WHITE PAPER / FORECASTING

What does yesterday tell us about tomorrow?

A planner needs a number before the day begins. We tested three forecasts that a dispatcher could reconstruct from one column of a rental ledger.

Which past days belong in the forecast?

Monday often looks different from Sunday. A natural first estimate for next Monday is therefore last Monday's count. Another planner might use yesterday's count because it captures the latest change in demand. Both rules make sense. An operations team still needs to know which one deserves a place on its planning sheet.

Research on non-stationary time series asks how the relevance of old observations changes over time. We took that question into a daily planning exercise. We compared three fixed rules on publicly recorded Capital Bikeshare rentals. Each rule had access to the same history at the same cut-off.

One forecast per morning

We used the 731 daily totals from 2011 and 2012. For each of the 366 dates in 2012 we predicted rentals using yesterday's count, the count seven days earlier, or the mean of the preceding seven daily counts. The December 2011 records supply the history for the first week in January. Every input predates the day being predicted.

We measured absolute error in rentals per day. The mean describes a typical planning burden across the year. The 90th percentile helps a planner see how far off the estimate can be on a difficult day. We kept both measures in the original units.

The weekly average earned the first benchmark

Daily rental forecast error across 2012, rounded to whole rentals

Forecast rule	Mean absolute error	90th percentile
Yesterday	870	1,956
Same weekday last week	1,110	2,485
Preceding seven-day mean	848	1,731

The seven-day mean had the smallest error on both measures. Compared with yesterday alone, the difference was about 22 rentals in mean error and 225 rentals at the 90th percentile. The same-weekday rule performed worst over this year despite its plausible calendar logic.

What the planner should ask next

A network-wide daily total can help with broad capacity planning. A depot needs finer information, including where bikes start the morning and which stations empty first. The next benchmark should match the geography and time window of the decision it supports.

For a weather-aware model, replay the forecasts that dispatchers actually had before each shift. Using the weather observed later in the day would answer a different question. Keep the seven-day average in the daily report so a more complicated model has a visible comparator through each change of season.

Research & data

- Learning Theory and Algorithms for Forecasting Non-stationary Time Series · NeurIPS 2015
- Capital Bikeshare daily rentals · UCI Machine Learning Repository