From noisy measurements to an explainable forecast

Ravi Okafor, Elena Silva, and Noor Patel

Open Methods Group hello@example.org

Motivation

Short-horizon forecasts are often accurate enough to act on, but difficult to audit. We ask whether a sparse, monotonic feature summary can expose the reason for each forecast without sacrificing calibration.

Data

18,240 hourly observations from three monitored facilities. We reserve the final six weeks at each site for an untouched test set.

Contributions

1. A stable feature summary for irregular sensor streams.
2. A calibrated forecast with a human-readable rationale.
3. An evaluation split that tests transfer between facilities.

Method

Three-stage pipeline

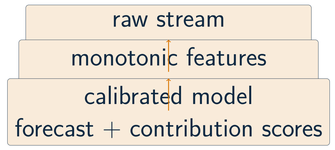

Results

Method	MAE	90% coverage	Transfer MAE
Seasonal baseline	1.84	0.82	2.41
Black-box ensemble	1.12	0.89	1.76
Our approach	**1.09**	**0.91**	**1.48**

Interpretation

The forecast says why

The strongest positive contribution is the recent slope; the strongest negative contribution is the recovery after the previous peak. The model declines to issue a confident forecast when both signals disagree.

Limitations

- Three facilities are not a universal deployment test.
- Contribution scores describe the model, not causality.
- Extreme missingness remains an abstention case.

Take-home

A small, stable feature vocabulary can make a useful forecast easier to challenge, transfer, and improve.

Contact example.org/forecast

Acknowledgements Open Data Methods Fund