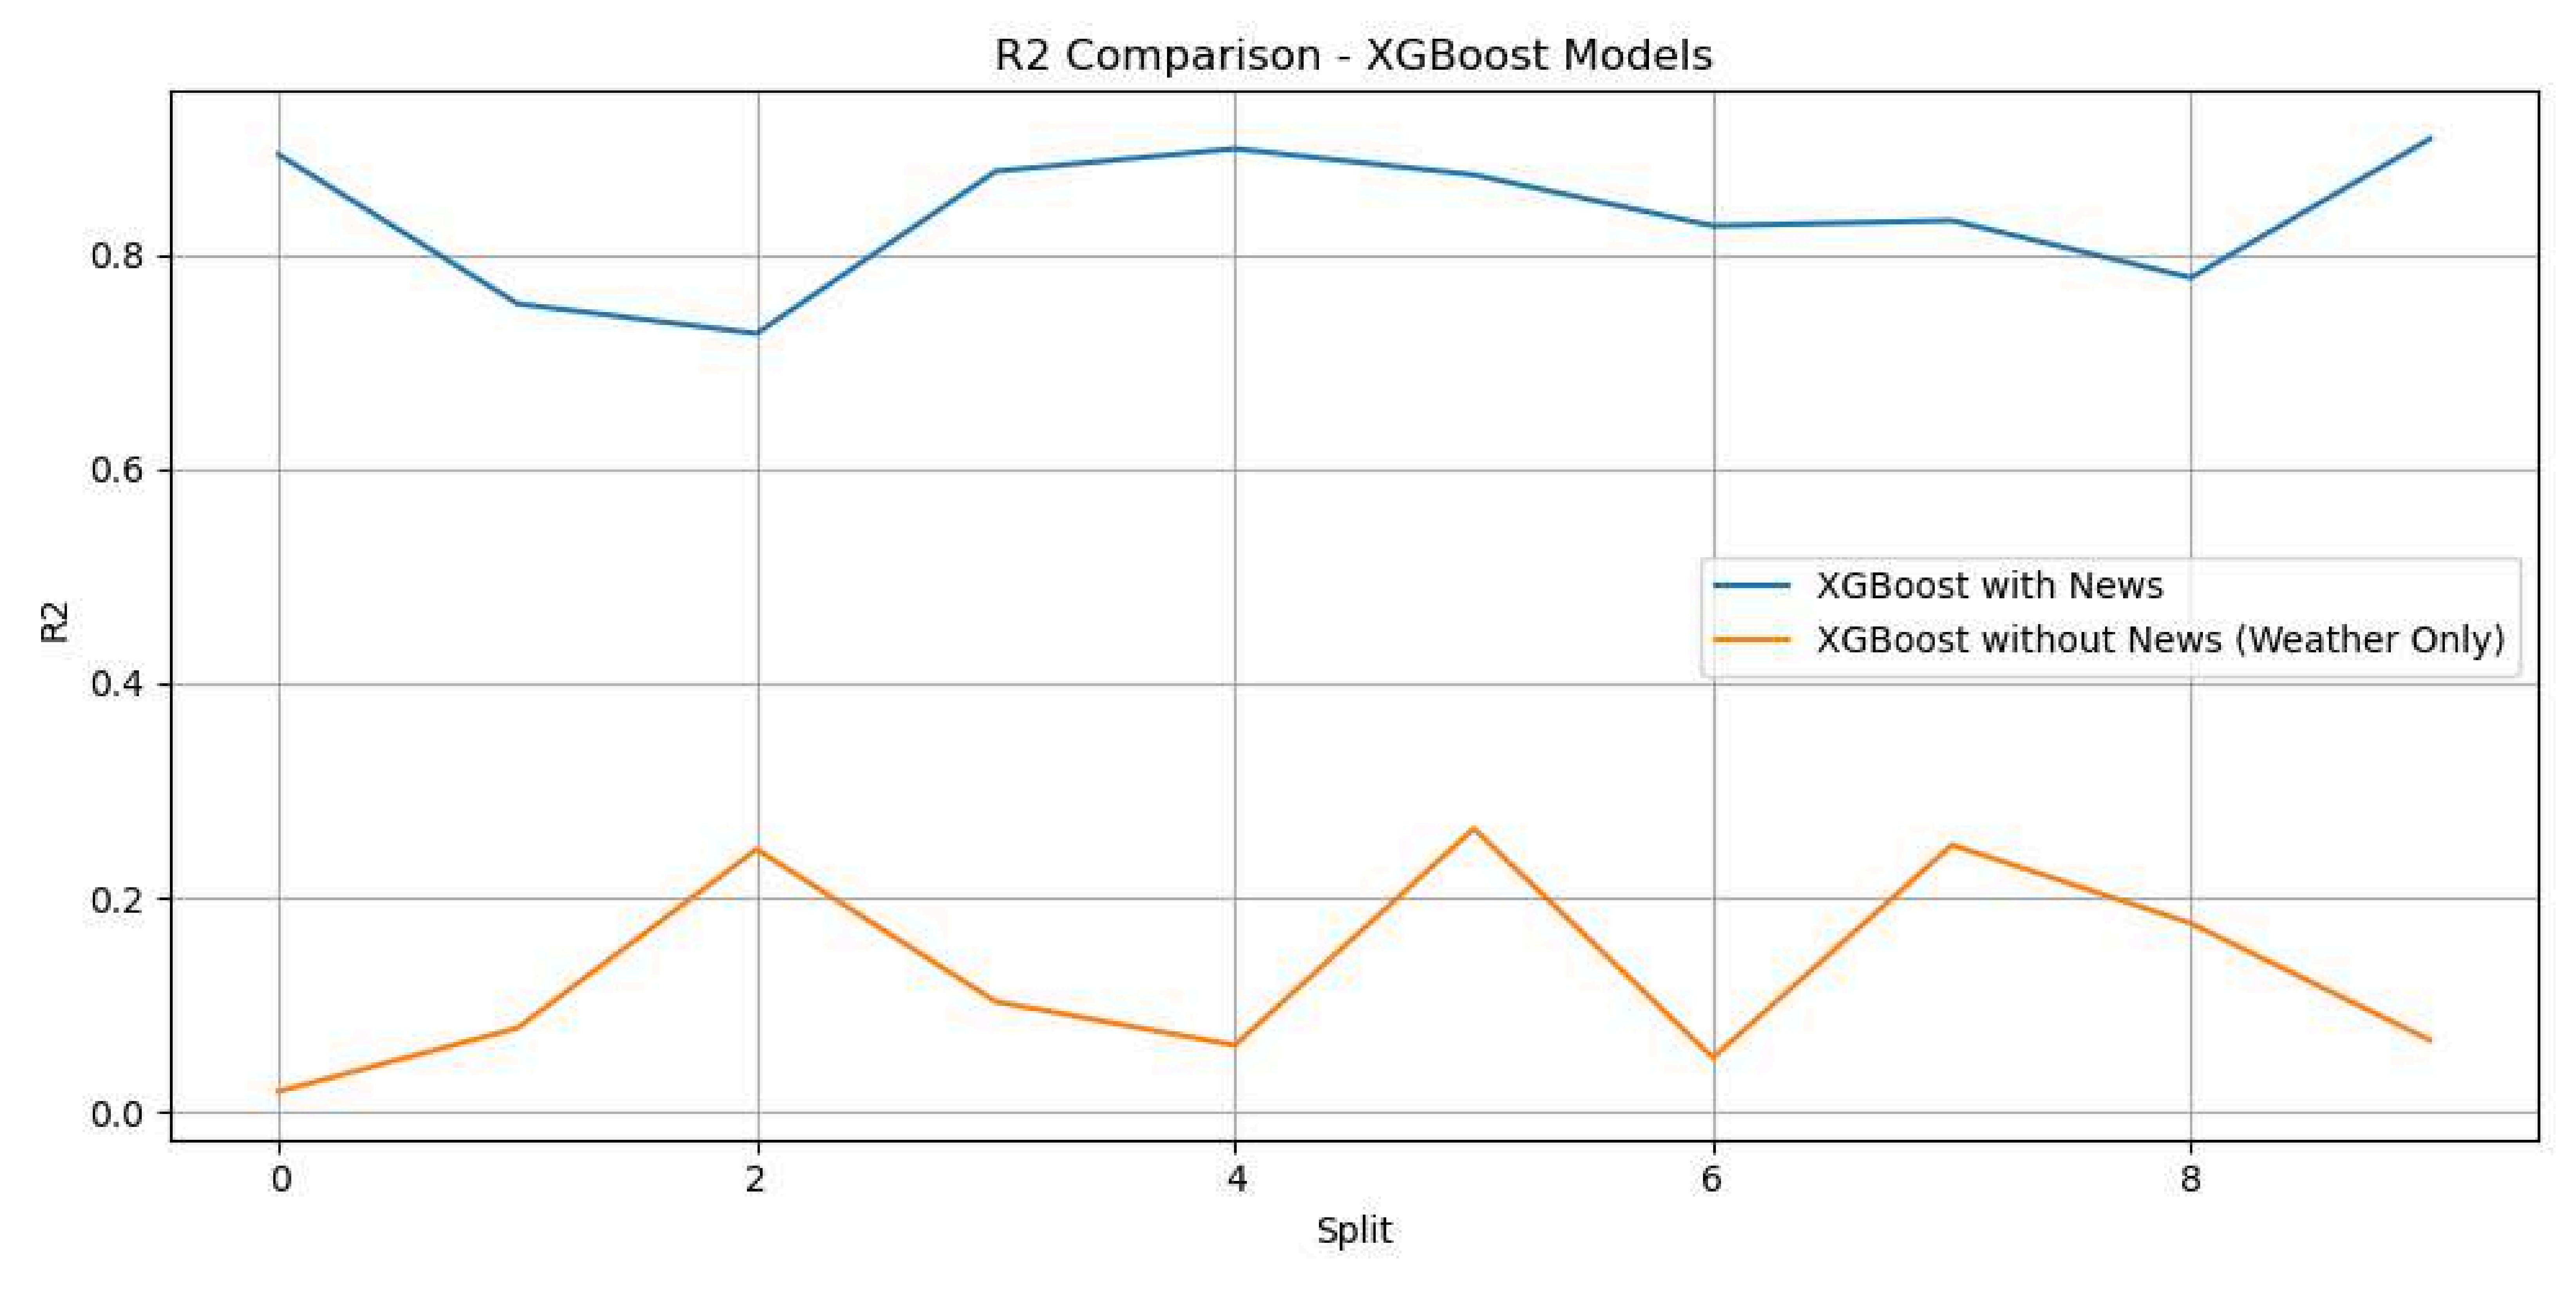
R2 Comparison - XGBoost Models
R2
0.0
0.2
0.4
0.6
0.8
0
2
4
6
8
Split
XGBoost with News
XGBoost without News (Weather Only)

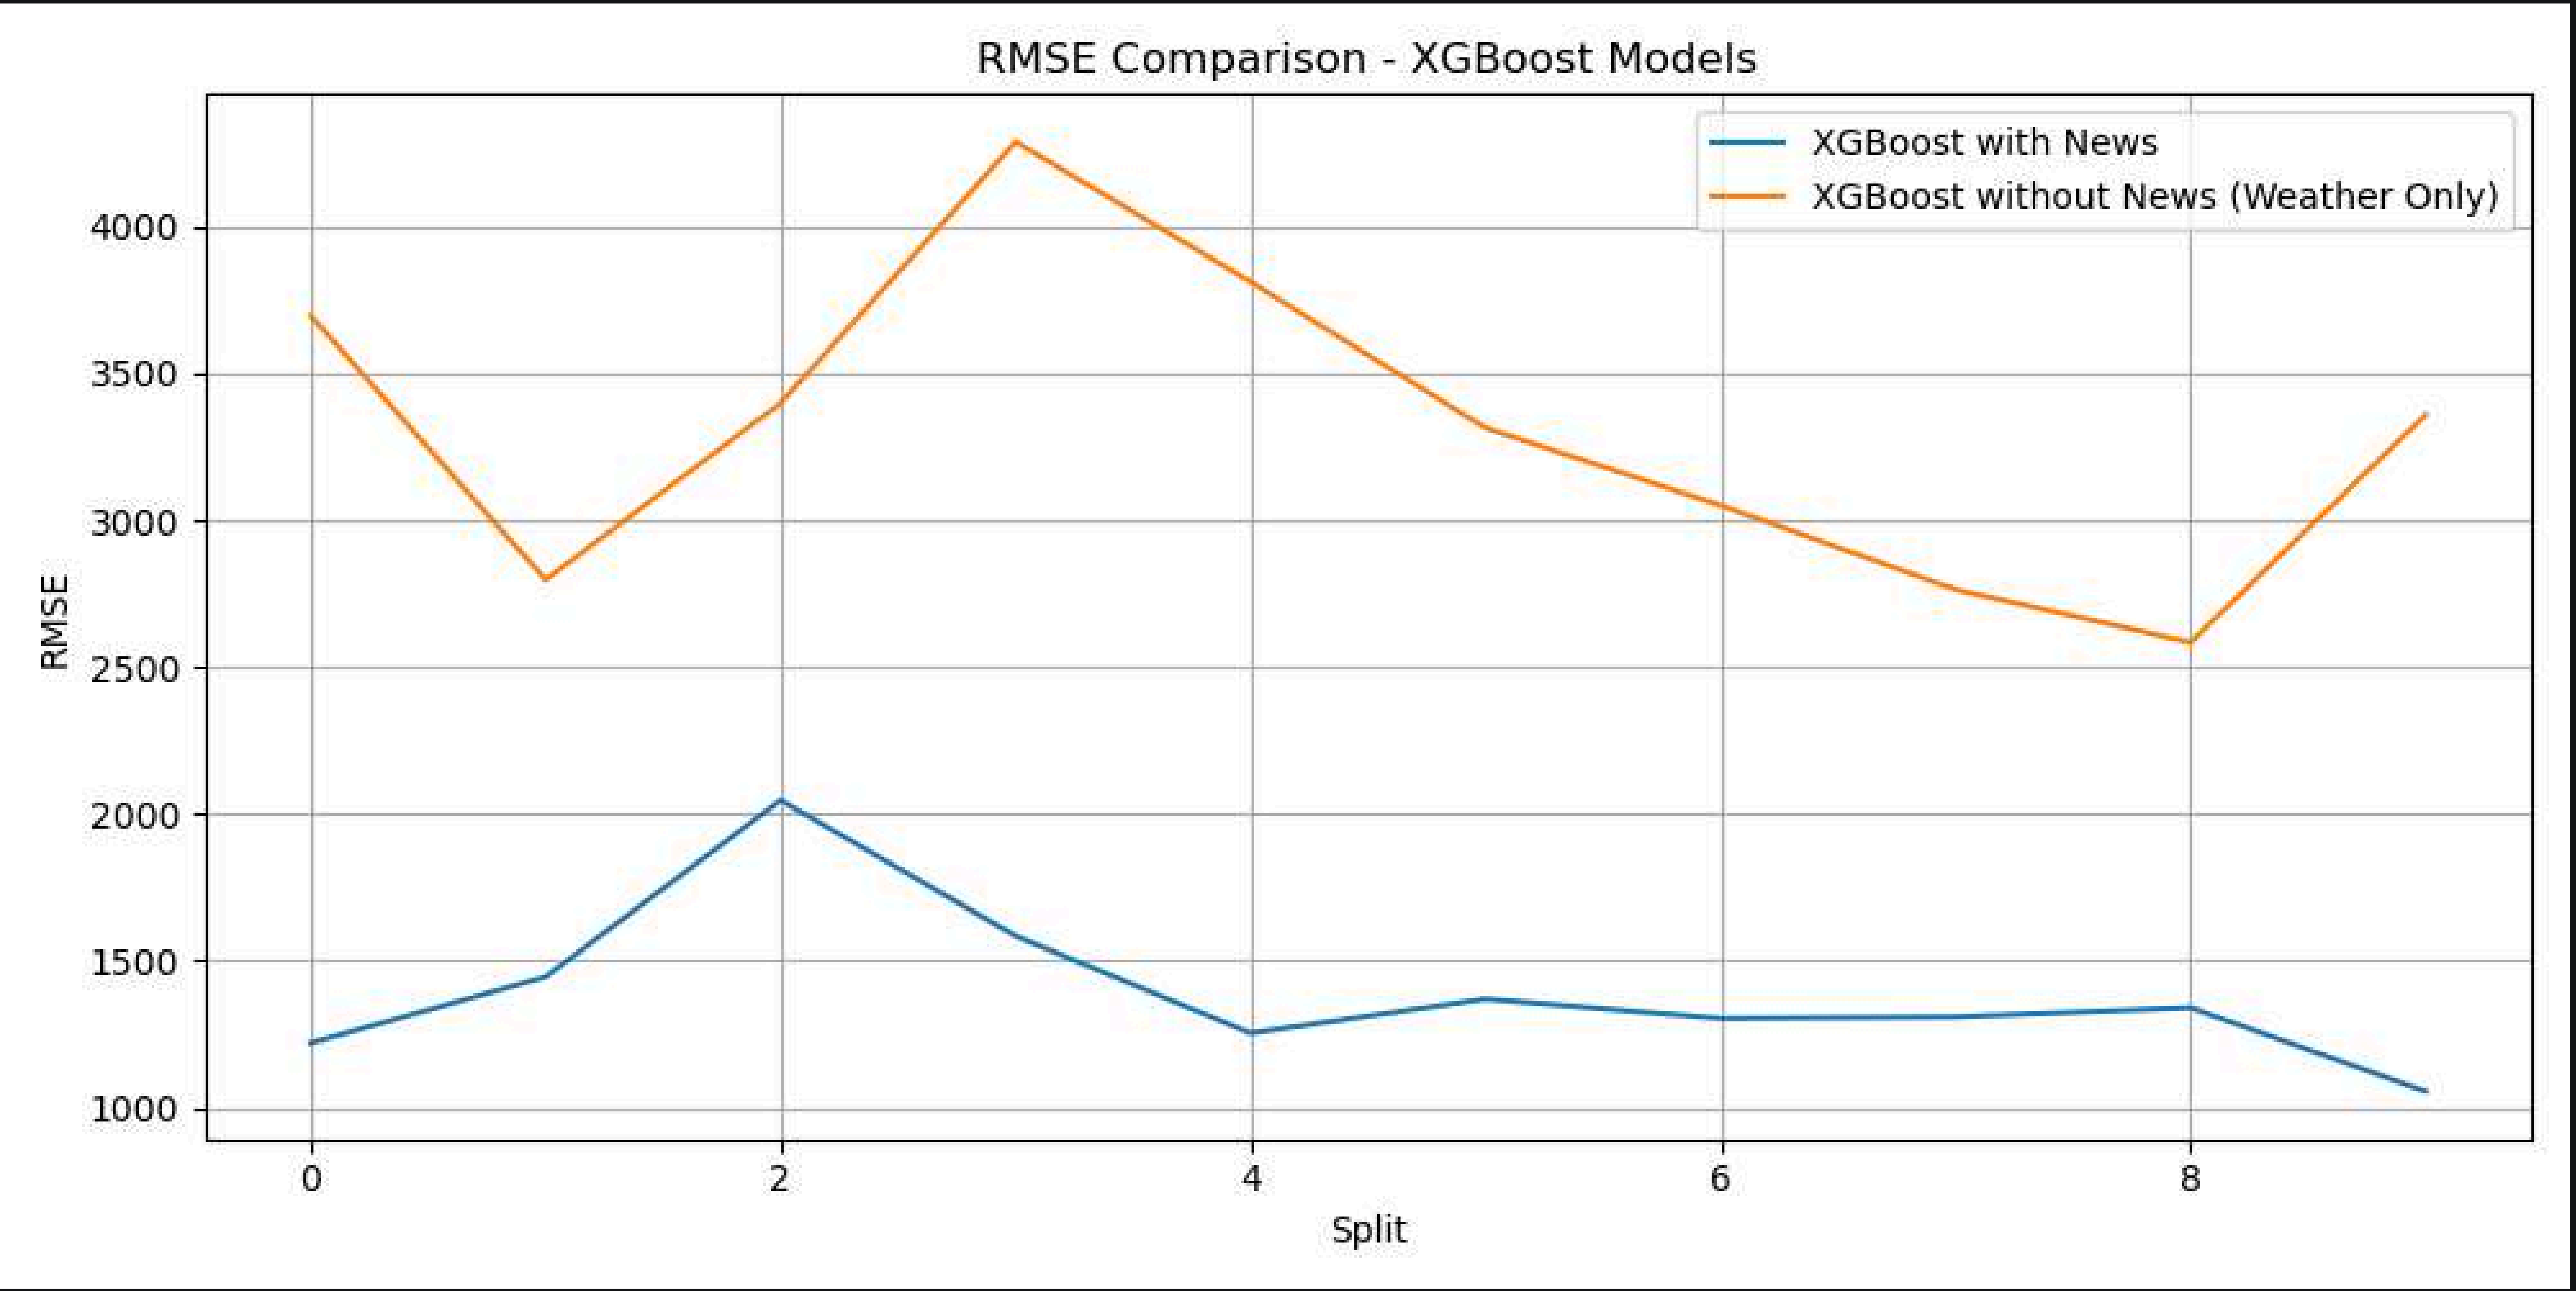
RMSE Comparison - XGBoost Models
XGBoost with News
XGBoost without News (Weather Only)
RMSE
4000
3500
3000
2500
2000
1500
1000
0
2
4
6
8
Split

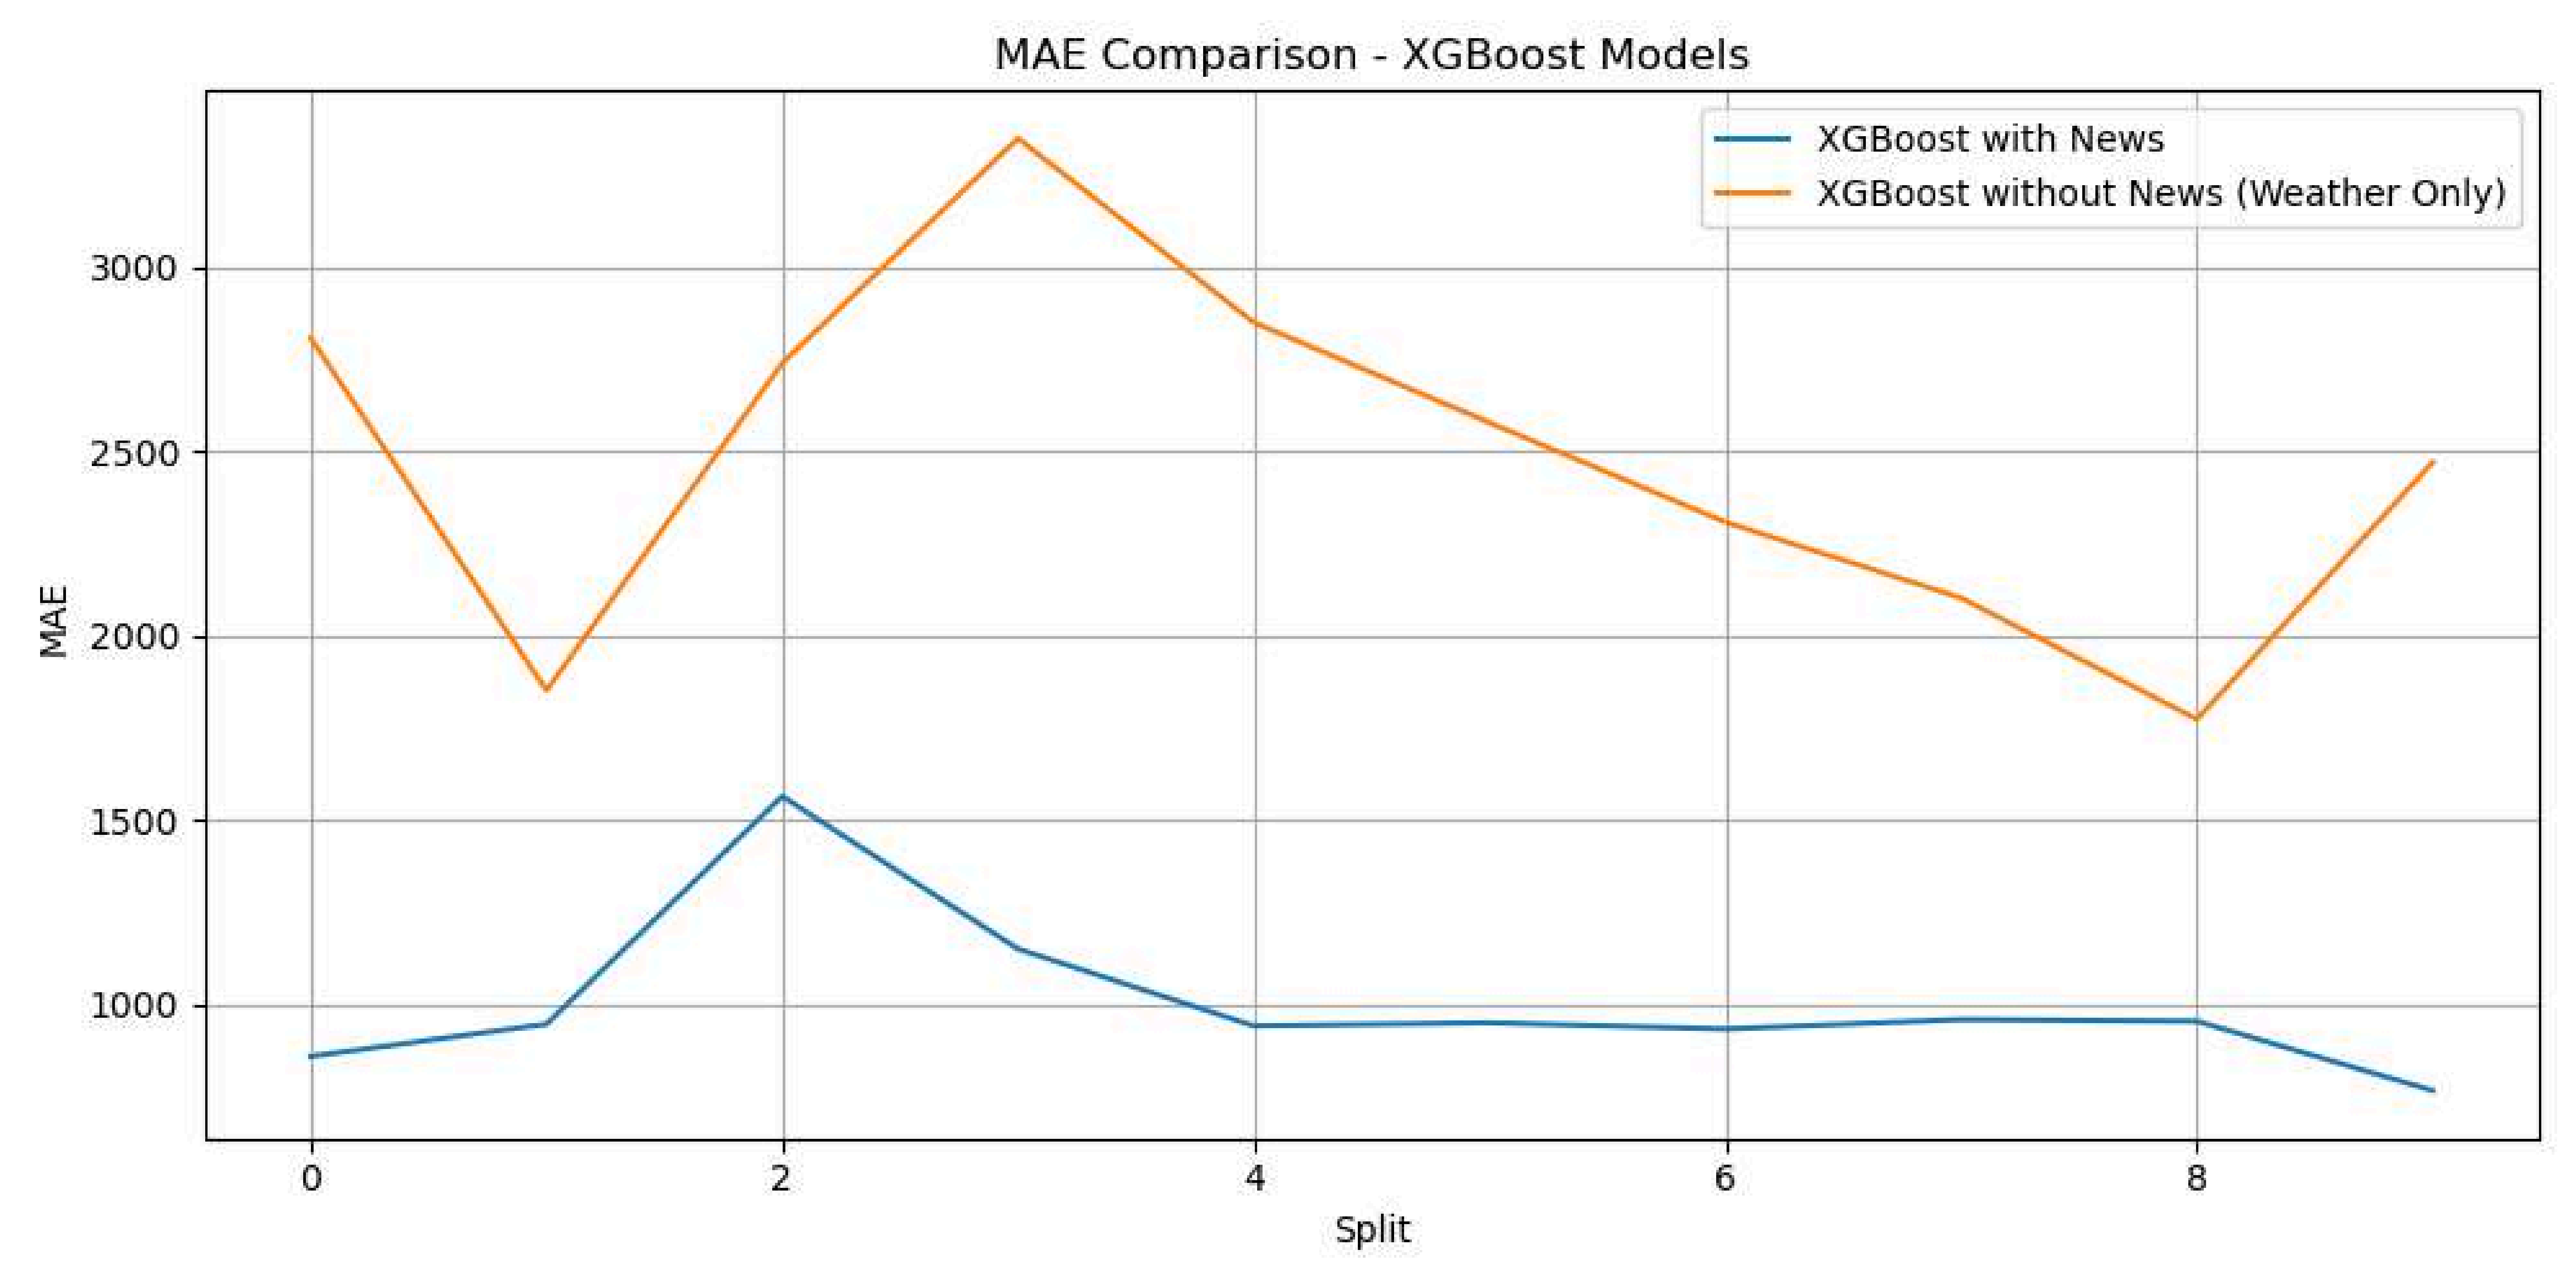
MAE Comparison - XGBoost Models
XGBoost with News
XGBoost without News (Weather Only)
MAE
3000
2500
2000
1500
1000
0
2
4
6
8
Split

Deployability

1. Computational Time
2. Less Data Needed
3. Computational Power

Limitations

1. More data would build a better LSTM Model
2. XGBoost expects training and future data to have similar patterns.
3. XGBoost is not inherently sequential and doesn't account for time order.

Our Team

Soumya Pandey

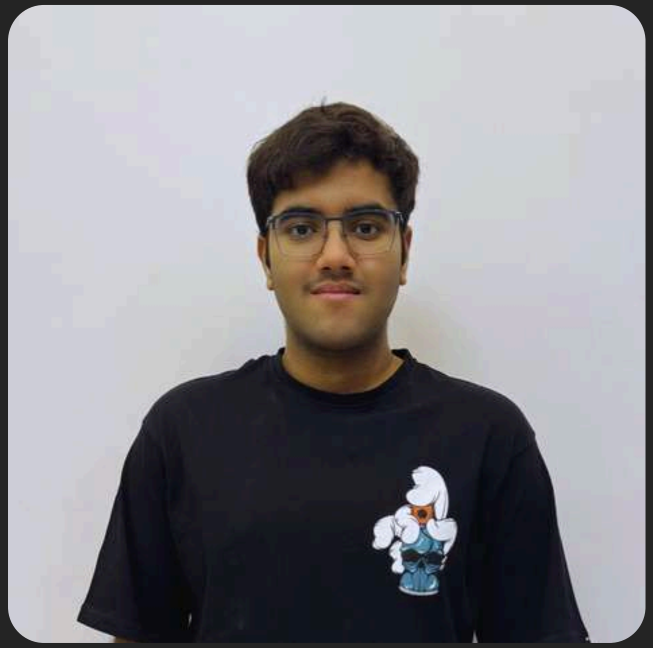

Nikunj Agrawal

Arsh Arora

Thank You

Questions?