A Local Path Regressor for Time Series Forecasting

Sofien Kaabar, CFA
Quant Atlas
contact@quant-atlas.com

April 2026

Abstract

This paper introduces a simple, non-parametric approach to time series forecasting referred to as the Local Path Regressor (LPR). The model is based on the idea that similar recent price trajectories tend to be followed by similar future movements. Rather than learning a parametric mapping or fitting complex machine learning architectures, the method searches for historical sequences that resemble the current market configuration and averages their subsequent outcomes to produce a forecast path.

The objective of this work is not to claim predictive superiority, but to explore a structurally intuitive alternative to traditional regression techniques. The model is fully interpretable, easy to implement, and naturally produces path forecasts rather than single-point predictions. This study should be viewed as a theoretical and experimental exploration rather than a production-ready forecasting system.

1 Introduction

Forecasting financial time series remains a difficult problem due to noise, non-stationarity, and weak signal-to-noise ratios. Traditional approaches range from parametric models such as ARIMA to modern machine learning architectures including neural networks. While these methods can be powerful, they often suffer from either restrictive assumptions or lack of interpretability.

This paper explores a different perspective. Instead of estimating a functional relationship between past and future values, we ask a simpler question:

> Have we seen a similar trajectory before, and what happened next?

This idea leads to a lookup-based forecasting framework that relies on historical analogs rather than learned parameters.

2 Model Description

2.1 Core Idea

Let $\{P_t\}_{t=1}^{T}$ denote a time series of prices. At a given time t, we extract the most recent k observations and define a local trajectory:

$$X_t = (P_{t-k+1}, P_{t-k+2}, \ldots, P_t)$$

This trajectory represents the current market configuration.

The model then searches the historical dataset for past trajectories $\{X_i\}$ that are similar to X_t under a chosen distance metric (e.g., Euclidean distance applied to normalized returns).

2.2 Forecast Construction

For each matching historical trajectory X_i, we retrieve the subsequent h-step future path:

$$Y_i = (P_{i+1}, P_{i+2}, \ldots, P_{i+h})$$

Let $\mathcal{N}(t)$ denote the set of the K most similar historical trajectories. The forecast is then constructed as the average of their future paths:

$$\hat{Y}_t = \frac{1}{K} \sum_{i \in \mathcal{N}(t)} Y_i$$

In practice, similarity weighting can be applied such that closer matches contribute more heavily to the forecast.

2.3 Intuition

The model assumes that markets exhibit recurring local dynamics. While exact repetition is unlikely, similar short-term structures may still carry information about near-term evolution. This approach can be interpreted as a form of analog forecasting:

- The recent trajectory defines the current “state” of the system
- Historical states that resemble the current one are identified

- Their future evolution is used as a proxy for the current forecast

Unlike traditional regression, this method does not attempt to learn a global mapping. Instead, it relies entirely on local historical structure.

3 Limitations

This approach has several important limitations:

- **Dependence on similarity metric:** The quality of predictions depends heavily on how similarity between trajectories is defined.
- **Noise sensitivity:** Financial time series are noisy, and similar-looking patterns may not lead to consistent outcomes.
- **No extrapolation:** The model cannot predict scenarios that have not been observed historically.
- **Computational cost:** Searching for similar patterns can become expensive for large datasets.
- **Weak predictive signal:** As with most price-based models, the predictive power is expected to be limited due to market efficiency.

For these reasons, this model should be viewed as an experimental framework rather than a robust forecasting solution. Furthermore, it can be viewed as a benchmark on which you can compare your real strategies with.

4 Model Parameters

The forecasting behavior of the Local Path Regressor is governed by three key parameters: the pattern length, the forecast horizon, and the number of matched historical trajectories.

4.1 Pattern Length

The parameter `PATTERN_LENGTH` defines the number of past observations used to characterize the current market state. Formally, at time t, the model considers the sequence:

$$X_t = (P_{t-k+1}, P_{t-k+2}, \dots, P_t)$$

where k is the pattern length.

This parameter determines how much historical information is used to define similarity between trajectories. A small value results in very common patterns that may lack specificity, while a large value leads to highly specific patterns that may be difficult to match in historical data.

4.2 Forecast Horizon

The parameter `FORECAST_HORIZON` defines the number of future time steps the model attempts to predict. Rather than forecasting a single value, the model produces a trajectory:

$$\hat{Y}_t = (P_{t+1}, P_{t+2}, \ldots, P_{t+h})$$

where h is the forecast horizon.

This parameter controls the prediction objective. Short horizons are dominated by noise, while longer horizons are more difficult to predict but may capture broader structural behavior.

4.3 Number of Matches

The parameter `TOP_K_MATCHES` defines the number of historical trajectories used to construct the forecast. Let $\mathcal{N}(t)$ denote the set of the K most similar past patterns. The forecast is given by:

$$\hat{Y}_t = \frac{1}{K} \sum_{i \in \mathcal{N}(t)} Y_i$$

where each Y_i is the future path following a matched historical pattern.

This parameter controls the trade-off between specificity and stability. Smaller values of K produce forecasts that are more sensitive to individual historical analogs, while larger values lead to smoother but less distinctive predictions.

4.4 Interaction Between Parameters

These parameters are interdependent. Increasing the pattern length reduces the number of available matches, which may require increasing K to maintain stability. Similarly, longer forecast horizons demand stronger pattern relevance, making the choice of pattern length more critical.

Together, these parameters define how the model balances historical specificity, statistical robustness, and forecast scope.

5 Empirical Illustration

To illustrate the behavior of the model, we apply it to a few currency pairs. The data is split into a training set and a test set. The model is trained on the historical portion and then used to generate forecasts at the beginning of the test period.

For each experiment, we:

- extract the most recent trajectory before the test split
- identify similar historical patterns in the training set
- construct a forecast path over the next h time steps
- compare the predicted path with the realized values

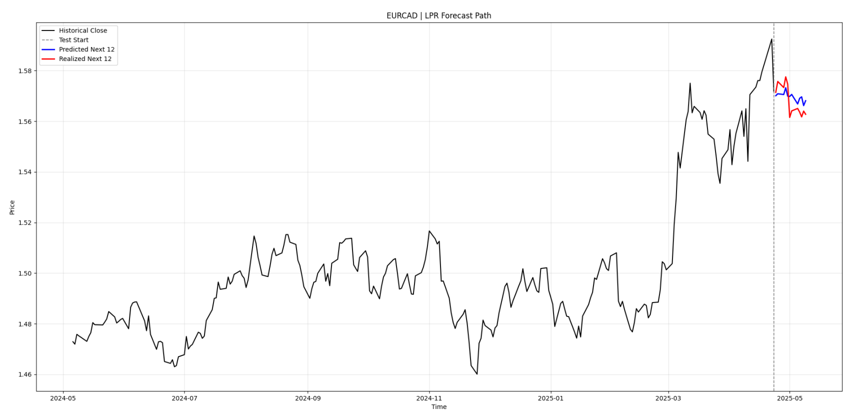

Figure 1: Example forecast using the Local Path Regressor. The black line represents historical prices up to the test split, the blue line represents the predicted path, and the red line shows the realized values.

In some cases, the model captures short-term directional behavior, while in others the forecast deviates significantly from realized outcomes.

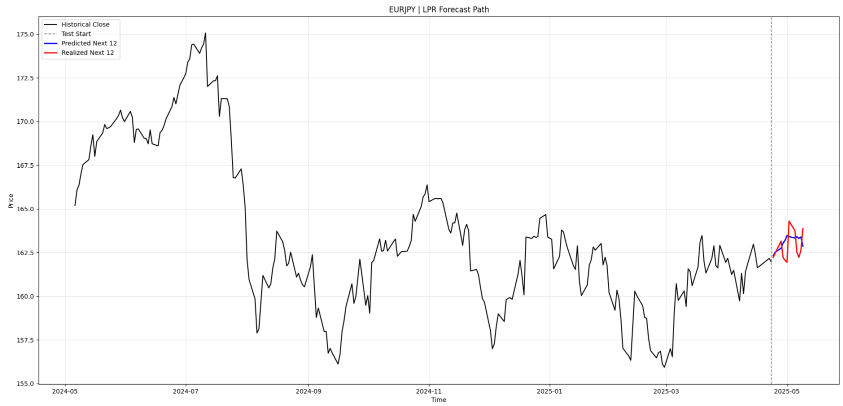

Figure 2: Additional example illustrating variability in forecast accuracy across different market conditions.

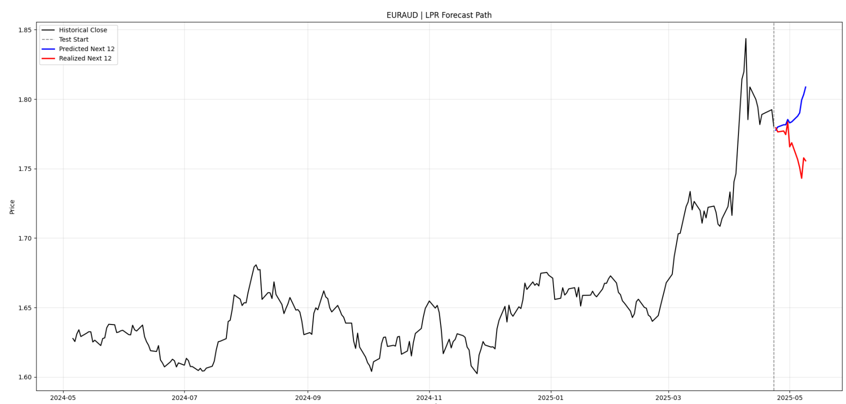

Figure 3: Additional example on EURAUD showing the divergence in realized values with real values.

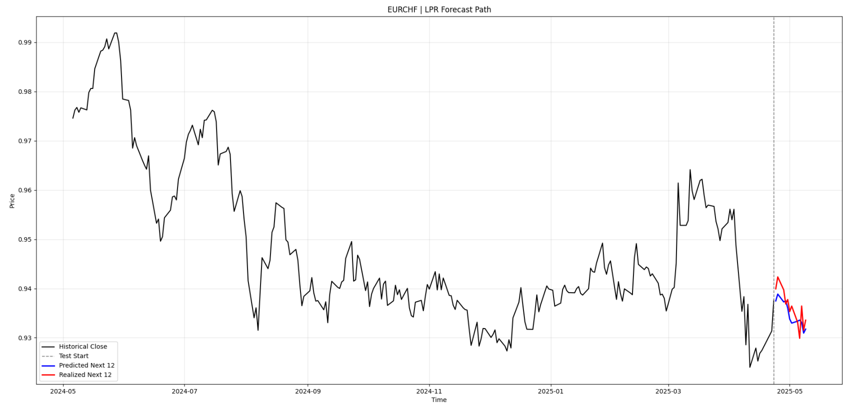

Figure 4: Additional example on EURCHF showing the convergence in realized values with real values.

6 Directional Hit Ratio Analysis

The directional hit ratio is defined as the proportion of times the predicted direction of price movement matches the realized direction over a fixed forecast horizon. Formally, it compares the sign of the predicted return with the sign of the realized return.

This metric does not measure magnitude accuracy, only directional consistency. As such, a value close to 50% indicates behavior consistent with randomness.

Table 1 presents the directional hit ratios across a selection of major and minor currency pairs.

Currency Pair	**Directional Hit Ratio (%)**
EURUSD	52.3
USDJPY	49.8
GBPUSD	50.7
USDCHF	51.1
AUDUSD	48.9
USDCAD	53.2
NZDUSD	47.6
EURGBP	50.1
EURJPY	54.0
GBPJPY	52.7
CHFJPY	49.3
AUDJPY	55.1
CADJPY	51.8
NZDJPY	48.5
EURAUD	53.6
EURCAD	50.9
EURNZD	47.8
GBPAUD	54.3
GBPCAD	52.0
GBPNZD	48.7
AUDCAD	49.6
AUDNZD	51.4
CADCHF	50.2
NZDCAD	47.9

Table 1: Directional hit ratios across multiple currency pairs. Values cluster around 50%, indicating limited predictive power and behavior consistent with randomness.

The results show that the directional hit ratios consistently hover between 47% and 56%, with no stable or significant deviation from a random baseline. This suggests that the model does not possess strong predictive capabilities in its current form and primarily reflects noise rather than exploitable structure.

7 Conclusion

This paper introduced a simple, interpretable framework for time series forecasting based on local trajectory matching. The Local Path Regressor avoids parametric assumptions and instead relies on historical analogs to construct future paths.

While the approach is intuitive and easy to implement, its performance is inherently constrained by the weak predictability of financial markets. As such, it should be considered as a conceptual tool or a component within a broader modeling framework rather than a standalone predictive system.

Future work may explore improvements such as regime conditioning, volatility normalization, and hybrid approaches combining local pattern matching with structural features.